

NBER WORKING PAPER SERIES

MOBILITY AND CONGESTION IN URBAN INDIA

Prottoy A. Akbar
Victor Couture
Gilles Duranton
Adam Storeygard

Working Paper 25218
<http://www.nber.org/papers/w25218>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2018

This work is supported by the World Bank, the Zell Lurie Center for Real Estate at the Wharton School, the Fisher Center for Urban and Real Estate Economics at Berkeley Haas, and we also gratefully acknowledge the support of the Global Research Program on Spatial Development of Cities at Ise and Oxford funded by the Multi Donor Trust Fund on Sustainable Urbanization of the World Bank and supported by the UK Department for International Development. We appreciate the comments from Leah Brooks, Ben Faber, Michael Gechter, Ejaz Ghani, Ed Glaeser, Vernon Henderson, KiJoon Kim, Gabriel Kreindler, Emile Quinet, Christopher Severen, Kate Vyborny, and participants at conferences and seminars. Hero Ashman, Xinzhu Chen, Allison Green, Xinyu Ma, Gao Xian Peh, and Jungsoo Yoo provided us with excellent research assistance. We are immensely grateful to Sam Asher, Geoff Boeing, Arti Grover, Nina Harari, and Yue Li for their help with the data. The views expressed here are those of the authors and not of any institution they may be associated with, nor of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Prottoy A. Akbar, Victor Couture, Gilles Duranton, and Adam Storeygard. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Mobility and Congestion in Urban India

Prottoy A. Akbar, Victor Couture, Gilles Duranton, and Adam Storeygard

NBER Working Paper No. 25218

November 2018

JEL No. R41

ABSTRACT

We develop a methodology to estimate robust city level vehicular mobility indices, and apply it to 154 Indian cities using 22 million counterfactual trips measured by a web mapping service. There is wide variation in mobility across cities. An exact decomposition shows this variation is driven more by differences in uncongested mobility than congestion. Under plausible assumptions, a one standard deviation improvement in uncongested speed creates much more mobility than optimal congestion pricing. Denser and more populated cities are slower, only in part because of congestion. Urban economic development is correlated with better (uncongested and overall) mobility despite worse congestion.

Prottoy A. Akbar
University of Pittsburgh
Dietrich School of Arts and Sciences
Department of Economics
4901 Wesley W. Posvar Hall
230 South Bouquet Street
Pittsburgh, PA 15260
prottoyamanakbar@pitt.edu

Victor Couture
University of California, Berkeley
Haas School of Business
2220 Piedmont Avenue
Berkeley, CA 94720-1900
couture@berkeley.edu

Gilles Duranton
The Wharton School
University of Pennsylvania
3620 Locust Walk
Philadelphia, PA 19104
and NBER
duranton@wharton.upenn.edu

Adam Storeygard
Department of Economics
Tufts University
Braker Hall
8 Upper Campus Road
Medford, MA 02155
adam.storeygard@tufts.edu

1. Introduction

Using a popular web mapping and transportation service, we generate information for more than 22 million counterfactual trip instances in 154 large Indian cities.¹ We then use this information to estimate indices of mobility (speed) of motorized vehicle travel in these cities. We first assess the robustness of our indices to a wide variety of methodological choices. Second, we provide a novel decomposition of overall mobility into uncongested mobility and the congestion delays caused by traffic. This decomposition allows us to perform a simple welfare analysis that compares the gains from improving uncongested mobility with the gains from introducing optimal congestion pricing. Third, we examine how indicators of urban economic development and other city characteristics correlate with mobility, uncongested mobility, and congestion delays. Finally, we provide additional mobility indices for walking and transit trips.


To the best of our knowledge, our paper provides the first systematic investigation of urban mobility across cities in a developing country.² Our main substantive findings are the following. First, there are large differences in mobility across Indian cities. A factor of nearly two separates the fastest and slowest cities. To illustrate this, figure 1 plots the speed of travel by motorized vehicles throughout the day in a particularly fast Indian city, Chandigarh, and in a particularly slow city, Kolkota.

Second, variation in mobility across cities is driven primarily by uncongested mobility, not by congestion delays. In figure 2, there is on average little variation across hours of the day for our sample of 154 cities. Due to slow uncongested mobility, a very poor city like Varanasi (Benares) is slower than average at all times, even at night in the absence of traffic. We see much greater intra-day differences in mobility in large cities, particularly the largest ones close to their center, as illustrated by Delhi in the figure. More generally, an index of uncongested mobility explains more than 50% of the variance in overall mobility across cities. A simple welfare analysis also suggests much larger gains from a 10% improvement in uncongested mobility than from implementing optimal congestion pricing in all of urban

¹By counterfactual, we mean trip instances that have not been actually taken by a household. As we show below, these trips were selected to mimic some characteristics of trips that are taken by households in other contexts.

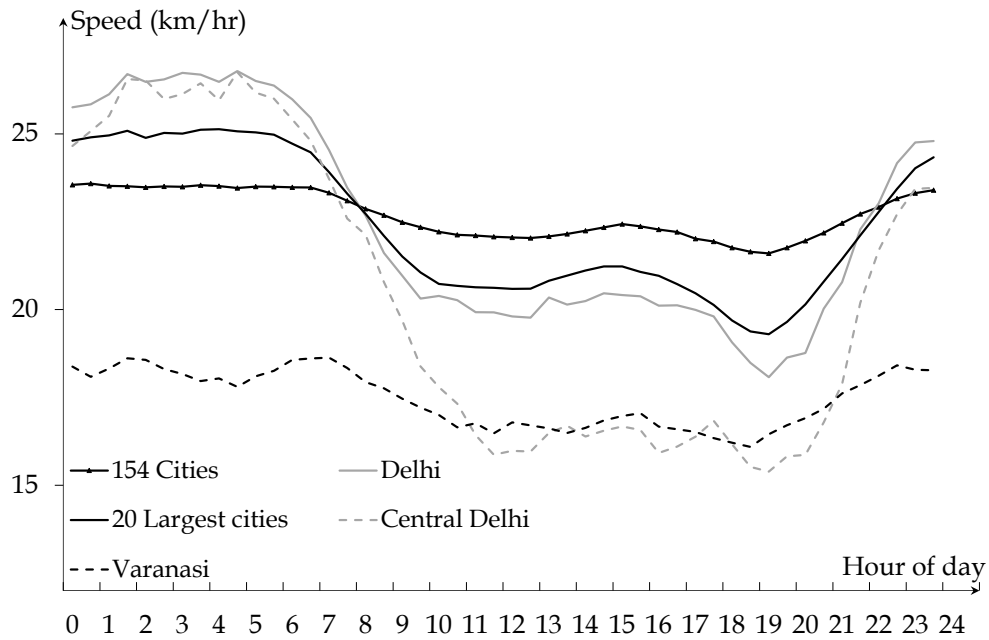
²Two new studies focusing on a single developing city complement our cross-city investigation: Kreindler (2018) studies the welfare impact of congestion pricing in Bangalore, and Akbar and Duranton (2018) measure the cost of congestion in Bogotá.

Figure 1: Mobility throughout the day, two select Indian vs. two us cities



Mean speed for trips with length between 5 and 10 kilometers. We consider trips in the same length range for reasons made clear below. Central Kolkota and Chicago refer to trips that take place on average within 5 kilometers of the center of their city. See section 2 for further details.

Figure 2: Mobility throughout the day in India



Mean speed for trips with length between 5 and 10 kilometers. Central Delhi refers to trips that take place on average within 5 kilometers of the center of Delhi. See section 2 for further details.

India, even after accounting for the potential of congestion pricing to reduce the significant level of travel time unreliability that we document. These findings challenge the conventional wisdom that traffic congestion is the main reason why some cities are slow and some are fast. To take one prominent example, a recent report by the Boston Consulting Group (BCG, 2018) claims that Kolkata has the most traffic congestion among the four largest Indian cities. We find that Kolkata is in fact the least congested of the four, but the slowest because of low uncongested mobility. This distinction has important policy implications, because uncongested speed cannot be improved by congestion pricing, ride-sharing promotion or restriction, or other policies often proposed to combat congestion.

Third, travel is generally slow in Indian cities, even outside peak hours. In addition to Chandigarh and Kolkata, figure 1 also plots comparable speed data for a fast us city, Grand Rapids, and a slow one, Chicago.³ Even the central part of Chicago, one of the most congested locations in the us, is generally faster than one of the fastest Indian cities, Chandigarh. Finally, we find that denser, more populated cities are slower, that there is a hill-shaped relationship between city per capita income and mobility, and that a city's mobility is related to characteristics of its road network.

This investigation is important for three reasons. First, there is an extreme paucity of useful knowledge about urban transportation, especially in developing countries. As a first building block towards a more serious knowledge base on urban transportation, some stylized facts are needed.⁴ For instance, we need to know how slow travel is in developing cities beyond the anecdotal evidence offered by disgruntled travelers. Equally important objects of interest are the differences between cities, between different parts of the same city, and across times of day within the same city.⁵ We hope that our results, methodology, and data sources can

³Couture, Duranton, and Turner (2018) rank Grand Rapids and Chicago as the second slowest and second fastest cities among the 50 largest metropolitan areas in the us. In the ranking below, Chandigarh is among the fastest cities in India and a natural counterpart of Grand Rapids in terms of population. Kolkata, the third largest city in India, is the natural counterpart to Chicago. It is also the slowest city in India.

⁴In richer countries, much of our knowledge stems from representative surveys of household travel behavior. These surveys nonetheless have clear limitations, including a lack of precision in what travelers report. They are also prohibitively expensive to carry out broadly in developing countries. For the us, the Bureau of Transportation Statistics reports a cost per household of perhaps \$300 to produce the National Household Transportation Survey or about \$40 million in total (see <http://onlinepubs.trb.org/onlinepubs/reports/nhts.pdf>, last accessed, 6 September 2018.)

⁵Several software and data services such as Inrix and TomTom propose popular measures of congestion for a large sample of world cities. These services do not make the details of their methodology public. It seems that they monitor either specific roads or average traffic speed. We show below that measures of average speed are problematic and perform poorly. Uber Mobility also provides data, including travel times for four Indian cities that, we argue below, lead to substantial overestimates of mobility.

help guide policy and future research on urban transportation in developing countries. We devote much of the last section of our paper to providing such guidance.

Second, there is a popular view that urbanization and economic development lead to ever larger cities and increased rates of motorization. According to this view, these two features will eventually lead to complete gridlock. We do find evidence of congestion in the largest Indian cities and a strong association between congestion and household access to motorized vehicles. However, economic development also brings about better travel infrastructure which facilitates uncongested mobility. In fact, indicators of urban economic development such as faster recent population growth, higher income levels (except at the very top), and higher motorization rates are generally associated with better overall mobility despite worse congestion.

Third, urban transportation in developing countries is prioritized for massive investments. For instance, transportation is the largest sector of lending by the World Bank and represents more than 20% of its net commitments as of 2016. Among the many problems that these investments are trying to remedy, the lack of urban land devoted to the roadway is widely perceived to be a chief cause behind slow mobility and urban congestion. Providing an assessment of the determinants of mobility to guide policy is thus fundamental. For instance, we find suggestive evidence that better mobility is associated with a more regular grid network and more primary roads.

Our investigation raises three challenges. The first is methodological. We propose a new approach to measure various forms of mobility from trip information, and to decompose them into uncongested mobility and delays caused by congestion. The second is a travel data challenge. There is no comprehensive source of data about urban transportation in Indian cities. Our approach is to collect data on predicted travel time from a popular website, Google Maps (GM).⁶ For each city, we designed a sample of trips and sampled each trip at different times on different days. Our main worry is that these counterfactual trips may not be representative of the actual travel conditions faced by city residents. To address this worry, we use four different trip design strategies. These strategies aim to replicate some characteristics of actual trips taken by urban households in other countries. We show that our

⁶https://en.wikipedia.org/wiki/Google_Maps, last accessed, 6 September 2018. A number of new studies, which we discuss later in the paper, also use Google Maps to measure traffic in a developing city, notably Kreindler (2016), Hanna, Kreindler, and Olken (2017), and Akbar and Duranton (2018). Alternative approaches include direct GPS records for particular vehicles such as taxis (Mangrum and Molnar, 2017) or sensors, which usually track traffic only on the most important arteries (Geroliminis and Daganzo, 2008).

city mobility indices vary little across sampling strategies, type of trip destinations, origin and direction of travel, or time of day. Finally, we face the challenge of consistently defining and measuring the cities in which we sample counterfactual trips. To answer this challenge, we rely on a wide variety of sources including the census of India, OpenStreetMap, and satellite imagery.

2. Data collection

In this section we provide an overview of our data. Further details are available in Appendix A.

2.1 City sample

United Nations (2015) reports the names and locations of 166 cities in India that reached a population of 300,000 by 2014. Following Harari (2016) and Ch, Martin, and Vargas (2018), we initially define the spatial extent of these cities using nightlights. Within these light boundaries, we restrict attention to 40-meter pixels defined as built-up in 2014 according to the Global Human Settlements Layer (GHSL) of the European Commission’s Joint Research Centre (JRC). After dropping cities for which no appropriate light exists, aggregating multiple cities within the same contiguous light, and dropping cities for which the relevant GHSL data are missing, we are left with an estimation sample of 154 cities.

2.2 Trips data

We define a *trip* as an ordered pair of points (origin and destination) within the same city as defined above. A *trip instance* is a trip taken at a specific time. Our target sample for city c is $15 \sqrt{Pop_c}$ trips, where Pop_c is the projected 2015 population of city c from United Nations (2015), and 10 trip instances per trip, to ensure variation across times of day. For a city of population, say, one million, our sampling strategy thus targets 15,000 trips (7,500 in each direction between the trips’ endpoints) and 150,000 trip instances. All trips are restricted to be at least one kilometer between origin and destination because Google results are less reliable for very short trips, few of which we expect to be motorized anyway. We sample across times of day to roughly match the weekday distribution of actual trips in Bogotá from Akbar and Duranton (2018). We oversample sparse overnight periods, and sample weekends at half the rate of weekdays.

We sample across four broad classes of trips, each designed to reflect key aspects of urban travel: radial, circumferential, gravity, and amenity trips.

Radial trips join a randomly located point within 1.5 kilometers of a city's center (as defined by United Nations, 2015) with another point in the city, either approximately 2, 5, 10, or 15 kilometers away, or at a distance percentile drawn from a uniform distribution. These trips are those predicted by the standard monocentric model of cities (Alonso, 1964, Mills, 1967, Muth, 1969). The monocentric model provides a reasonable first-order characterization of the distribution of population, density, and land and house prices in cities of many countries (see Duranton and Puga, 2015, for a survey).

Circumferential trips, orthogonal to radial trips, join a randomly located origin at least 2 kilometers from the city center with a destination at approximately the same radius but displaced approximately 30 degrees clockwise or counterclockwise.

Gravity trips join a random origin with a destination in a random direction, at a distance that is drawn from a truncated Pareto distribution with shape parameter 1 and support between one kilometers and 250 kilometers. Both commutes and city trips in general have been shown to reflect this distribution in many contexts (Ahlfeldt, Redding, Sturm, and Wolf, 2015, Akbar and Duranton, 2018).

Amenity trips join a random origin with an instance of one of 17 amenities (e.g. shopping malls, schools, train stations) as recorded in Google Places. The particular establishment selected is based on a combination of proximity and "prominence" assigned by Google. The weighting across these amenity types is based on a mapping of amenities to trip purposes whose share we draw from the 2008 US National Household Transportation Survey (NHTS) (Couture *et al.*, 2018).

Using the sampling scheme above, we simulated 22,661,818 trip instances in Google Maps, covering 1,166,738 locations pairs and, hence, 2,333,476 trips across all cities and strategies, over 40 days between September and November of 2016. For each trip, we record origin, destination, trip type, length in hundreds of meters, and estimated duration in minutes of Google's recommended route under current traffic conditions (which we sometimes refer to as real-time travel time), as well as the time required for the same route without traffic.

We collected a second sample, 2,115,275 instances of 97,212 trips, in August–September 2018, with several key differences to focus on assessing data quality and measuring travel

time unreliability.⁷ Each trip was collected on at least ten different weekdays in the same 5-minute time-of-day window, in the morning (9 a.m. - 12 p.m.) and/or evening (5 p.m. - 9 p.m.) peak, with durations and distances reported in seconds and meters.

Google's route selection and speed estimates are based on the location and speed of mobile phones using the Android operating system, as well as other phones running Google software, especially Google Maps. Accurate measurement thus requires that drivers are providing information. Kreindler (2018) shows that trip speeds from Google Maps are very close to speeds for actual trips of both cars and motorcycles in Bangalore, measured with a custom-designed smartphone app.

In Appendix A, we also provide a comparison between our data and information provided by Uber Movement about travel speeds in four large cities in our sample. This comparison is complicated by the manner in which Uber Movement aggregates its information. Instead of travel times of actual trips, Uber Movement reports times between zones by averaging travel times from Uber trips that pass through these zones. As we show, this greatly undersamples the beginnings and ends of trips, and these beginnings and ends are considerably slower than the middle parts. Because of this, Uber Movement reports speeds that are substantially faster than our trips. We show that we can nonetheless closely approximate the speed figures obtained from Uber Movement data once we appropriately distort our treatment of Google Maps trip instances and, in particular, focus on travel speeds for their middle parts.

It is however possible that travel times predictions are worse in cities with lower mobile phone penetration. This is unlikely to affect our results. There were 300 million smartphone users in India as of the 4th quarter of 2016 (see Appendix A for details of all web-based sources). In December of 2015, 71% of mobile internet users were urban. Given a 1.324 Billion population of India in 2016, and a 31% urbanization rate from the 2011 Census, a naive calculation implies that 52% of urban residents, including residents of smaller cities, and children, have smartphones. In setting up their phones, users may choose to opt out of sending information to Google. However, the opt-out rate, which Google does not publish, would have to be extremely high to affect our results. Crucially, to estimate slowed traffic on a block, Google only needs one vehicle with a phone, and by definition, time-varying

⁷We collected a further 115,733 trip instances for Bokaro Steel City in December 2017 as the UN database initially reported its location incorrectly. However, we exclude Bokaro Steel City from all results in section 6. In August 2018, we also collected 250,307 trip instances for two American cities, Chicago and Grand Rapids, to produce figure 1 and 526,360 trip instances for four large Indian cities to help compare counterfactual trips from Google Maps to actual trips from Uber rideshares. Finally, we collected data on transit and walking trips in January 2018. We describe them in section 7.

congestion implies many vehicles. Put together, this suggests that all cities have enough phones to generate high-quality speed estimates.

Even if Google has access to all of these data, it might not use them to provide real-time traffic data, reporting instead modeled averages, perhaps especially in smaller cities. In order to test for this, we looked for variation in trip duration (reported in minutes) and trip length (reported in hundreds of meters) across instances of the same trip occurring at the same time of day (within a 10-minute band) on different days, using our main sample. We find that on average 76% of these pairs of trip instances have some variation in either duration or length. Interestingly, this percentage has a correlation across cities of 0.72 with our preferred congestion index. In Bangalore, the most congested city in India, 91% of pairs of trip instances taking place within 10 minutes of each other on different days exhibit variation in either travel time or trip length. This percentage is the lowest in small cities, where we expect trips experiencing little traffic to be especially common and our preferred index of congestion is the lowest. However, even for these cities, we still find substantial variation in the travel times and travel routes suggested by Google Maps. In all cities, at least 60% of pairs of trip instances taking place within 10 minutes of each other on different days exhibit variation in either travel time or trip length.

We believe that this is strong evidence that Google is using real-time traffic information to calculate travel speed, even in small and poor cities. If trips with zero variance reflect bad data, we can also calculate a mobility index on a sample that discards them. The correlation with our preferred mobility index below is 0.998. The trips with no time variation are clearly not driving our results. The 2018 sample, with more precise resolution in meters and seconds, shows this even more clearly. Over 98% of trips in the average city show variation across instances on different days within a 5-minute time-of-day window in evening peak, morning peak, or both. In no city does this value fall below 83% (and 91% excluding Jammu and Srinagar).

Finally, we also estimate differences in travel speed during public holidays that take place in some Indian states during weekdays of our main data collection period from mid-September to early November 2016. This includes, for instance, the Onam (harvest) festival celebrated in Kerala, and Diwali (festival of lights) celebrated in several states at slightly different dates. We find evidence of modestly faster mobility in a large majority of cases and evidence of slower mobility for festivals with major outdoor celebrations drawing a large number of visitors. For instance, a public holiday in Punjab celebrating Guru Ram Das, a

major figure of Sikhism, is associated with modestly faster mobility in Punjabi cities except in Amritsar, which experiences slower mobility. Guru Ram Das is the founder of Amritsar and major celebrations take place in this city during this public holiday. Importantly, we find evidence of mobility effects during public holidays even in smaller cities with a population below half a million. We provide further details of our holiday analysis in Appendix A.

2.3 City-level data

Several pieces of information were derived from administrative data. Daily labor earnings by district and gender are from the Employment and Unemployment Survey of the National Sample Survey (NSS-EUE) 2011–12. Population, and share of population with access to a car or motorcycle by “town” (fourth administrative level), and average commute distance for urban non-agricultural workers by mode and district, are from the 2011 Census. We assign city populations as follows. The population of those towns falling completely within a city light are fully included. Towns falling partially within a city light contribute a share of their population defined by the share of the town’s land area falling in the light. The other census variables (earnings, commutes, share of households with access to a car, motorcycle) are analogously aggregated using the resulting town population shares.

Weather data are from Weather Underground. Data were available for 112 of 154 cities, for a median of eight periods per day. Population growth from 1990 to 2015 is from United Nations (2015). We also use variables that characterize ‘urban shape’ computed by Harari (2016). Data on characteristics of the road network within a (lights-based) city are from OpenStreetMap via GeoFabrik, and processed through OSMnx.

3. A methodology for measuring mobility

3.1 A general conceptual framework

Consider the following general travel problem faced by a household. Its members work and conduct errands at several destinations, selected from a potentially large choice set. Potential destinations are costly to reach. To maximize utility, the household will choose to undertake some trips and not others. Some important decisions like household location and car purchases may also be made simultaneously with local mobility and accessibility. Fully modeling this presents overwhelming theoretical challenges and data requirements.

This travel problem is clearly not tractable unless we drastically simplify it. As a starting point, we note that the household travel problem is not unlike the standard consumption problem where consumers choose their basket from a large number of goods. We often simplify this consumption problem by considering a price index. We can do the same thing for the choice of destinations made by households. In each city, we can consider a number of residential locations and attempt to measure the cost of a ‘typical’ trip. The data requirements are still considerable but no longer overwhelming. The pitfalls of this approach are the same as those associated with typical price indices. Not knowing the preferences of households, it is unclear how travel costs (i.e., the prices) should be aggregated, keeping in mind that different households with different preferences face different price indices.

To minimize these pitfalls, we show that our mobility indices do not depend on how we weight different kinds of trips. In particular, our indices vary little by sampling strategies, type of trip destinations, origin and direction of travel, or time of day. This is because slower cities are often slower at all times, for all types of trips, and throughout the city. As a result, we need not rely on a particular utility specification to tell us how to weight, say, a trip to the train station at peak hour on a weekday relative to a trip to a shopping destination on the weekend.⁸

3.2 *Measuring mobility*

We want to measure the ease of going from an origin to a destination in cities. We focus on the speed of road travel using a motorized vehicle. Data from the 2011 Indian census suggests that 46% of urban commutes, and 55% of urban commutes longer than 1 kilometer, are by motorized road transport. Measuring the speed of travel in a city raises a number of challenges since trips differ considerably in their length, location of origin and destination, time and day of departure, and mode.

The simplest approach is to compute a measure of mean speed for a given city:

$$S_c^m = \frac{\sum_{i \in c} D_i}{\sum_{i \in c} T_i}, \quad (1)$$

⁸While generalized transportation costs involve money, time, and several dimensions of travel comfort and travel conditions (Small and Verhoef, 2007), here we focus on travel time and reliability. This generalization is not as extreme as it seems. First, if we think of travel time as home production and value it at half the wage as is customary in the literature, it represents a large share of the overall cost of travel. Second, many other components of travel costs such as gas consumption and vehicle depreciation are also correlated with travel distance and thus with travel time.

where c denotes a city and i is a trip instance. Because we sum the length D_i of all trip instances in city c and divide by the sum of trip durations T_i , the ratio S_c^m is a length-weighted measure of travel speed. It is straightforward to define the corresponding unweighted mean.

Means are attractive because of their simplicity and ease of computation. However, in our case means may not be comparable across cities. Most importantly, trip length and distance to the center differ systematically across cities. As we show below, these characteristics are important determinants of trip speed. We can condition them out by estimating the following type of regression:

$$\log S_i = \alpha X_i' + s_{c(i)}^{fe} + \epsilon_i, \quad (2)$$

where the dependent variable is log trip speed ($S_i = D_i/T_i$), X_i is a vector of characteristics for trip instance i , $s_{c(i)}^{fe}$ is a fixed effect for city c , and ϵ_i is an error term.

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

Equation (2) does not specify the exact content of the vector of characteristics X . In addition to the city within which a trip takes place, we expect the main variables that determine the speed of a motorized trip in our data to be its length, time of departure, distance to the center, and perhaps the type of the trip. We also expect trip speed to be affected by weather conditions. We derive our benchmark index from equation (2) and we will test the robustness of our estimates of the city fixed effects with respect to which variables are included in the regression and how.

Travel conditions may also vary across cities in ways that may not be well captured by equation (2). For instance, we find below that peak hours are relatively slower and last longer in more congested cities. To capture this, we also estimate a more flexible version of equation (2) where we allow both the constant and the vector of coefficients to vary across cities:

$$\log S_i = \alpha_{c(i)} X_i' + s_{c(i)} + \epsilon_i. \quad (3)$$

Equation (3) includes many coefficients for each city. Comparing for instance the time of day effect for traffic between 9.30 and 10 p.m. across 154 cities will not be insightful. Rather than keep all these coefficients separate, we aggregate them into index measures of mobility for each city.

More specifically, we proceed as follows. We first estimate equation (3) for each city separately. Each of these 154 regressions can be used to generate a predicted speed for all trips in the data, telling us how fast trip i would be if it were taken in city c : $\hat{S}_{ci} = \exp(\hat{\alpha}_c X'_i + \hat{\phi}_c^2/2)$. We also predict speeds from an analogous ‘national’ regression using all trip instances by imposing common coefficients regardless of the city of travel: $\hat{S}_i = \exp(\hat{\alpha} X'_i + \hat{\phi}^2/2)$.

Then, we compute a predicted duration for each trip i if it were to take place in city c ($\hat{T}_{ci} = D_i/\hat{S}_{ci}$) or ‘nationally’ ($\hat{T}_i = D_i/\hat{S}_i$). Finally we can compute a relative speed index for each city:

$$L_c = \frac{\sum_i \hat{T}_i}{\sum_i \hat{T}_{ci}}. \quad (4)$$

The index L_c represents the time it would take to conduct all trip instances in the data at the estimated speed for city c relative to the predicted time it would take to conduct these trips at the average estimated ‘national’ speed. L_c is a unitless scalar, but we can multiply it by $\sum_i D_i / \sum_i \hat{T}_i$, the average national speed, to transform it into a predicted speed for city i .

We note that the index L_c defined in equation (4) resembles a Laspeyres price index in the sense that we compare the speed of trips across Indian cities for the same national bundle of trip instances. Like a standard Laspeyres index, L_c may be sensitive to sampling error or to out-of-sample predictions.

Alternatively, we can compute the predicted time it takes to undertake all city c trips in city c relative to the predicted time it takes to undertake all city c trips from a national regression. That is, we can compute:

$$P_c = \frac{\sum_{i \in c} \hat{T}_i}{\sum_{i \in c} \hat{T}_{ci}}. \quad (5)$$

This alternative speed index is analogous to a Paasche price index. Because we compare city trips at predicted city speed to city trips at predicted national speed, this Paasche index will be less sensitive to the problems of out-of-sample predictions that may afflict the Laspeyres index above. It is also straightforward to compute the corresponding Fisher index: $F_c = \sqrt{L_c \times P_c}$.

Finally, we can compute a broad class of mobility indices derived from logit or CES utility specifications. In the logit case of Ben-Akiva and Lerman (1985), the travel decision is a discrete choice over a set of trip destinations. In Appendix B, we derive the following mobility index, which resembles the (inverse of) the familiar CES price index:

$$G_c = \left(\frac{\sum_{i \in c} b_{ci} T_{ci}^{1-\sigma}}{\sum_{i \in c} b_{ci} \bar{T}_i^{1-\sigma}} \right)^{1/(\sigma-1)}, \quad (6)$$

where b_{ci} is a quality parameter for the destination of trip i in city c , and σ is an elasticity of substitution between trip destinations. In this standard utility maximization framework, cheaper (shorter) trips receive more weight, with the strength of that relationship governed by the elasticity of substitution σ . To construct the denominator of G_c , we use a non-parametric procedure to compute, from the national sample, the average duration \bar{T}_i of trips with approximately the same length as trip i in city c . This procedure delivers a pure mobility index that depends only on speed differences across cities.⁹

Instead of tackling the difficult problem of estimating the parameters of G_c , we show that for a wide range of values of σ and b_{ci} , G_c is highly correlated with our benchmark index from equation (2). We also experiment with richer nesting structures, in which trips to similar destination types (e.g., work, shopping, medical/dental, etc) are more substitutable.¹⁰

It is important to keep in mind that the observations used to estimate equations (2) and (3) and to compute the indices in equations (4), (5), and (6) are counterfactual trips, not actual trips. This presents both benefits and costs. The main advantage of our approach is that trips are exogenously chosen. Unlike Couture *et al.* (2018), we do not need to worry about the simultaneous determination of some variables such as trip length and speed, which could affect the estimates of city fixed effects in equations (2) and (3).¹¹ Conceptually, our approach is similar to measuring price indices from store price tags instead of from consumers' transactions.

This exogeneity is also a potential limitation of our method. The trip instances that we query do not correspond to actual trips and may not be representative of the travel conditions faced by urban travelers when they demand to travel. If our trips are far enough from representative, and if the relative speed of various types of trips varies across cities, then our mobility indices will be mismeasured.

To this criticism, we have four answers. The first is that some of the trips we created were designed to resemble actual trips in other cities, with respect to either their direction and length, or their destination type and frequency. Second, we show below that when

⁹To see this, note that both the city-level numerator and the national-level denominator of G_c have the same number of trips, and the same distribution of trip lengths. The index in each city is therefore free of gains from variety and gains from closer proximity to travel destinations, and determined only by speed differences relative to a national sample.

¹⁰As another example, consider a utility function with limited scheduling flexibility, as in Kreindler (2018). Such a function would increase the weight of trips during peak time. Our approach is to show that mobility indices based on only peak time trips are highly correlated with those based on all trips.

¹¹For instance, as mobility gets better travelers may choose to travel to further destinations. In addition, the (counterfactual) trip instances that we query do not affect real traffic conditions.

we introduce a comprehensive set of controls for other trip characteristics, the economic significance of the trip type indicators in equation (2) is small. Third and most important, our large sample allows us to estimate mobility indices for each trip type, destination, time of day, distance to city center, and various other subsamples. These indices are all highly correlated with our baseline index. As argued earlier, this result implies that our indices do not depend in an important way on the particular utility weight that each counterfactual trip could receive. Finally, Akbar and Duranton (2018) use Google Maps in Bogotá to measure the speed of actual trips reported in a transportation survey and counterfactual trips designed using the same strategy as here. Within short time intervals within days, the speeds of the two types of trips are virtually indistinguishable.

3.3 *Disentangling two sources of mobility: uncongested mobility and congestion.*

Mobility can naturally be decomposed into two components: an uncongested or “free flow” speed, and a congestion factor. To separate the “intrinsic” slowness of a city from its congestion, we can adapt the approach proposed above. To measure mobility, we use as dependent variable in equation (2) the log of actual trip speed and estimate city fixed effects \hat{s}_c^{fe} that we can interpret as an index of mobility. To measure mobility *in the absence of traffic*, we repeat the same estimation as with actual speed but use as dependent variable the log of speed in the absence of traffic returned by Google Maps for each query. The resulting city fixed effects \hat{nt}_c^{fe} are our index of uncongested mobility.¹²

To measure congestion, we repeat the same estimation using the difference between log trip duration with traffic and log trip duration without traffic, $\log T_i - \log T_i^{nt} = \log(T_i/T_i^{nt})$, as the dependent variable. While strictly speaking, the city fixed effects, \hat{f}_c^{fe} , that we estimate are a measure of delay, we can interpret them as a broad index of congestion, which we refer to as the congestion factor.

The dependent variable when estimating mobility is $\log S_i = \log D_i - \log T_i$. The dependent variable when estimating mobility in the absence of traffic is $\log S_i^{nt} = \log D_i - \log T_i^{nt}$. It then follows that when estimating the congestion factor we have $\log T_i - \log T_i^{nt} = -(\log S_i^{nt} - \log S_i)$. Our third regression thus uses as dependent variable the difference between the dependent variables of the first two regressions. Because we estimate these

¹²Alternatively, recall that we observe each trip an average of ten times and oversample times in the middle of the night when we expect very little traffic. We can treat the speed of the fastest trip instance as an estimate of uncongested speed. In practice, these two methods yield city congestion indices with a Spearman correlation coefficient of 0.96.

three regressions for the same trip instances using the same set of covariates, it follows directly from simple econometrics that a city’s congestion factor is the difference between its uncongested mobility factor and its overall mobility factor:

$$\hat{f}_c^{fe} = \hat{n}t_c^{fe} - \hat{s}_c^{fe}. \quad (7)$$

This result is useful on two counts. First, it provides us with an exact decomposition which we exploit below. Second, when we regress these three city fixed effects on the same set of city determinants below, the estimated coefficients will also conveniently add up. For instance, the estimated effect of city population on mobility will be equal to the estimated effect of city population on mobility in the absence of traffic minus the estimated effect of city population on the congestion factor.

3.4 *Measuring travel time unreliability*

Empirical studies in other contexts find that travelers care not only about travel speed, but also about reliable travel times (Brownstone and Small, 2005). For example, unexpected late arrival at work has a cost distinct from that of a predictably long commute. Measuring unreliability for a large sample of trips over many routes is challenging using traditional methods (loop detectors, GPS devices, or recall diaries). Our empirical design using Google Maps is uniquely well-suited for this exercise, because we can query the same trip at the same time (within a five minute time window) on different weekdays.

As in Brownstone and Small (2005), we measure unreliability for a given trip departing at a given time using the percentiles of the travel time distribution across different weekdays. In particular, we compute the unreliability of a trip as the ratio of the 90th to the 50th percentile of its travel time distribution net of city-specific effects for each weekday, to account for, say, systematically faster Thursdays in one city and Mondays in another.¹³ We then compute unreliability indices for each city, using the unreliability ratio of each trip as a dependent variable in the regression in equation (2) with all other controls except weather (which is part of what we want to capture in measuring unreliability). We perform this analysis using our August-September 2018 sample of morning and evening peak trips, collected specifically for the purpose of this unreliability analysis. The scope of our unreliability analysis is broader

¹³The standard deviation of arrival times is also a common measure of unreliability. Using a high percentile of the trip time distribution has the advantage of capturing travelers’ special aversion to being very late (Brownstone and Small, 2005).

than any previous attempts in the literature, and offers the first cross-city evidence on travel time unreliability.

4. Trip-level results

4.1 Descriptive statistics

We queried 22,777,551 unique trip instances. After eliminating a small fraction of trips for which trip length is not well measured or larger than the haversine distance between origin and destination by more than 50 kilometers, we are left with 22,744,156 observations, 14.8% of which are weekend trips.¹⁴

Some basic trip statistics are reported in table 1. Average travel speed is 22 kilometers per hour. While the interquartile range is fairly small at only about 8 kilometers per hour, the tails of the distribution are quite long. Similar observations can be made for trip duration and length. The average trip under actual traffic conditions lasts about 13% longer than its counterpart without traffic. Keeping in mind that we oversampled trips taken at night, we return to this issue below. Finally, the average trip is about 50% longer than its “effective” (haversine) length.

Table 2 reports summary statistics for the 154 cities in our sample. They are on average large, with a mean population above 1.3 million, and fast growing, having doubled in population since 1990.¹⁵ Variation across cities in rates of access to personal motorized transportation and road infrastructure stocks is substantial.

Table 3 reports descriptive statistics for various naive measures of mean city travel speed. Mean travel speed across cities is 24.4 kilometers per hour, which is rather slow in light of the speeds reported for Central Chicago and Grand Rapids in figure 1.¹⁶ This said, 24.4

¹⁴Google Maps often provides problematic routes for motorized travel on short trips. Furthermore, Google Maps rounds trip lengths, and moves our origin and destination points to the nearest road. In extreme cases, such as when a sampled origin is in the middle of a large park, this can lead to routes that are shorter than the haversine distance between the sampled origin and destination. To limit these problems we query only trips longer than one kilometer. These problems still sometimes arise beyond one kilometer.

¹⁵The two sources of population differ both because of the target year and because they are based on slightly different boundaries. In most cases differences are small, but a few cities in Kerala are substantially smaller using our lights-based definition than in the UN database. These cities appear to have a particularly expansive urban agglomeration as defined by the Indian census.

¹⁶This cross-city mean is slightly larger than the overall population mean of 22.1 kilometers per hour reported in table 1 because travel speed is faster in smaller cities for which we have fewer observations.

Table 1: Trip statistics

	Mean	St. dev.	percentile:						
			1	10	25	50	75	90	99
Speed	22.1	7.1	11.5	14.7	17.1	20.6	25.4	31.6	45.8
Duration	20.0	17.6	4	7	9	14	23	40	93
Duration (no traffic)	17.2	14	4	6	9	13	20	33	76
Trip length	8.2	10	1.3	1.9	2.9	4.7	8.9	17.9	54.1
Effective length	5.4	7.0	1.0	1.2	1.8	2.9	5.5	11.9	39.6

Note: 22,744,156 observations. Durations are in minutes, lengths in kilometers; and speeds in kilometers per hour.

Table 2: Summary statistics for Indian cities

	Mean	St. dev.	Min.	Max.
Population ('000, Census/lights, 2011)	1,328	3,031	19	23,889
Population ('000; UN, 2015)	1,545	3,179	307	25,703
Population growth 1900-2015 (%)	106	65	31	399
Total area (km ²)	238	414	5.91	3,569
Total roads length (km)	1,421	3,519	10	33,120
Motorways (km)	66.7	118	0	800
Primary roads (km)	56.1	114	0	687
Share households with car access (%)	9.99	5.78	2.33	31.5
Share households with motorcycle access (%)	41.3	11.7	5.83	73.4
Mean daily earnings (\$)	4.91	1.93	2.00	12.28

Notes: Cross-city averages not weighted by population. 153 cities except for vehicle registrations for which one city is missing.

Table 3: Summary statistics for travel speed in Indian cities

	Mean	St. dev.	Min.	Max.
All trips	24.4	3.79	16.2	34.9
Radial trips	22.2	3.79	14.8	32.8
Circumferential trips	20.6	3.23	14.3	29.5
Gravity trips	22.6	3.42	14.7	30.9
Amenity trips	26.9	6.08	16.6	42.0
All trips, unweighted by length	21.8	2.90	15.7	31.4
All trips, in absence of traffic	26.8	4.49	16.3	38.1
All trips, effective speed	16.4	2.77	11.6	24.0

Notes: 154 cities. Speed in kilometers per hour.

kilometers per hour is higher than the sometimes apocalyptic descriptions found in the popular press.

We note considerable differences in mean speed across cities. The standard deviation across cities is 3.8 kilometers per hour, more than half the standard deviation of 7.2 across trips in table 1. Mean speed for the slowest city is 16.2 kilometers per hour whereas it is more than twice as high for the fastest city at 34.9. We show below that these wide raw speed differences remain once we adequately control for features of our sampling strategy.

The second to the fifth rows of table 3 report mean speed for each type of trip separately. Circumferential trips are slower whereas amenity trips are faster. As we show below, these differences are mostly caused by differences in the length and location of these trips.

The sixth row of table 3 reports a measure of mean speed by city, which, unlike the other rows, is not weighted by trip length. Because this increases the influence of shorter trips that are also slower, this unweighted mean of 21.8 kilometers per hour is slightly lower than the length-weighted mean of 24.4 reported in the first row.

The seventh row of table 3 exploits the information provided by Google Maps regarding trip duration in the absence of traffic. As expected, mean speed in the absence of traffic is higher but the difference is small. At 26.8 kilometers per hour, mean speed in the absence of traffic is only about 10% above the mean of actual speed reported in the first row. Interestingly, the variation across cities is not smaller for mean speed in the absence of traffic than for actual mean speed. If anything, it becomes slightly larger. We return to this intriguing finding below.

Finally, the last row of table 3 reports a measure of mean effective speed. Rather than trip length, we use the haversine distance between the origin and destination. Since the ratio between mean trip length and effective trip length is about 1.5 in table 1, we unsurprisingly find a roughly similar ratio between actual and effective trip speed.

4.2 *Trip regressions*

Before an in-depth analysis of mobility indices and their correlates, we first estimate a number of variants of the generic regression described by equation (2).

A first series of results is reported in table 4. Column 1 regresses log trip speed on city fixed effects controlling for log trip length, an indicator for each type of trip, each day of the week, and each thirty-minute period during the day. Column 2 introduces further controls:

Table 4: Determinants of log trip speed

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
log trip length	0.24 ^a (0.0036)	0.14 ^a (0.012)	0.14 ^a (0.012)	0.24 ^a (0.0036)	0.14 ^a (0.012)	0.14 ^a (0.012)	0.13 ^a (0.015)
log trip length ²		0.014 ^a (0.0034)	0.014 ^a (0.0035)		0.014 ^a (0.0034)	0.014 ^a (0.0034)	0.016 ^a (0.0045)
log distance to center		0.15 ^a (0.042)	0.15 ^a (0.042)		0.14 ^a (0.041)	0.14 ^a (0.041)	0.098 (0.063)
log distance to center ²		0.025 (0.023)	0.025 (0.023)		0.031 (0.022)	0.031 (0.022)	0.041 (0.034)
Type: circumferential	-0.015 ^a (0.0020)	-0.0039 ^b (0.0016)	-0.0040 ^b (0.0016)	-0.015 ^a (0.0020)	-0.0037 ^b (0.0016)	-0.0038 ^b (0.0016)	-0.0017 (0.0019)
Type: gravity	0.077 ^a (0.0065)	-0.0032 (0.0032)	-0.0032 (0.0032)	0.079 ^a (0.0066)	-0.0027 (0.0032)	-0.0027 (0.0033)	0.00098 (0.0043)
Type: amenity	0.082 ^a (0.0058)	0.0064 ^c (0.0036)	0.0063 ^c (0.0036)	0.083 ^a (0.0057)	0.0066 ^c (0.0036)	0.0065 ^c (0.0036)	0.0087 (0.0054)
City effect	Y	Y	Y	Y	Y	Y	Y
Day effect	Y	Y	Y	weekd.	weekd.	weekd.	Y
Time effect	Y	Y	Y	Y	Y	Y	Y
Weather	N	N	Y	N	N	Y	only
Observations	22,744,156	-	-	19,385,656	-	-	10,319,939
R-squared	0.48	0.53	0.53	0.48	0.53	0.53	0.51
Cities	154	154	154	154	154	154	107

Notes: OLS regressions with city, day, and time of day (for each 30 minute period) indicators. Log speed is the dependent variable in all columns. Robust standard errors in parentheses. *a*, *b*, *c*: significant at 1%, 5%, 10%. All trip instances in columns 1-3. Only weekday trip instances in columns 4-6. Sample sizes for columns 1 and 4 apply to columns 1–3 and 4–6, respectively. Only weekday trip instances for which we have weather information in column 7. Weather in column 3 and 6 consists of indicators for rain (yes, no, missing), thunderstorms (yes, no, missing), wind speed (13 indicator variables), humidity (12 indicator variables), and temperature (8 indicator variables). These variables are introduced as continuous variables in column 7.

the square of the log trip length, log distance to the center (defining a trip's location as the midpoint between its origin and destination), and its square. Column 3 further adds weather variables (and indicators for missing weather data). Columns 4 to 6 repeat the specifications of columns 1 to 3 on a sample of only weekday trips. Column 7 is restricted to observations with non-missing weather data.

Table 4 reports selected coefficients. Longer trips are faster: the elasticity of trip speed with respect to trip length is 0.24 in columns 1 and 4, and larger for longer trips in the other columns where we introduce a quadratic term. This is a prominent feature of urban

transportation data in other contexts.¹⁷ Regressing log trip speed on log trip length without any further control yields an R^2 of 0.40.

Unsurprisingly, trips further from the center are also faster. The elasticity of trip speed with respect to distance from the center of 0.15 is quite large, implying that a trip at 10 kilometers from the center of a city is about 40% faster than one a kilometer away.

In column 1, we find fairly large differences of up to 10% in speed between different types of trips. These differences become mostly insignificant and economically small when controls for trip location are added in column 2. In the end, amenity trips are slightly faster while circumferential trips are slower but the speed difference between them is only about 1%. We also note that regressing log trip speed solely on trip type indicators yields an R^2 of only about 0.003. These two results are reassuring, and suggest that the design of our hypothetical trips is not driving our results. In Appendix D, we report versions of table 4 for each type of trip. While the non-linearities for the effect of trip length and distance to the center slightly differ, the results overall are similar to those in table 4, suggesting that the simple additive specification of table 4 is not obscuring deeper differences between trip types.

We now turn to the regression coefficients not reported in table 4. Starting with the weather, we find that characteristics associated with bad weather such as rain, high levels of humidity, high temperatures, and more windy conditions tend to be associated with slightly *higher* travel speeds. For instance, in columns 3 and 6, trips in rain are 2–3% faster.


To explain this contrast, we conjecture that roads in many Indian cities are ‘multi-purpose’ public goods used by various classes of motorized and non-motorized vehicles to travel and park as well as a wide variety of other users such as street-sellers, animals, or children playing. Non-transportation uses of the roadway arguably slow down motorized vehicles. Worse weather may reduce these activities and thus make travel faster. We provide further indirect evidence for this conjecture below.¹⁸

As expected, we also observe fluctuations in travel speed across times of day. In figure 3, which mirrors figure 2 but isolates hour effects, the dark continuous line plots the speed relative to 3 - 3.30 a.m. for each thirty-minute period estimated in column 5 of table 4 for all

¹⁷Couture *et al.* (2018) estimate a larger elasticity close to 0.40 using self-reported us data where the measure of trip duration also includes a fixed cost of getting into one’s vehicle and getting into traffic. Using self-reported data, Akbar and Duranton (2018) find an even larger elasticity for Bogotá travelers, because their sample also includes transit trips, with even larger fixed costs. Using analogous Google Maps data for the same Bogotá trips, Akbar and Duranton (2018) find an elasticity of 0.21, very close to the elasticity estimated here.

¹⁸However, it is important to note that our data collection period did not include monsoon season. Extreme weather conditions may affect mobility negatively, including for a period of time after they end.

Figure 3: Estimated time effects for weekday travel




Time effects (in percentages, relative to 3 - 3.30 a.m.) as estimated in column 5 of table 4. All the coefficients from 7am to midnight are significant at 1%.

cities. The gap between the fastest time in the middle of the night and the slowest at 6.30 p.m. - 7 p.m. is just 13%. We also note that morning peak hours are more muted than the evening peak hours.¹⁹ The figure also plots the same time profile estimated only on the twenty largest cities. The patterns are much more marked. The slowest periods in the evening are now more than 25% slower than the fastest in middle of the night. In addition, travel speed starts declining earlier in the morning and recovers later in the evening.

While larger, this difference remains less important than that estimated by Akbar and Duranton (2018) for Bogotá where the slowest period is about half as fast as the fastest. These mild within-day fluctuations may mask a lot of heterogeneity across Indian cities. To investigate this, we repeat the same exercise using only observations from the city of Delhi. Although Delhi is slow, we purposefully do not take the slowest city or a pathological case. The pattern is the same as for the 20 largest cities but more pronounced. The slowest time is now 35% slower than the fastest. Restricting attention further to trips taking place on average within five kilometers of the center of Delhi generates even more extreme patterns with the

¹⁹Although we do not report the results here, we can also estimate time of day effects more accurately using trip fixed effects. The resulting estimates for time of day effects are virtually indistinguishable.

Figure 4: Kernel density for estimated city effects



The city effects are as estimated in column 5 of table 4 for all 154 cities. Epanechnikov kernel with bandwidth of 0.031.

slowest time now being more than 40% slower than the fastest.²⁰

If we take the difference between the fastest and slowest time as a summary measure of congestion, we can draw several lessons from figure 3. First, in many cities, there may not be that much congestion. Travel speed is slow and does not vary much throughout the day as the demand for travel changes. It is only in the largest cities and more particularly in their centers that travel speed experiences considerable variation during the day. We return to this below. Second, the evolution of travel speeds during the day reflects more than standard commuting patterns. Travel speed declines from roughly 5.30 a.m. to midday, the lowest speed are observed around 6.30 - 7 p.m., and only slowly recover late into the evening. This is consistent with the conjecture raised above that the roadway is used for multiple purposes from late in the morning until well into the evening.

We finally turn to city effects. As argued above, we can interpret them as mobility index values. They measure (log) trip speed in cities after conditioning out log trip length and its square, log trip distance to the center and its square, and day and time of day effects. Figure

²⁰Since India is a vast country with a single time-zone, attenuated within-day fluctuations could be due to the timing of sunrise and sunset, which differs across cities. Within our sample, there is range of up to a 98 minutes in sunrise and 126 minutes in sunset. To assess whether cities experience peak hours at different official hours, we produced a variant of figure 3 that defines the time of each trip as a fraction of the time between local sunrise and sunset (or between local sunset and sunrise). It is virtually indistinguishable from figure 3.

4 represents a kernel density estimate of the distribution of city fixed effects from column 5 of table 4. The standard deviation is 0.106. The slowest city is 28% slower than the mean while the fastest city is 42% faster. This gap of a factor of two between the slowest and fastest city is extremely large. Using traveler-reported data and a different methodology, Couture *et al.* (2018) find that among the largest 50 US metropolitan areas, the fastest is less than 40% faster than the slowest. The analogous difference for the largest 50 cities in India is 80%. These large differences are unlikely to be due to sampling bias. All cities have at least 70,000 observations, and the largest cities have more than half a million.

Table 5 reports the 20 slowest and 20 most congested cities. First, we note that seven of the 10 largest cities by population in 2015 are among the 20 slowest. The three exceptions are Ahmadabad and Surat in Gujarat and Jaipur in Rajasthan. The state of Gujarat stands out in India for its innovative and more efficient urban planning practices (Annez, Bertaud, Bertaud, Bhatt, Bhatt, Patel, and Phata, 2016). The list of the 20 slowest cities also contains 6 cities from the state of Bihar (among 8 in our data). Bihar is the poorest state in India. Most of the other slow cities are from the neighboring states of Jharkhand and Uttar Pradesh, which are also among the five poorest states in India.

For the 20 most congested cities, two key attributes stand out. First, there are only seven cities with a congestion factor of more than 20%; that is, cities for which the trips we queried were on average slower by 20% or more relative to their speed in the absence of traffic. These cities are all among the largest cities of India. Second, the correlation between congested and slow appears fairly low. Only 9 of the twenty most congested cities are among the twenty slowest, and one is actually among the ten fastest.

The list of the fastest cities is more heterogeneous. Many are small and in more developed parts of India. Others are exceptional in different ways. The fastest, Ranipet, is an independent city based on our delineation procedure. However, it may be viewed more meaningfully for our purposes as a suburb of the city of Vellore, located about 20 kilometers away. Chandigarh hosts a population above a million, but unlike most Indian cities, it is a planned city characterized by a regular grid pattern laid out by the French architect Le Corbusier.²¹ Srinagar and Jammu are both in the disputed state of Jammu and Kashmir. While they were among the fastest cities in 2016, our second data collection reveals that they were among the 20 slowest in 2018. We believe that the high levels of mobility we observe for

²¹Figure A.5 in appendix shows Chandigarh's road network, which has the most regular grid of all Indian cities in our sample.

Table 5: Ranking of the 20 slowest and 20 most congested cities, slowest and most congested at top

Rank	Slowest			Most congested		
	City	State	Index	City	State	Index
1	Kolkata	West Bengal	-0.33	Bangalore	Karnataka	0.33
2	Bangalore	Karnataka	-0.25	Mumbai	Maharashtra	0.30
3	Hyderabad	Andhra Pradesh	-0.25	Delhi	Delhi	0.30
4	Mumbai	Maharashtra	-0.24	Chennai	Tamil Nadu	0.27
5	Varanasi	Uttar Pradesh	-0.23	Kolkata	West Bengal	0.27
6	Patna	Bihar	-0.22	Hyderabad	Andhra Pradesh	0.24
7	Delhi	Delhi	-0.22	Pune	Maharashtra	0.23
8	Bhagalpur	Bihar	-0.22	Jaipur	Rajasthan	0.19
9	Bihar Sharif	Bihar	-0.19	Coimbatore	Tamil Nadu	0.18
10	Chennai	Tamil Nadu	-0.17	Ahmadabad	Gujarat	0.17
11	Muzaffarpur	Bihar	-0.16	Patna	Bihar	0.16
12	Aligarh	Uttar Pradesh	-0.15	Chandigarh	Chandigarh	0.16
13	Darbhangha	Bihar	-0.14	Kochi	Kerala	0.16
14	English Bazar	West Bengal	-0.14	Thiruvananthapuram	Kerala	0.15
15	Gaya	Bihar	-0.13	Begusarai	Bihar	0.15
16	Allahabad	Uttar Pradesh	-0.13	Kozhikode	Kerala	0.15
17	Ranchi	Jharkhand	-0.12	Nagpur	Maharashtra	0.14
18	Dhanbad	Jharkhand	-0.12	Ranchi	Jharkhand	0.14
19	Akola	Maharashtra	-0.12	Vijayawada	Andhra Pradesh	0.14
20	Pune	Maharashtra	-0.11	Bhopal	Madhya Pradesh	0.14

Notes: Mobility index is measured by the city effect estimated in column 5 of table 4 and is centered around its mean. Congestion factor is measured from a similar regression using log trip duration minus log trip duration in absence of traffic as dependent variable, and scaled to reflect log deviation from uncongested speed.

Table 6: Ranking of the 10 fastest cities, fastest at top

Rank	City	State	Index
1	Ranipet	Tamil Nadu	0.35
2	Srinagar	Jammu and Kashmir	0.26
3	Kayamkulam	Kerala	0.24
4	Jammu	Jammu and Kashmir	0.23
5	Thrissur	Kerala	0.19
6	Palakkad	Kerala	0.16
7	Chandigarh	Chandigarh	0.16
8	Alwar	Rajasthan	0.15
9	Thoothukkudi	Tamil Nadu	0.15
10	Panipat	Haryana	0.15

Notes: Mobility index is measured by the city effect estimated in column 5 of table 4.

2016 are an indirect consequence of the 2016-17 unrest in Kashmir which occurred following the killing of the commander of a militant organization. Most of the state remained under shutdown, curfew, and/or heavy police surveillance during our main data collection.

Table 7 reports a number of variants of our benchmark specification in table 4 column 5. Column 1 uses log effective speed (haversine length divided by time) instead of actual speed as dependent variable. The increase in effective speed with trip length and with trip distance to the center is even more pronounced than the increase in actual speed. This is consistent with shorter and more central trips being more tortuous. Column 2 uses speed under “typical” traffic conditions at the time of day of the request (also provided by Google Maps) as dependent variable. Results are very similar to those for the corresponding specification using actual speed in column 5 of table 4. Column 3 uses the same specification to predict speed with no traffic. Interestingly, trips taking place further from the center remain faster. While figure 3 above suggests that central parts of Delhi are more congestible, the bulk of the difference in speed between more central and more peripheral trips remains in the absence of traffic. This is plausibly caused by the expected greater density of intersections and narrower streets in more central parts of cities in India (and many other countries).

The second part of table 7 reports our preferred specification of table 4 for different times of day: off peak in column 4, low peak in column 5, high peak in column 6, and radial trips at peak hours going towards the center in the morning and back towards the periphery in the evening in column 7. This last specification is meant to mimic archetypal commuting patterns. While again the curvature of the effect of trip length and distance to the center varies slightly, the results are generally very similar to those we obtained before.

4.3 Comparing mobility indices

We now turn to comparing mobility indices. Because many different variants of equations (2) and (3) are available and many different samples of trips can be selected, many mobility indices are possible. To explore these possibilities, we compute a wide variety of such indices. To avoid hard-to-digest matrices of pairwise correlations, we form our benchmark mobility index from the city fixed effects estimated from the specification reported in column 5 of table 4, and compare all our other indices to this one. We also report the standard deviation, maximum and minimum of each variant. Standard deviations vary very little, except for the

Table 7: Determinants of log trip speed, variants

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	effective length	typical traffic	no traffic	off peak	low peak	high peak	peak radial
log trip length	-0.18 ^a (0.012)	0.13 ^a (0.012)	0.16 ^a (0.012)	0.14 ^a (0.011)	0.13 ^a (0.013)	0.13 ^a (0.012)	0.040 (0.030)
log trip length ²	0.085 ^a (0.0031)	0.017 ^a (0.0039)	0.019 ^a (0.0032)	0.019 ^a (0.0031)	0.013 ^a (0.0039)	0.0098 ^a (0.0034)	0.065 ^a (0.010)
log distance to center	0.57 ^a (0.036)	0.16 ^a (0.046)	0.22 ^a (0.046)	0.23 ^a (0.048)	0.12 ^a (0.042)	0.087 ^b (0.036)	0.15 ^a (0.051)
log distance to center ²	-0.13 ^a (0.015)	0.014 (0.026)	-0.037 (0.025)	-0.047 ^c (0.027)	0.054 ^b (0.023)	0.083 ^a (0.019)	-0.12 ^a (0.044)
City effect	Y	Y	Y	Y	Y	Y	Y
Day effect	weekd.	weekd.	weekd.	weekd.	weekd.	weekd.	weekd.
Time effect	Y	Y	Y	Y	Y	Y	Y
Weather	N	N	N	N	N	N	N
Observations	19,385,656	19,385,656	19,385,656	4,910,731	10,469,622	2,375,960	826,539
R-squared	0.34	0.56	0.54	0.54	0.54	0.53	0.54
Cities	154	154	154	154	154	154	154

Notes: OLS regressions with city, day, and time of day (for each 30 minute period) indicators. Log effective speed is the dependent variable in column 1. Log speed under “typical” traffic conditions is the dependent variable in column 2. Log speed under ‘no traffic’ is the dependent variable in column 3. Log speed is the dependent variable in all subsequent columns. All columns only consider weekday observations. Column 4 considers observation from only off-peak hours (before 7.30 a.m. and after 10.30 p.m.). Column 5 considers observation from only low peak hours (from 8.30 a.m. to 5.30 p.m. and from 8 p.m. to 10 p.m.). Column 6 considers observations from only high peak hours (from 5.30 p.m. to 8 p.m.). Finally, column 7 considers only radial observation (going towards the city center in the morning and back in the evening) from peak hours (both low and high peak). Robust standard errors in parentheses. *a*, *b*, *c*: significant at 1%, 5%, 10%.

mean speed indices, which are constructed on a different (linear) scale.

The results are reported in table 8. Panel A compares our benchmark mobility index to the analogous indices estimated in the other columns of table 4 that includes various trip level controls. All these correlations are above 0.98 when we include the square of trip length and distance to center and fall to about 0.92 when we do not.

Panel B compares our benchmark index to the analogous indices estimated using the same specification but considering different types of trips separately. The correlations are again high. The lowest at 0.90 is with perhaps our most artificial type of trips, circumferential trips, and the highest is with perhaps our most realistic, amenity trips. In results not reported here,

Table 8: Pairwise Spearman rank correlations with our benchmark mobility index

Index	Corr.	Std. Dev.	Min	Max
Panel A: Columns from table 4				
(1)	0.916	0.100	-0.232	0.332
(2)	>0.999	0.105	-0.321	0.347
(3)	0.992	0.108	-0.337	0.355
(4)	0.918	0.101	-0.240	0.332
(6)	0.991	0.109	-0.347	0.356
(7)	0.983	0.115	-0.329	0.374
Panel B: Trip subsamples				
Radial	0.926	0.117	-0.318	0.373
Circumferential	0.900	0.113	-0.286	0.330
Gravity	0.966	0.112	-0.375	0.308
Amenities	0.966	0.107	-0.345	0.373
Interact time/day with trip type	>0.999	0.105	-0.322	0.347
Panel C: Mean speeds				
Simple mean	0.476	3.790	16.212	34.903
Mean unweighted by length	0.619	2.899	15.7	31.4
Mean of "typical" traffic speed	0.452	3.814	16.2	35.1
Mean of uncongested speed	0.340	4.494	16.3	38.1
Mean effective speed	0.410	2.768	11.6	24.0
Panel D: Table 7 variants				
Effective speed	0.864	0.118	-0.430	0.392
"Typical" traffic	0.997	0.102	-0.301	0.345
No traffic	0.850	0.100	-0.242	0.339
Fastest trip instance	0.851	0.101	-0.261	0.298
Off peak	0.881	0.099	-0.255	0.316
Low peak	0.991	0.113	-0.388	0.361
High peak	0.948	0.130	-0.430	0.367
Peak radial	0.915	0.133	-0.450	0.405
Panel E: Full indices				
Laspeyres	0.794	0.151	0.105	1.478
Paasche	0.941	0.107	0.767	1.478
Fisher	0.910	0.126	0.322	1.478
Logit/CES ($\sigma = 0$)	0.923	0.107	0.797	1.482
Logit/CES ($\sigma = 2$)	0.836	0.113	0.819	1.441
Logit/CES ($\sigma = 4$)	0.687	0.142	0.846	1.542
Panel F: Distance to center				
Trips within 5 km of center	0.970	0.108	-0.278	0.350
Trips within 3 km of center	0.918	0.111	-0.268	0.356
Trips within 2 km of center	0.827	0.116	-0.261	0.336
Weight by inverse dist. to center	0.959	0.106	-0.293	0.341
Panel G: Weight by powered congestion factor				
$\lambda = 0.2$	0.910	0.137	-0.533	0.378
$\lambda = 0.3$	0.927	0.123	-0.396	0.373
Panel H: 2018 data				
Baseline index for 2018	0.817	0.119	-0.334	0.329
Baseline index for 2018 (152 cities)	0.933	0.118	-0.334	0.329

Notes: 154 cities in all rows except in the last row of panel A which uses 107 and the last row of panel H which uses 152. The first column reports the Spearman rank correlation between the index at hand and our preferred index from column 5 of table 4. The second column reports the standard deviation. The third and fourth column report the maximum and minimum respectively.

we note that even indices computed separately for each of our 17 individual amenity classes, which represent less than 3% of a city's trips in nearly all cases, are highly correlated. Fifteen of them are correlated with the baseline index at 0.87 or higher. Finally, allowing time of day and weekend indicators to vary by trip type (radial inward, radial outward, circumferential, gravity, and 17 amenity types), so that, for example, trips to a temple on the weekend might be different than those on a weekday, also makes essentially no difference in rankings.

Next, panel c compares our benchmark index to various measures of mean speed computed above. The correlations are much lower than in the previous two panels. For instance, the correlation between our benchmark mobility index and mean speed computed as total travel length divided by total travel time is only 0.48. As noted in Couture *et al.* (2018) for us metropolitan areas, means of speed do not provide good descriptions of mobility in cities. This is because trip length, which varies systematically across locations, has a large explanatory power on trip speed. As a result, mean speeds are sensitive to sampling strategies, unlike our preferred mobility indices that control for trip length.

Panel d reports correlations between our benchmark mobility index and mobility indices computed from the estimations reported in table 7. The correlation of our benchmark mobility index with an index that measures speed using effective (haversine) rather than traveled trip length is 0.87. The 20 slowest cities reported in table 5 using our benchmark mobility index are all among the 30 slowest cities by effective speed. We can thus rule out the possibility that slow cities are more efficient at transporting travelers farther for the same number of straight line kilometers traveled. Slow cities are just slow.

Still in panel d, the correlation of our benchmark index with an uncongested mobility index, computed using travel times in the absence of traffic, is also relatively high at 0.85. This strongly suggests again that poor mobility is largely the outcome of generally slow travel. While congestion plays a role, it may not be the main driver of poor mobility in Indian cities. We return to this issue below. Interestingly, when ranking cities by uncongested mobility, we find that the five slowest cities in the absence of traffic are all in Bihar and 17 of the 20 slowest cities are in the poor northeastern part of India. Except for Kolkata which also ranks among the cities that are slow in the absence of traffic, most major Indian cities are in the middle of the distribution of uncongested mobility indices. For these cities, congestion is arguably an important determinant of why they are slow. Eight of the 10 fastest cities reported in table 6 are also among the 10 fastest cities in the absence of traffic.

The second part of panel d reports correlations between our benchmark index and mobility

indices computed in the same manner as our benchmark but from observations taken at specific hours of the day. The correlation of our benchmark index with an index of low peak hour trips, taken between 8.30 a.m. and 5.30 p.m. and between 8 p.m. and 10 p.m., is extremely high. It is still high with an index computed only during the most extreme hours of the early evening, between 5.30 and 8 p.m., when traffic is generally at its slowest. The correlation is still 0.92 with an index computed using only the 5% of sample composed of radial trips at peak hours that go towards the center in the morning and away from the center in the evening.

Panel E reports correlations between our benchmark index and more sophisticated Laspeyres, Paasche, Fisher, and logit/CES indices computed as described by equations (3), (4), (5), and (6). Row 1 uses a Laspeyres index computed from the same specification as for our benchmark index which allows all 58 regression coefficients to vary across cities. The correlation is still fair at 0.79. It jumps to 0.89 when we focus only on the 50 largest cities. The lower full-sample correlation is due to flawed out-of-sample predictions in small cities for long trips far from the center. The Paasche index in row 2, which does not suffer from the same problem, has a much higher correlation with our benchmark index.

Rows 4 to 6 of panel E report correlations with the logit/CES index for different values of the elasticity of substitution σ . The correlation for $\sigma = 0$, the perfect complement case for which all trips receive equal weight, is very high at 0.92, and only declines slightly to 0.84 for $\sigma = 2$. The correlation with our benchmark index remains relatively high at 0.69 even for an extreme value of $\sigma = 4$, which gives a two-kilometer trip about 400 times the weight of a longer 15-kilometer trip.²² In Appendix B, we describe simulations showing that correlations remain invariably high across a wide range of random quality draws b_{ci} . In the same appendix, we describe mobility indices from models of travel demand with richer substitution patterns. These nested indices put less weight on destination types (e.g., shopping trips) that are relatively slower in a given city, because they allow travelers in each city to substitute away from costlier travel destination types. We find that such nested indices are highly correlated with our benchmark index. This finding further confirms that our benchmark index provides a robust characterization of travel cost differences across cities, because slow cities tend to be slow at all times, for all types of trip destinations, and across the city.

²²Atkin, Faber, and Gonzalez-Navarro (2018) estimate an elasticity of substitution across retail stores slightly smaller than 4 for poor Mexican households. This is almost certainly an upper bound: the index considered here covers a much broader set of destinations that are unlikely to be as substitutable as retail stores.

Panel F considers indices based on trips progressively closer to the center of the city. Correlations fall as expected, but even limiting to trips centered within 2 kilometers of the center, the correlation is still 0.83. Using an index that weights trips close to the center more heavily, while still including more peripheral trips, yields an index much more similar to the benchmark.

In panel G we try to weight each trip by how likely it is to be taken. Although this information is not directly available to us, we can use the implicit density of vehicles along the route as a proxy. To do so, we assume that (i) the speed of a trip instance is reduced from the maximum for that trip solely by vehicle congestion, (ii) the elasticity of trip speed with respect to the density of vehicles, λ , is constant, and (iii) the density of vehicles is constant along the route. Under these assumptions, we can weight each trip i by its length, D_i , times the implicit density of vehicles, $(T_i/T_i^{nt})^{1/\lambda}$. While these assumptions are unlikely to be strictly true, they manage to capture the fact that more vehicles slow down traffic and thus slower trip instances should receive a higher weight given that they represent more travelers. The question is of course which value to use for λ . We use $\lambda = 0.2$ and $\lambda = 0.3$. The value $\lambda = 0.2$ is a standard value in the traffic modelling literature (Small and Verhoef, 2007). The higher value $\lambda = 0.3$ reduces the weight put on slow trips since slower speeds in India may not be caused only by more vehicle traffic. With both values, the indices are highly correlated with our benchmark index.

Finally, in panel H, we re-estimate our baseline index using the data we collected in the Summer of 2018. Although we collected these data with a different set of questions in mind, we can, with minor caveats, use them to estimate our preferred mobility index. Overall, the correlation between our baseline index for 2016 and 2018 is 0.82. This correlation jumps to 0.93 if we discard Srinagar and Jammu which experienced a large change in mobility, arguably due to changes in their security situation as suggested above. We also find minor discrepancies for a number of cities in Kerala that were subject to flooding during our data collection in August 2018. We also note that these results are not sensitive to the exact samples used to compare the datasets but leave further longitudinal analysis to future work.

We draw two important conclusions from this analysis. First, because trip length is such an important determinant of trip speed, and because trip length varies across cities of different sizes, appropriately estimating a city mobility index requires accounting for trip-length differences. This finding highlights the importance of using entire trip instances as units of analysis, instead of trip segments or travel speed at discrete locations. Second, we find that

once trip length is conditioned out, the mobility indices that we estimate for each city are not sensitive to the exact sample being used, and therefore to the weight that different kinds of trips receive. Although we use a variety of trips that reflect important differences in traveler behavior, these differences do not appear to matter when estimating city mobility.

4.4 Unreliability results

To complete our description of trip-level results, we document significant unreliability in travel time across urban India. Across different weekdays, trip times at the 90th percentile are on average 6% slower than median trip times. In more than 80% of cities in our sample, average unreliability is between 4% and 7%. Trip time is more unreliable in larger cities, with an average unreliability of 9% in the 20 largest cities.²³ The most congested cities are also the most unreliable, with a correlation between our congestion and unreliability indices of 0.84. Evening peak trips are 1 percentage point more unreliable than morning peak trips, and we do not find evidence that trips closer to the city center are more unreliable. Unreliability is uncorrelated with uncongested mobility. We note that our unreliability estimates from Google Maps assume optimal re-routing depending on road conditions, and Indian travelers without the benefit of information on current traffic or unwilling to re-route may face even more unreliable travel times. In the welfare analysis below, we interpret these results using existing empirical evidence on how individuals value reliable travel time.

5. Decomposition: uncongested mobility and congestion

In this section, we decompose our indices of mobility into mobility in the absence of traffic (uncongested mobility) and the congestion factor following equation (7). This relationship allows us to perform two useful exercises: first, an exact decomposition of the variance in our mobility index, and second, a simple analysis to compare the welfare gains from faster uncongested mobility with those from reduced congestion and reduced unreliability.

²³Four of the five most unreliable cities are in Kerala, likely due to severe flooding (the state's worst since 1924) during our sampling period. The most unreliable city outside of Kerala is Mumbai with an average 90th/50th percentile ratio of 12%. Removing the least or most reliable cities has little impact on average unreliability.

5.1 Variance decomposition

The variance of the mobility index is equal to the sum of three terms: the variance of the index of uncongested mobility, the variance of the congestion factor, and minus twice the covariance between the index of uncongested mobility and the congestion factor. As shown in the first row of Table 9 Panel A, the variance of the uncongested mobility index accounts for 88% of the variance of our benchmark mobility index while that of the congestion factor accounts for only 32%. This is a striking finding. Differences in mobility between Indian cities are mostly driven by differences in their uncongested mobility, not by differences in how congested they are. As we show in the rest of this section, this finding is explained by both pervasive differences in uncongested mobility between cities and the fact that congestion remains modest in most cities. However, the finding is different when we focus on the largest cities. These cities face more similar uncongested mobility but are congested to different degrees.

This said, a possible caveat here is that our data collection oversamples trips at night and this may bias our mobility index towards uncongested mobility. Performing the same exercise with indices computed only from trips taken at peak hours, we find that the uncongested mobility index still represents 75% of the variance of the mobility index during peak hours whereas the congestion factor represents only 48%.

We repeat the same exercise focusing on cities with population above the median. For these cities, the role of uncongested mobility falls, but remains larger than the congestion factor, and the covariance term essentially goes to zero. For cities below the median population, the explanatory power of the congestion factor is very low. For cities in the top population quartile, the covariance term becomes negative, but the uncongested mobility still represents a larger share of the variance. Only in the top decile do the two factors have approximately even shares.

In the next two panels of Table 9, the role of congestion expands as we limit attention to city centers, especially at peak hours and in larger cities. Variance in uncongested mobility still however represents a substantial share of overall variance across cities in all samples. In the final panel, we repeat the same decomposition for each type of trip separately and find roughly similar results for the respective roles of uncongested mobility and congestion.

Table 9: Variance decompositions of our baseline mobility index

Sample	Cities	All trips			Peak trips		
		Uncongested mobility	Congestion factor	Covariance	Uncongested mobility	Congestion factor	Covariance
Panel A: Full trip sample							
All	154	0.884	0.318	0.101	0.745	0.479	0.112
Largest 50%	77	0.646	0.346	-0.004	0.515	0.508	0.012
Smallest 50%	77	1.305	0.126	0.215	1.348	0.178	0.263
Largest 25%	38	0.526	0.287	-0.093	0.414	0.415	-0.085
Largest 10%	15	0.357	0.376	-0.134	0.262	0.497	-0.121
Panel B: Distance to city center less than 5 km							
All	154	0.946	0.387	0.167	0.788	0.578	0.183
Largest 50%	77	0.725	0.449	0.087	0.566	0.645	0.105
Smallest 50%	77	1.300	0.127	0.213	1.336	0.178	0.257
Largest 25%	38	0.558	0.459	0.009	0.412	0.623	0.017
Largest 10%	15	0.452	0.784	0.118	0.295	0.920	0.107
Panel C: Distance to city center less than 3 km							
All	154	1.027	0.407	0.217	0.872	0.618	0.245
Largest 50%	77	0.809	0.447	0.128	0.646	0.656	0.151
Smallest 50%	77	1.306	0.133	0.220	1.345	0.186	0.266
Largest 25%	38	0.583	0.513	0.048	0.425	0.685	0.055
Largest 10%	15	0.595	0.922	0.258	0.383	1.078	0.231
Panel D: By trip type							
Radial	154	0.960	0.369	0.164	0.803	0.560	0.181
Circumferential	154	1.034	0.397	0.216	0.879	0.604	0.241
Gravity	154	0.789	0.223	0.006	0.678	0.334	0.006
Amenities	154	0.841	0.302	0.071	0.705	0.448	0.076

Note: 154 cities. Peak hours as in table 7.

5.2 Valuing improvements in mobility and congestion.

We now use economic theory to guide our valuation of faster uncongested mobility and reduced congestion in urban India. In particular, we model inverse travel speed as the price of travel distance, which has a downward sloping aggregate demand curve. The congestion externality drives a wedge between equilibrium travel where demand intersects the average travel cost curve, and optimal travel where demand intersects the marginal travel cost curve. Figure C.1 illustrates this model for the purpose of motivating our valuation exercise, and Appendix C contains detailed derivations of the results below.

We first discuss how to value faster uncongested mobility. To fix ideas, consider first the case without any congestion externality. An uncongested speed improvement corresponds to a downward shift of a flat average cost curve. The product of total distance traveled and the decline in travel price (inverse speed) is a lower bound for the consumer surplus created, since additional gains from induced travel demand are ignored.²⁴ Panel A of figure C.1 illustrates this case.

Now consider the case with rising average cost. In this case, an increase in uncongested mobility generates a proportionally smaller peak hour speed improvement, because of congestion. Panel B of figure C.1 illustrates this. This crowding out effect is large when the average travel cost curve is steep or when the demand curve is flat. Our empirical results, however, suggest that this is not the case. In a bivariate cross-city regression, a 1% increase in uncongested mobility is associated with a 0.85% increase in peak-hour mobility, with a standard error of 0.05% and an R^2 of 0.54. In other words, $1 - 0.85 = 15\%$ of any given uncongested speed improvement is crowded out by congestion at peak travel demand.²⁵ Hence, ignoring all gains from induced demand, a 10% improvement in uncongested mobility (approximately one standard deviation) implies time saving gains on the order of 8.5% of the cost of travel. As an example, uncongested travel in Delhi is 15% faster than in Kolkota, and both cities experience considerable peak travel delay. However, peak travel speed remains 10% faster in Delhi than in Kolkota.²⁶

We next value the deadweight loss from congestion, shown in panel C of figure C.1 as the loss in consumer surplus from the wedge between the marginal and average cost of travel. (Not all congestion is a deadweight loss, so there can be considerable travel delay at peak hours even at a social optimum.) Our finding that peak demand congestion crowds out only 15% of uncongested mobility implies that the slope of the average cost curve is much flatter than that of the demand curve, and therefore that the deadweight loss triangle is small. A more precise statement requires an estimate of the price elasticity of travel demand, which

²⁴The same argument holds in discrete choice models like the one in Appendix B, which we can make more general by adding a nest for all non-travel consumption.

²⁵Both Akbar and Duranton (2018) and Kreindler (2018) find that the road network is modestly congestible. These results, amounting to a modestly increasing average travel cost curve, are consistent with our key empirical result here that differences in uncongested mobility are reflected in overall mobility.

²⁶In fact, keeping only the largest cities or weighting our regression by population makes our estimate larger than 1% but more imprecise, suggesting that even in the largest cities, uncongested speed advantages are mostly preserved at peak travel demand. Our 0.85% estimate is biased downward if cities with higher uncongested speed also experience higher travel demand. We do find a correlation of 0.2 between uncongested speed and peak travel congestion. We ignore this correlation in our computation, because it is small and because we prefer to obtain a conservative lower bound.

we lack in our context. In Appendix C, we show that assuming a unit elasticity at equilibrium implies a deadweight loss from congestion that represents about 1% of travel costs. This is similar to what Couture *et al.* (2018) find in the United States for a unit elasticity. Akbar and Duranton (2018) and Kreindler (2018) find a deadweight loss from congestion that is smaller than 1% of travel costs in two highly congested developing cities, Bogotá and Bangalore. As a comparison, our computations above suggests that a lower bound for the gains from a 10% improvement in uncongested mobility is on the order of 8.5% of travel time costs.

Our data show that congested networks are also more unreliable. We may therefore underestimate the gains from optimal congestion pricing if we ignore the fact that unreliable travel also contributes to a rising average cost curve. The most authoritative empirical evidence on how travelers value unreliability comes from Brownstone and Small (2005). They conclude that Californian morning commuters value the 90th to 50th percentile difference in travel time 95% to 140% as highly as they value median travel time. Here we use a value of unreliability equal to 100% of the value of travel time.

We can now evaluate the contribution of unreliability to the average cost of travel. In our data, the population-weighted average peak travel delay is 20%, and average unreliability is 8%. This suggests that the cost of unreliability is 40% of that of travel delay in equilibrium, so the average cost curve is 40% steeper when accounting for unreliability.²⁷ With this at hand, we can correct the bias in our welfare estimates above due to ignoring unreliability. Appendix C shows that the gains from a 10% uncongested speed improvement decline to 7.9% of travel cost, and the deadweight loss eliminated by optimal congestion pricing rises to 2.3% of travel costs.²⁸ Although taking unreliability into account raised our estimated gains from congestion pricing, these gains remain small. Clearly, our exact numbers here are tentative. However, as long as the optimal quantity of travel is close to the equilibrium quantity—which is what Kreindler (2018) and Akbar and Duranton (2018) also find—our conclusion that congestion pricing generates relatively small reductions in unreliability and travel delay is likely to hold. Overall, we conclude that the gains from achieving a one standard deviation improvement in uncongested mobility are many times larger than the gains from introducing optimal congestion pricing in urban India.

²⁷This assumes that unreliability rises linearly with the quantity of travel. This is conservative because unreliability rises somewhat less than proportionally with travel delay in our data.

²⁸We only consider improvements in reliability via congestion pricing. The gains from eliminating travel time unreliability (instead of only reducing it with congestion pricing) are potentially large, but we lack empirical evidence on the effectiveness of policies targeted at reducing unreliability, such as better handling of traffic incidents.

6. Correlation of mobility with city characteristics and urban development

We now explain mobility using city characteristics. We first consider basic characteristics like population and area. We then consider indicators of urban economic development, such as income levels, car ownership rates, and urban population growth. In addition, we consider road network measures that reflect urban development, such as the availability of primary roads and conformity to a regular grid pattern.

We report results for our benchmark mobility index in table 10. Table 11 panels A and B report the same specifications predicting the benchmark uncongested mobility and congestion indices, respectively. Because the mobility index is equal to the uncongested mobility index minus the congestion factor and we estimate the same specifications for all three dependent variables in each column, a given coefficient in table 10 is equal to the analogous coefficient in table 11 panel A minus the analogous coefficient in table 11 panel B.

In column 1 of table 10, we consider a simple specification with only log city population and log city area as explanatory variables. Because our dependent variable is a measure of log speed, we can interpret the coefficients as elasticities. For city population, we estimate an elasticity of -0.18. For city area, the elasticity is of opposite sign and equal to 0.15. These two variables explain more than half of the variation in mobility across Indian cities. Further controls added in subsequent columns change these results little. The robustness of these results is further confirmed in appendix tables D.2 and D.3 where we use alternative measures of mobility as dependent variables.

These results suggest a large “gross density” effect since an increase in population keeping land area constant is an increase in population density. An elasticity of the cost of travel with respect to population density of 0.16 to 0.18, as implied by the results of table 10, is comparable to the analogous elasticity of housing price at the center of French cities of 0.21 estimated by Combes, Duranton, and Gobillon (2019). To the extent that we can compare India with France, these results are suggestive that much of the urban costs measured by Combes *et al.* (2019) reflect worse mobility in denser cities.

At the same time, the mostly offsetting nature of the coefficients on population and urban land area in table 10 suggest that “net scale” effects are small, once we allow for land area to adjust to a larger population. Consistent with this, we estimate an elasticity of about -0.05 when regressing our preferred mobility index on log city population alone. Small net scale effects are comparable to analogous density elasticities for measures of urban productivity

such as wages (Combes and Gobillon, 2015).

In panels A and B of table 11, we estimate the same specifications as in table 10 using our preferred index of uncongested mobility and congestion factor as dependent variables. Consistent with our earlier decompositions of overall variance, we find that most of the effect of city population and city area on mobility works through uncongested mobility. For the congestion factor, we find an elasticity of city population of 0.02 in column 1. This coefficient remains between 0.02 and 0.03 in subsequent specifications. For the effect of city area on the congestion factor, we estimate small and insignificant elasticities in most specifications. Putting these results together, it appears that gross density mostly affects uncongested mobility, while the negative net scale effects are mostly about congestion.

Column 2 of tables 10 and 11 adds the log of primary roads length. Here and in subsequent specifications, we estimate a small but robust elasticity of mobility with respect to primary road kilometers of about 0.01. Other roadway measures are not robustly associated with mobility.²⁹ Interestingly, we find that the effect of primary roads on mobility mostly occurs through uncongested mobility while the effect of primary roads on the congestion factor is a precisely estimated zero. We think these findings reflect two facts. First, primary roads are intrinsically faster than secondary or tertiary roads. Second, the absence of an effect on the congestion factor is consistent with the fundamental law of ‘primary roads’ congestion: more primary roads attract new traffic and eventually leave congestion unchanged (Duranton and Turner, 2011).

Column 3 of table 10 further includes log city income and its square.³⁰ We find evidence of a hill shape where mobility first increases with income and then declines. The turning point corresponds to a city slightly below the top quartile of income. This finding is consistent with our rankings of the fastest and slowest cities in tables 5 and 6. Many of the fastest are middle-income cities, while the slowest are either among the poorest or richest cities in the country. When we examine the separate effects of income on uncongested mobility and the congestion factor in table 11 we find that the overall shape of the income-mobility relationship reflects two opposing forces. Uncongested mobility improves with income, perhaps because

²⁹Surprisingly, more motorways—high capacity dual carriage roads equivalent to freeways in the United States—are not robustly associated with improvement in mobility. We note that many Indian cities do not have any motorways in our sample.

³⁰Our income measure is log daily earnings for men. Since it is measured at the district level, it is subject to substantial measurement error. We exclude women due to lower labor force participation.

Table 10: Correlates of city mobility indices, benchmark mobility index

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log population	-0.18 ^a (0.016)	-0.18 ^a (0.015)	-0.17 ^a (0.016)	-0.17 ^a (0.017)	-0.17 ^a (0.016)	-0.17 ^a (0.018)	-0.17 ^a (0.016)	-0.16 ^a (0.017)
log area	0.15 ^a (0.016)	0.14 ^a (0.017)	0.13 ^a (0.017)	0.12 ^a (0.018)	0.12 ^a (0.017)	0.12 ^a (0.019)	0.12 ^a (0.017)	0.11 ^a (0.019)
log roads		0.013 ^a (0.0043)	0.012 ^a (0.0043)	0.013 ^a (0.0040)	0.011 ^b (0.0044)	0.014 ^a (0.0045)	0.013 ^a (0.0042)	0.014 ^a (0.0042)
log income			0.22 ^b (0.10)	0.23 ^b (0.10)	0.20 ^c (0.11)	0.23 ^b (0.11)	0.22 ^b (0.10)	
log ² income			-0.064 ^b (0.031)	-0.066 ^b (0.031)	-0.055 ^c (0.033)	-0.064 ^c (0.034)	-0.065 ^b (0.032)	
Network / shape				0.26 ^a (0.096)	0.11 ^b (0.049)	0.055 ^c (0.029)		
Pop. growth 90-10							0.052 ^c (0.030)	
share w. car								0.21 (0.14)
share w. motorcycle								0.11 ^b (0.054)
Observations	153	153	153	153	153	142	153	152
R-squared	0.54	0.56	0.57	0.59	0.58	0.60	0.58	0.59

Notes: OLS regressions with a constant in all columns. The dependent variable is the city fixed effect estimated in the specification reported in column 5 of table 4. Robust standard errors in parentheses. *a, b, c:* significant at 1%, 5%, 10%. Log population is constructed from town populations from the 2011 census. Log roads is log kilometers of primary roads within the city-light. Income is measured with male earnings from the 2011 census. The network / shape variable used in column 4 measures the share of edges in the road network that conform to the grid's main orientation, i.e., whose compass bearing are within 2 degrees of the modulo 90 modal bearing in the network. The network / shape variable in column 5 is a Gini index for the distribution of edge compass bearings in the road network. It also measures how grid-like the city is. The network / shape variable used in column 6 uses Harari's (2016) measure of the average distance between the centroid of the city and all the points that define its periphery. It measures the compactness of the city. The measure of population growth between 1990 and 2010 was constructed UN data. The share of households with access to a car or to a motorcycle is from the 2011 census.

of better roads. The congestion factor also increases with income, consistent with our findings on car ownership, which also rises with income. This second force appears to kick in at higher levels of income as evidenced by the fact that it is captured by the squared log income term in the regression. This is also consistent with our earlier findings that congestion is important in only a small number of cities.

In columns 4 and 5 of table 10, we consider two different measures of how well the road

Table 11: Correlates of city mobility indices, uncongested mobility and congestion factor

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Uncongested mobility								
log population	-0.16 ^a (0.017)	-0.15 ^a (0.016)	-0.15 ^a (0.015)	-0.14 ^a (0.016)	-0.14 ^a (0.015)	-0.15 ^a (0.017)	-0.14 ^a (0.014)	-0.13 ^a (0.017)
log area	0.17 ^a (0.017)	0.15 ^a (0.017)	0.13 ^a (0.018)	0.13 ^a (0.018)	0.13 ^a (0.017)	0.14 ^a (0.019)	0.13 ^a (0.017)	0.11 ^a (0.020)
log roads		0.013 ^a (0.0050)	0.013 ^a (0.0049)	0.014 ^a (0.0045)	0.012 ^b (0.0050)	0.015 ^a (0.0053)	0.015 ^a (0.0044)	0.017 ^a (0.0045)
log income			0.15 ^c (0.084)	0.15 ^c (0.083)	0.13 (0.089)	0.14 (0.093)	0.14 (0.086)	
log ² income			-0.026 (0.026)	-0.029 (0.025)	-0.022 (0.027)	-0.023 (0.028)	-0.029 (0.026)	
Network / shape				0.21 ^c (0.11)	0.047 (0.063)	0.020 (0.032)		
Pop. growth 90-10							0.11 ^a (0.030)	
share w. car								0.55 ^a (0.14)
share w. motorcycle								0.016 (0.051)
R-squared	0.43	0.45	0.48	0.50	0.49	0.50	0.53	0.53
Panel B: Congestion factor								
log population	0.024 ^b (0.0096)	0.024 ^b (0.0096)	0.026 ^a (0.0095)	0.025 ^a (0.0097)	0.024 ^b (0.0099)	0.023 ^b (0.010)	0.029 ^a (0.0095)	0.030 ^a (0.0098)
log area	0.018 ^b (0.0086)	0.017 ^c (0.0095)	0.0071 (0.010)	0.0084 (0.011)	0.0096 (0.011)	0.011 (0.011)	0.0044 (0.010)	0.0068 (0.011)
log roads		0.00068 (0.0028)	0.00094 (0.0029)	0.00068 (0.0029)	0.0017 (0.0029)	0.00081 (0.0031)	0.0017 (0.0027)	0.0028 (0.0028)
log income			-0.079 (0.060)	-0.080 (0.060)	-0.064 (0.061)	-0.091 (0.065)	-0.081 (0.057)	
log ² income			0.037 ^b (0.019)	0.038 ^b (0.019)	0.032 ^c (0.019)	0.041 ^b (0.020)	0.036 ^b (0.018)	
Network / shape				-0.046 (0.045)	-0.061 ^c (0.035)	-0.035 ^b (0.016)		
Pop. growth 90-10							0.054 ^a (0.018)	
share w. car								0.34 ^a (0.091)
share w. motorcycle								-0.096 ^b (0.042)
R-squared	0.54	0.54	0.60	0.60	0.61	0.61	0.63	0.63
Observations	153	153	153	153	153	142	153	152

Notes: OLS regressions with a constant in all columns. The dependent variable is the city fixed effect estimated for uncongested mobility in panel A, and the congestion factor in panel B. Robust standard errors in parentheses. *a, b, c*: significant at 1%, 5%, 10%. See the footnote of table 10 for further details about the explanatory variables.

network of a city conforms to a regular grid.³¹ Both measures suggest a positive association between a more grid-like pattern and better mobility in cities. The magnitude of the coefficients reported in the table for these measures is hard to interpret directly. A normalization indicates that a standard deviation in our grid variable is associated with 0.16 (in column 4) or 0.11 (in column 5) standard deviation in log mobility. This finding provides preliminary evidence in support of calls for more regular grid patterns for the roadway of emerging cities (Angel, 2008, Fuller and Romer, 2014).

We also experimented with the measures of urban form constructed by Harari (2016) and found a robust association between mobility and her measure of urban sprawl. The results are reported in column 6. That more sprawl is positively correlated with mobility is consistent with earlier results by Glaeser and Kahn (2004) for the us.

In column 7 of tables 10 and 11, we introduce a measure of past population growth. Cities that experienced faster population growth between 1990 and 2010 enjoy both faster uncongested mobility and more congestion. Overall the positive effect happening through uncongested mobility appears to dominate. While we leave a deeper investigation of these results for future research, we emphasize that they are inconsistent with typical claims that rapid urban population growth in developing countries is necessarily associated with worse mobility. Congestion may worsen with population growth but this negative effect is more than offset by faster roads.

Finally, in column 8 we no longer consider income but instead introduce two measures for the share of population with access to a car or equivalent, and (separately) a motorcycle. The insignificant positive coefficient for cars in explaining mobility in table 10 results from two offsetting effects where more cars are strongly and positively associated with both uncongested mobility and congestion in table 11. Motorcycles are associated with faster travel via less congestion, consistent with them taking up less room than cars, but inconsistent with them being a response to congestion. Again, causal identification is beyond our scope here but we would like to highlight that standard indicators of urban economic development such as higher incomes, faster population growth, and more cars are generally associated with better mobility outcomes despite higher congestion.

³¹The first measure captures the share of edges in the network that conform to the grid's main orientation i.e., whose compass bearing are within 2 degrees of the modulo 90 modal bearing in the network. The second measure is a Gini index for the distribution of edge compass bearings. Appendix A provides details. We also experimented with measures of the density of intersections and the length and circuitry of road segments but failed to uncover any robust association with our measures of mobility.

Although our findings above are generally stable across a wide variety of specifications, they may be subject to bias due to omitted city-level variables. In results reported in Appendix E, we control for city fixed effects, using within-city variation in population, area, and roads, at the level of concentric rings (0 to 2 kilometers from the center, 2 to 5, 5 to 10, 10 to 15, and 15 and beyond) to gain further insights about variation in mobility. Within cities, rings with more population and less urban area are slower, just as in the across-city results above.

7. Transit and walking

While roughly half the households in the average city in our data have access to a private vehicle – sometimes a car but more often a motorcycle – we recognize that city dwellers in India also often walk and use transit. To investigate these two alternative modes of travel, we also collected travel time data for walking and transit for all our trip instances.

For walking trips, speeds typically do not vary much across our trips and remain constant within trip. Mean walking speed is 4.8 kilometers/hour with a standard deviation of 0.1 kilometers/hour. We first estimate a city effect for walking trips in the same spirit as our baseline mobility index above. The standard deviation for the city effects is unsurprisingly tiny at 0.006. When we try to explain city effects for walking trip using the same approach as in table 10, the only robust correlate of our walking mobility index is a measure of average slope in the city. As Google Maps' algorithm reflects, steeper slopes slow down walking.

As described in Appendix A, we also collected transit data. These data have two important limitations. Google Maps only appears to return transit information for formal transit, and it bases its information on official timetables. This ignores informal transit and delays or missed services in formal transit. With these caveats in mind, we first note that only about 20% of our trip instances have a transit alternative that we define as 'viable': it requires less than an hour wait, and is strictly faster than walking. Despite this selection, viable transit trips take on average 2.3 times as long as trips with private vehicles. In regressions not reported here, we additionally find that, unsurprisingly, the transit time penalty is higher for shorter trips, trips further from the center, and nighttime trips.

Next, for 141 cities we can estimate an index analogous to our baseline mobility index for transit. Unlike with walking, there is a lot of cross-city variation for transit. The standard deviation for our transit mobility index is about twice that of our baseline mobility index for

private vehicles. This variation does not seem to be due to sampling problems as these indices are precisely estimated and alternative transit mobility indices are all highly correlated.

The correlation between our mobility index for transit and our baseline mobility index (for private vehicles) is extremely low at 0.02. This correlation even becomes negative when we focus on the largest cities. However, when we re-estimate our mobility index for private vehicles on a sample limited to trip instances for which a viable transit alternative is possible, this correlation increases from 0.02 to 0.17. This difference suggests a fair amount of selection regarding which trip instances have a viable transit alternative known to Google. To confirm the low correlation between transit and vehicle travel times we regress log transit travel time on log private vehicle travel time and log walking time. In this regression, the coefficient on log vehicle travel time is only 0.19 while the coefficient on log walking time, which is essentially a measure of trip length, is 0.52.

Finally, we also replicated the regressions of table 10 for our transit mobility index. We did not find any robust correlates of transit mobility at the city level. Given the sizable variation across cities in transit mobility, this may seem surprising. Nonetheless, this result is consistent with the weak correlation between (private vehicle) mobility and transit mobility. Although we must remain cautious given the caveats that apply to our transit data, taken together these results suggest to us that transit mobility depends much more on the coverage and frequency of transit than on driving speeds. This conclusion for the cross-section of cities is consistent with figure A.4 in Appendix A, which shows that transit mobility is highest at the slowest hours for vehicular mobility.

8. Conclusions and policy implications

We propose a novel approach to measuring vehicular mobility within cities, and decomposing it into uncongested mobility and a congestion factor. We apply it using novel large scale data on counterfactual trips in 154 Indian cities collected from Google Maps. After showing that various sampling and estimation strategies yield similar estimates of mobility, we document a number of important facts about mobility in Indian cities. Among the most important, we first highlight large mobility differences across cities. Second, slow mobility is primarily due to cities being slow all the time rather than congested at peak hours. We do nonetheless find an important role for congestion in the largest cities, especially close to their centers. Third, several city attributes are consistently correlated with mobility and its

components. We find that population and land area are key correlates of city mobility. Higher population density is strongly associated with slower uncongested mobility as well as more congestion. We also find that both recent population growth and a measure of cars per capita are positively associated with uncongested mobility but also with congestion. More primary roads and a more regular grid-pattern are associated with moderately faster mobility. Higher income cities have higher uncongested mobility, but also higher congestion, leading to a hill-shaped relationship between income and overall mobility. Overall, these indicators of urban economic development are associated with better mobility despite worse congestion, contrary to a conventional wisdom that urban growth and development condemns developing cities to complete gridlock. While in principle, variation in uncongested mobility could be due to many city attributes beyond those we consider here in our regressions, such as the state of the vehicle stock or driving culture, we interpret it as being primarily due to the quality of the road network. Most old cars can be driven 45 kilometers per hour (the 99th percentile of our trip speed distribution), and Google Maps' algorithm is likely to pick out a high percentile of the block speed distribution it observes in order to distinguish motorized from non-motorized vehicles.

We hope that this first set of cross-city evidence on urban mobility and congestion in a developing country can help guide policy and future research. We now review three of our findings that have research and policy implications. First, we document that congestion in India is not a nationwide problem, but rather is highly concentrated near the center of the largest Indian cities. Given their importance to the Indian economy, these areas with the highest levels of congestion, such as the center of Bangalore, should be the focus of policy effort to alleviate congestion, and of future research to identify the most effective policies.

Second, we compared travel patterns in India with those from more developed cities, and we uncovered important differences. The fastest Indian cities are slower than the slowest American cities, and in general Indian cities do not experience the familiar twin peak congestion patterns due to morning and evening commutes. There is a muted morning peak, and instead a slow buildup of congestion that often persists until late into the evening. Light rainfall appears to speed up traffic slightly. These unique patterns are consistent with Indian roads being multi-purpose public goods serving a wide variety of uses other than motorized transport that slow down travel. If this conjecture is correct, then further research on technologies and policies for separating roadway uses appears especially promising, with appropriate consideration for the costs of restricting non-vehicle uses. More generally, our

findings of unique Indian travel patterns imply that country-specific policies are necessary, and that using our data sources and methodology to study other countries individually may uncover distinctive patterns.

Third, our most surprising and perhaps controversial finding is that in most Indian cities travel is slow at all times, not just peak times. Related to this, a simple welfare analysis suggests sizable potential gains from improving uncongested speed in urban India, and comparatively small gains from eliminating the deadweight loss from congestion. As a result, standard policy recommendations like congestion pricing, HOV lanes, or other types of travel restrictions may do little to improve mobility. Recent empirical work on congestion in two developing cities by Akbar and Duranton (2018) and Kreindler (2018) reinforce this message, as they also conclude that the deadweight loss from congestion is small. We would therefore like to encourage researchers to also study policies and investments that generate faster uncongested speed. Our paper provides a first set of results suggesting a modest positive role for the design of a regular network grid and the presence of more primary roads, but much work remains to be done in terms of identifying cost-effective ways to build faster urban networks. On an optimistic note, we find that better uncongested mobility generally correlates with the process of economic development.

We believe a lot more can be learned from the data we use here. In an extension of this paper (Akbar, Couture, Duranton, and Storeygard, 2018), we provide complementary measures of urban accessibility in Indian cities, decompose accessibility into proximity to destinations and mobility, and provide an analysis of the urban correlates of accessibility and proximity. This sort of data can thus be used to learn about the fundamentals of urban travel beyond mobility and congestion. It can also potentially play an important role in our understanding of patterns of land use and property prices in cities in relation to transportation. Relative to more traditional travel surveys, the information used here is less complete but can be gathered at a small fraction of the cost, hundreds of dollars instead of tens of millions for a full travel survey. The type of data we used here is also much more versatile and can thus be targeted at narrower issues or areas without fear of losing statistical power. It can also be collected at much higher frequency than the typical 5 to 8 year gap between consecutive traditional travel surveys, and allows for the evaluation of policy changes in the short-run (Kreindler, 2016, Hanna *et al.*, 2017). We believe future studies of this type will shed useful light on many aspects of transportation policy in cities. Many other possible applications are possible. They include, for instance, the monitoring of city recovery after major natural

disasters.

We also hope that more data underlying the production of real-time travel information will be made available for research. The data that we use allow us to learn about mobility, and the price (time cost) of travel for all possible trips at all times. The analogous quantities (i.e., number of travelers) are potentially knowable from the same underlying data. With both prices and quantities, the detailed study of congestion, both on particular road segments and in larger areas, will be possible. Repeated observations of the same travelers would also enable a much better analysis of individual travel behavior. For instance, Kreindler (2018) uses a panel of trip-level data for 2,000 commuters from a smartphone app to learn about individual response to peak travel congestion, and to measure the welfare impact of various pricing policies to alleviate congestion in Bangalore. With appropriate regard for privacy, the availability of larger trip-level samples across cities would allow for a comprehensive analysis of the welfare consequences of better urban mobility and accessibility.

References

- 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. 2018. Accessibility in urban India. Work in progress, University of Pennsylvania.
- Akbar, Prottoy A. and Gilles Duranton. 2018. Measuring congestion in a highly congested city: Bogotá. Processed, University of Pennsylvania.
- Alonso, William. 1964. *Location and Land Use; Toward a General Theory of Land Rent*. Cambridge, MA: Harvard University Press.
- Anderson, Simon P., André de Palma, and Jacques-Francois Thisse. 1992. *Discrete Choice Theory of Product Differentiation*. Cambridge, ma: MIT Press.
- Angel, Shlomo. 2008. An arterial grid of dirt roads. *Cities* 25(3):146–162.
- Annez, Patricia Clarke, Alain Bertaud, Marie-Agnes Bertaud, Bijal Bhatt, Chirayu Bhatt, Bimal Patel, and Vidyadhar Phata. 2016. Ahmedabad. More but different government for “slum free” and livable cities. World Bank Policy Research Working Paper 6267.
- Atkin, David, Benjamin Faber, and Marco Gonzalez-Navarro. 2018. Retail globalization and household welfare: Evidence from Mexico. *Journal of Political Economy* 126(1):1–73.
- Ben-Akiva, Moshe E. and Steven R. Lerman. 1985. *Discrete Choice Analysis: Theory and Application to Travel Demand*. Cambridge, MA: MIT Press.
- Boston Consulting Group. 2018. Unlocking Cities. The impact of ridesharing across India. Report.
- 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.
- Ch, Rafael, Diego Martin, and Juan F. Vargas. 2018. Measuring cities with night-time light data. Processed, CAF - Development Bank of Latin America.
- Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon. 2019. The costs of agglomeration: House and land prices in French cities. *Review of Economic Studies* forthcoming.
- Combes, Pierre-Philippe and Laurent Gobillon. 2015. The empirics of agglomeration economies. In Gilles Duranton, Vernon Henderson, and William Strange (eds.) *Handbook of Regional and Urban Economics*, volume 5A. Amsterdam: Elsevier, 247–348.
- Couture, Victor, Gilles Duranton, and Matthew A. Turner. 2018. Speed. *Review of Economics and Statistics* Forthcoming.
- Duranton, Gilles and Diego Puga. 2015. Urban land use. In Gilles Duranton, J. Vernon Henderson, and William C. Strange (eds.) *Handbook of Regional and Urban Economics*, volume 5A. Amsterdam: North-Holland, 467–560.

- Duranton, Gilles and Matthew A. Turner. 2011. The fundamental law of road congestion: Evidence from US cities. *American Economic Review* 101(6):2616–2652.
- Fuller, Brandon and Paul Romer. 2014. Urbanization as opportunity. World Bank Policy Research Paper 6874.
- Geroliminis, Nikolas and Carlos F. Daganzo. 2008. Existence of urban-scale macroscopic fundamental diagrams: Some experimental findings. *Transportation Research Part B: Methodological* 42(9):759–770.
- Glaeser, Edward L. and Matthew E. Kahn. 2004. Sprawl and urban growth. In Vernon Henderson and Jacques-François Thisse (eds.) *Handbook of Regional and Urban Economics*, volume 4. Amsterdam: North-Holland, 2481–2527.
- Hanna, Rema, Gabriel Kreindler, and Benjamin A. Olken. 2017. Citywide effects of high-occupancy vehicle restrictions: Evidence from “three-in-one” in Jakarta. *Science* 357(6346):89–93.
- Harari, Mariaflavia. 2016. Cities in bad shape: Urban geometry in India. Processed, Wharton School of the University of Pennsylvania.
- Kreindler, Gabriel. 2016. Driving Delhi? Behavioural responses to driving restrictions. Processed, MIT.
- Kreindler, Gabriel. 2018. The welfare effect of road congestion pricing: Experimental evidence and equilibrium implications. Processed, MIT.
- Mangrum, Daniel and Alejandro Molnar. 2017. The marginal congestion of a taxi in New York City. Processed, Vamderbilt University.
- Mills, Edwin S. 1967. An aggregative model of resource allocation in a metropolitan area. *American Economic Review (Papers and Proceedings)* 57(2):197–210.
- Muth, Richard F. 1969. *Cities and Housing*. Chicago: University of Chicago Press.
- Sheu, Gloria. 2014. Price, quality, and variety: Measuring the gains from trade in differentiated products. *American Economic Journal: Applied Economics* 6(4):66–89.
- Small, Kenneth A. and Erik T. Verhoef. 2007. *The Economics of Urban Transportation*. New York (NY): Routledge.
- United Nations. 2015. *World Urbanization Prospects: The 2014 Revision*. New York (NY).

Appendix A. Further data description

City sample and extent

United Nations (2015) reports the population and location of 166 cities in India that reached a population of 300,000 by 2014. Following Harari (2016) and Ch *et al.* (2018), we define the spatial extent of these cities as sets of contiguous 30 arc-second pixels with a lights-at-night digital number (DN) of at least 35 whose boundaries reach within 3 kilometers of the UN's reported latitude and longitude. The lights data are the stable lights product from the F-18 satellite. They are available at <https://ngdc.noaa.gov/eog/dmsp/downloadV4composites.html> (last accessed 6 September 2018.) The UN database initially reported an incorrect location for one city (Bokaro Steel City); it has since been corrected. We resampled Bokaro Steel City in December 2017 once we discovered this problem.


We drop two cities (Cherthala and Malappuram) that are not within 3 kilometers of a $DN > 35$ light, one (Santipur) that belongs to a lit area with exactly one $DN > 35$ pixel, which is an implausibly small extent, and five cities that are too far east to be in the land use dataset described below (Agartala, Aizawl, Guwahati, Imphal, and Shillong). Four city-lights contain two cities each: Raipur and Durg-Bhilainagar, Mumbai and Bhiwandi, Asansol and Durgapur, and Bangalore and Hosur. We treat each of these four pairs as an individual city, with the center of the larger member of each pair kept as the center of the combined city. Our primary sample thus includes 154 cities.

We further restrict city boundaries for the purpose of defining trip origins and destinations by excluding water bodies and non-urban land using 40-meter resolution land cover classifications from the Global Human Settlement Layer (GHSL) of the European Commission's Joint Research Centre (JRC). Cells identified as at least partially built up or roads within a city light are retained. Panel A of figure A.1 shows the lit and built-up portions on a median-sized city, Jamnagar in Gujarat, which we use for illustrative purpose throughout this appendix.


Trip sample

This section describes how we determine the within-city trips to query on Google Maps. We define a *trip* as a pair of points (origin and destination) within the same city as defined above. A *trip instance* is a trip taken at a specific time on a specific day. A *location/point* refers to a pair of longitude-latitude coordinates identifying the centroid of a roughly 40-meter GHSL pixel. We require that trip location pairs are at least one kilometer apart in haversine length, for three reasons. First, the rounding of travel times and lengths introduce potentially non-classical measurement error in our computations of travel speed. Second, Google does not


Figure A.1: Illustrations for the city of Jamnagar




Panel A: Built-up categories within lit area (higher numbers reflect greater built-up intensity)




Panel B: Radial trips of absolute lengths 2 km, 5 km, 10 km, and 15 km from the center




Panel C: Radial trips over uniformly picked distance percentiles



Panel D: Circumferential trips around the center

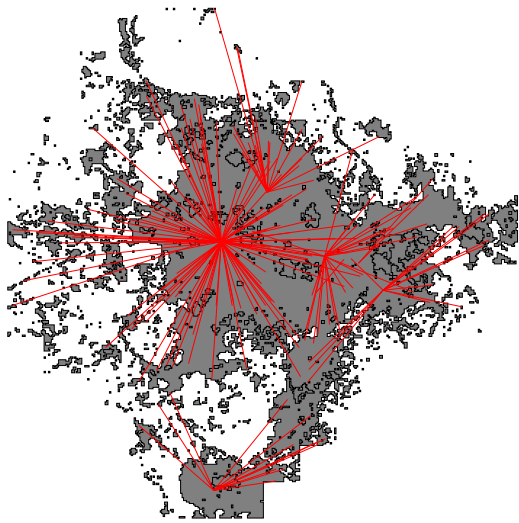


Panel E: Gravity trips




Panel F: Trips to school

Figure A.1 (continued): Illustrations for the city of Jamnagar



Panel G: Trips to shopping malls



Panel H: Trips to hospitals

always return a driving time under traffic conditions for very short trips. Even when it does, the travel times can sometimes be very inconsistent or require taking unnecessary detours. Third, walking is an easy alternative to driving for short trips, and sources of error such as the unobserved time cost of finding parking, etc. will be a more significant component of the trip.

Our target sample for city c is $15 \sqrt{Pop_c}$ trips, where Pop_c is the projected 2015 population of city c from United Nations (2015), and 10 trip instances per trip, to ensure variation across times of day. That is approximately 82,000 trip instances for the smallest of our cities, 116,000 instances for a median-sized city, and 760,000 instances for the largest city (Delhi).³²

We define four types of trips: radial (2/9 of all trips), circumferential (1/9), gravity (1/3), and amenity (1/3).

Radial trips

Radial trips are defined in a polar coordinate system with respect to a city center. They have one end at a randomly located point within 1.5 kilometers of the city centroid as defined by United Nations (2015). Distance from the centroid is drawn from a truncated normal distribution with mean 0, standard deviation 0.75 kilometer and support $[0,1.5]$ kilometers. For convenience, we call this the destination, but in practice trips in both directions are sampled. For each destination, the point of origin is determined using two methods with equal probability:

³²By comparison, in the 2008 US National Household Transportation Survey (NHTS), the 187th, 100th, 50th, 10th and 1st most sampled US metro areas have about 200, 800, 2,200, 12,000, and 29,000 trips, respectively.

1. *Absolute distances* of $AbsDist \in \{2,5,10,15\}$ kilometers (equally weighted) are drawn. For each of these four distances, we (uniform) randomly pick a point of origin within the lit-up area of the city that is between $(AbsDist - 0.2)$ kilometers and $(AbsDist + 0.2)$ kilometers from the given destination. See panel B of figure A.1 for illustration with the city of Jamnagar. Darker shades of red distinguish longer trips.
2. *Distance percentiles* relative to the largest possible distance for any trip from a lit-up area of the city to that destination are drawn from a uniform distribution from the 1st to 99th percentile (excluding distances less than 1 kilometer). See panel C of figure A.1 for illustration with the city of Jamnagar.

If a city has no valid trips for a given absolute distance ± 0.2 kilometer, the trips assigned to that distance are reallocated to the distance percentiles sample.³³ Similarly, if there are not enough unique 40 m pixel centroids ± 0.2 kilometer from the center destination to fill a given absolute distance's quota, the remainder of the quota is filled with randomly drawn distance percentiles instead.

Circumferential trips

Like radial trips, circumferential trips are also defined in a polar coordinate system with respect to a city center. Circumferential trips originate at a random origin at least 2 kilometers away from the city centroid. The analogous destination is at the same distance (± 0.2 kilometer) from the centroid, 30 (± 3) degrees clockwise or counter-clockwise from the origin. For three small cities, the city centroid according to United Nations (2015) is far from the geographic center of the city-light, so it was not possible to fill the circumferential trip quota. See panel C of figure A.1 for illustration with the city of Jamnagar.

Gravity trips

Gravity trips are designed to match the length profile of trips sampled in the US NHTS and the Bogotá Travel Survey. We identified each location-pair using the following algorithm:

1. Consider a uniformly randomly picked initial point (*GravityPoint*) and a length (*GravityLength* kilometers) drawn from a truncated Pareto distribution with shape parameter 1 and with support between 1 kilometer and 250 kilometers (corresponding to a mean of roughly 5.52 kilometers).³⁴

³³Only 43 cities have a maximum distance to centroid of 15 kilometers or more. 78, or roughly half, of the cities have a maximum distance of 10 kilometers or more. 132 cities have a maximum distance of 5 kilometers or more, and all cities have a maximum distance greater than 2 kilometers (with the smallest maximum distance being 2.8 kilometers).

³⁴This mean of 5.52 kilometers is slightly smaller than the mean of 6.51 kilometers for Bogotá from Akbar and Duranton (2018).

2. Choose a point randomly from among all points at a straight-line length between $(GravityLength - 0.2)$ kilometers and $(GravityLength + 0.2)$ kilometer from the point $GravityPoint$. If there are no such points, start over from (1) with a new pair of $(GravityPoint, GravityLength)$.

See panel E of figure A.1 for illustration with the city of Jamnagar. Darker shades of red distinguish longer trips.

Amenity trips

Amenity trips join a random origin with an instance of one of 17 amenities (e.g. shopping malls, schools, train stations) as recorded in Google Places. The particular instance we used is based on a combination of proximity and “prominence” assigned by Google. The weighting across these amenity types is based on a mapping of amenities to trip purposes for the 100 largest MSA in the US from the 2008 US National Household Transportation Survey (NHTS) (Couture *et al.*, 2018). NHTS has nine categories of trip purpose (trip share in parentheses): Work (23.6%), Work-related business (3.3%), Shopping (21.8%), School & Religious practice (4.6%), Medical/dental (2.2%), Vacation & visiting friends/relatives (6.0%), Other social/recreational (13.8%), Other family/personal business (24.3%), and Other (0.5%).

The Google Places API classifies points of interest using one or more of roughly 100 Google-defined place “types”. We match each NHTS trip purpose to the most relevant Google Places types, using city hall for Work, under the assumption that employment is relatively concentrated near the city center. Since we cannot identify types associated with Other family/personal business, we reallocated its 24.3% share among the rest of the categories except Work using the following formula. If place type v gets $TripTypeShare_v\%$ of the trips otherwise, then they get an additional $\frac{24.3(23.6 - TripTypeShare_v)}{\sum_w(23.6 - TripTypeShare_w)}$. Less popular place types get a larger share of Other family/personal business as we do not want too few absolute trips in any category. The final allocation is shown below. The first number in each category is its initial allocation, and the second is its share of Other family/personal business.

- Work: city hall (23.6%+0%)
- Work-related business: gas station (3.3%+1.5%)
- Shopping: shopping mall (7.3%+1.2%), convenience store (7.3%+1.2%), grocery/ supermarket (7.2%+1.2%)
- Social/recreational: movie theater (5.7%+1.3%), park (5.7%+1.3%), stadium (2.4%+1.5%)
- School & religious practice: school (2.3%+1.6%), place of worship (2.3%+1.6%)
- Medical/dental: hospital (1.1%+1.7%), doctor (1.1%+1.7%)

- Vacation & visits: train station (3.0%+1.5%), airport (1.0%+1.7%), bus station (2.0%+1.6%)
- Other: police (0.25%+1.75%), post office (0.25%+1.75%)

We set a different maximum radius of the search around any initial point based on the place type:

- 50 kilometers radius: city hall, airport, stadium
- 20 kilometers radius: train station, bus station, hospital, doctor
- 10 kilometers radius: movie theater, school, police
- 5 kilometers radius: shopping mall, convenience store, grocery/supermarket, park, place of worship, gas station, post office

A query request to Google Places API specifies a search location and a 'type'. For each query, we randomly draw (without replacement) a new location within our city's lit-up boundary. We call a query to the API successful if it returns at least one place. For a given city, if a query by 'type' is unsuccessful more often than not after at least 50 unsuccessful queries, we switch to querying by 'keyword', which is more likely to return results but also more likely to include badly matched returns, e.g. return coordinates for some segment of a road named "Airport Road" instead of coordinates for the airport. If queries by keyword also continue to be unsuccessful more often than not, after 50 unsuccessful queries we reallocate the remaining share of the location pairs evenly among the rest of the place types under the same trip purpose category. For example, suppose we require 100 location pairs for 'convenience stores' and the first 50 queries by type return zero results. So we switch to querying by keyword. Suppose, the 80th query by keyword is the 50th unsuccessful one. Then we stop there, get 30 location pairs from the successful queries for 'convenience stores' and reallocate the remaining 70 required location pairs to 'shopping mall' and 'grocery/supermarket' (35 each). If all place types in the same trip purpose category yield zero place returns more often than not and we have yet to fulfil our quota of location pairs in the category, then we re-distribute the count of unqueried location pairs evenly across all the rest of the place types.

From each successful query, we collect only the first twenty places returned by Google in order of "prominence", as determined "by a place's ranking in Google's index, global popularity, and other factors". For each place, Google's Places API returns us: geographical coordinates, "name", "vicinity" (this might be either an address or nearby landmarks), and the "types" it is classified under. We only keep places that are at least one kilometer in straight-line distance from the random initial point. Then we use the "name", "vicinity" and "types" of the place to score the relevance/quality of each place return. We drop places below a minimum threshold (i.e. more likely to be a bad match), and use the highest scoring place, breaking ties first with length differentials over one kilometer (i.e. keeping the closest),

and then by "prominence" (i.e., the order in which they are reported by Google). This ensures that small differences in length are ignored in favor of Google's recommendation.

Since not all successful queries return good quality places, we make 50% more queries than needed. When choosing the final set of trips to query for traffic, we prioritise trips to places that scored highly on relevance. If we need to break ties here, we pick randomly. Panels F, G, and H of figure A.1 illustrate for the city of Jamnagar our selection of trips to schools, shopping malls, and hospitals, respectively.

Querying trips on Google Maps

Our target sample was 2,373,764 trips across all cities and strategies, corresponding to 1,186,882 locations pairs. Because of some overlaps between trips and because Google Maps did not return any route for few hundred trips, we ended up with 2,333,762 queried trips, or 98.3% of our target. Across cities, the mean is 98.7% with a coefficient of variation of 1.34%.

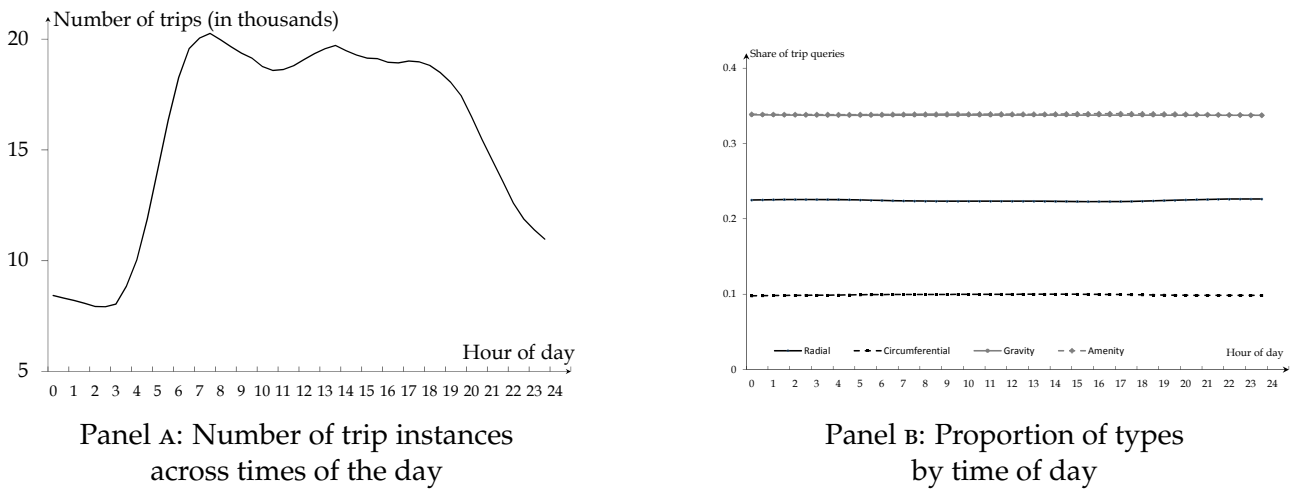
We simulated 22,766,881 trip instances across 40 days between September and November of 2016. This corresponds to 92.5% of our target on average in Indian cities with a coefficient of variation of 4.06%. The median (as well as the mean) trip was queried 10 times (with a standard deviation of 1.9) and 99% of the trips were queried at least 8 times. Missing trip instances are due mostly to empty returns from Google Maps or minor technical glitches such as early computer disconnections, formatting problems in the returns, etc.

We wanted the distribution of trip departure/query times to roughly resemble the distribution of departure times on a typical weekday.³⁵ However, we also wanted enough trip queries from each time period of the day for the fixed effects to be credible, so we oversampled the early morning. At any hour of the day, we had the following number of machines querying trips on Google: 12 a.m. - 4 a.m.: 15, 4 a.m. - 5 a.m.: 20, 5 a.m. - 6 a.m.: 35, 6 a.m. - 8 a.m.: 40, 8 a.m. - 12 p.m.: 35, 12 p.m. - 1 p.m.: 40, 1 p.m. - 5 p.m.: 35, 5 p.m. - 7 p.m.: 40, 7 p.m. - 9 p.m.: 30, 9 p.m. - 10 p.m.: 25, 10 p.m. - 12 a.m.: 20. All the machines had identical processing power, so the number of machines also reflects the distribution of our trip queries across hours of the day. Panel A of figure A.2 shows the realized distribution of query times across hours of the day.

We wanted to have an even spread of days and times across cities and trip types/strategies. So the order in which the trips were queried was randomized to alternate between strategies and cities (based on the size of the city, e.g. city A - with twice as many trips as city B - is queried twice between every city B query). Once we have run through the ordered list of trips, we start over at the beginning of the list. Panel B of figure A.2 shows the stable realized proportion of trip types across hours of the day.

³⁵We rely on a household transportation survey from Bogota, Colombia as a reference for this.

Figure A.2: Queries by time of the day



As the ordering of trips stays the same, one may worry that if the time it takes to cycle through the list is roughly a multiple of 24 hours, there will be too little variation in time of day across instances of the same trip. So we split the day into four 6-hour time slots (12 a.m. - 6 a.m., 6 a.m. - 12 p.m., 12 p.m. - 6 p.m., 6 p.m. - 12 a.m) and forced randomization within each of them by maintaining a separate trip query list for each slot. That means, at the end of each 6 hour slot we bookmarked our location on the query list and came back to it in 18 hours. This makes sure that no trip is randomly over- or under-queried at any given 6-hour slot of day. We managed to make sure that 95% of the trips were queried at all four 6-hour time slots, and every trip was queried at, at least, three of the four slots.


We sampled weekends at 50% of our weekday rate, using the same method. While we might prefer to oversample “Other family/personal business” trips on weekends, as discussed above we cannot narrow down the set of destinations for this category.

Travel lengths and speeds

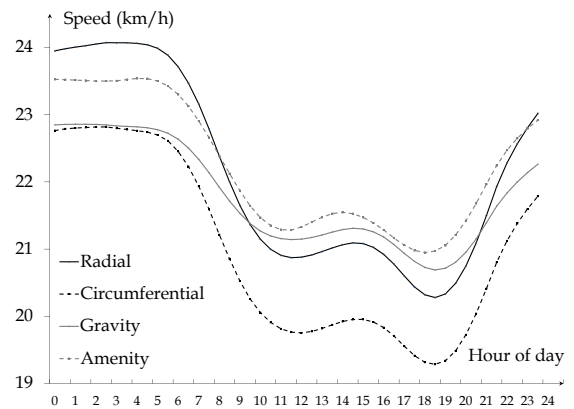
The median Google-reported travel length across all our trips is 5 kilometers (with a standard deviation of 10.5 kilometers). However, there are noticeable differences across our four trip selection strategies. Figure A.3 shows the distribution of travel lengths for the portfolio of trips under each strategy. Amenity trips are relatively shorter in length, with a median of 4.2 kilometers. This is understandable as our algorithm weakly prefers closer destinations for any given amenity. Radial trips are the longest, with a median of 6.6 kilometers. This is probably because we force a large share of the trips to be of fixed haversine lengths of 5, 10 and 15 kilometers, which translate to even larger actual travel lengths.³⁶ Recall that the

³⁶In fact, the ratio of total travel length to total haversine length is 1.53.

Figure A.3: Travel length and speed



Panel A: Distribution of travel lengths by trip class



Panel B: Travel speeds across time of day by trip class

gravity trips are designed to mirror the distribution of travel lengths that have been observed in other cities.

Panel B of figure A.3 shows how travel speeds through the day vary across our trip selection strategies. As we would expect, speeds are highest in the early hours of the morning and late at night and lowest during the day, in particular around the 6 - 7 p.m. evening rush hour. Some of the differences in speeds across strategies may be explained by the differences in trip lengths, as longer trips also tend to be faster. But, clearly there is more to it: circumferential trips experience the lowest speeds, and speeds for the radial and circumferential trips seem relatively more sensitive to daytime increases in traffic.

Walking and transit trips

We do not expect walking times for a given trip to vary by either the day or the hour of day. However, walking speeds do vary based on slope and the density of the network of streets and pedestrian paths. So, unlike for driving times, we query each location pair only once, in one direction, for walking times.

Google does not generally track transit in real time, but instead relies on public transportation schedules made available by transit authorities and open General Transit Feed Specification data. Thus, for any given trip, we do not expect any meaningful variation across weekdays in our travel times by transit. Scheduled transit frequency does however vary by time of day. We thus re-queried each weekday trip instance in our driving data as a transit trip, at its original time of day, but on 10 January 2018. This was a Wednesday that did not coincide with any public holidays in India to our knowledge.

Table A.1: Ranking of cities by transit network coverage

Rank	City	State	Coverage
1	Chennai	Tamil Nadu	0.74
2	Bangalore	Karnataka	0.73
3	Pune	Maharashtra	0.73
4	Mysore	Karnataka	0.69
5	Mumbai	Maharashtra	0.67
6	Ahmedabad	Gujarat	0.65
7	Chandigarh	Chandigarh	0.63
8	Rajkot	Gujarat	0.62
9	Kolkata	West Bengal	0.61
10	Jaipur	Rajasthan	0.61

Notes: Coverage refers to the share of trip instances with viable transit routes returned by Google Maps.

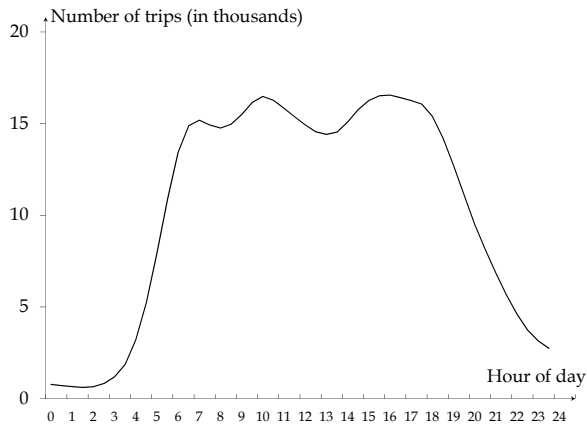
There are several important caveats to these data. First, 22% of queries, including all queries in 14 cities, returned no routes. Second, we do not expect the schedules to include informal transit providers, which own the large majority of India’s bus fleet.³⁷ Third, some returned routes are implausible. Specifically, we exclude routes that (1) require walking all the way, (2) require waiting over an hour to start the trip, or (3) are slower than their walking counterpart, which happens when Google uses inter-city rail, presumably because it is the only nearby transit alternative, to create highly convoluted itineraries. Following these exclusions, only 20% of our driving trip instances offer viable transit alternatives, and they are highly concentrated in the largest cities. In 133 of our 154 cities, less than 8% of trips are viable by transit. We cannot distinguish whether the absence of a viable transit route is due to limitations in the city’s transit network or limitations in Google Maps’ coverage of the transit network. With that in mind, we report the 10 cities with the largest share of our trip instances covered by Google Maps in Table A.1. Figure A.4 shows, for each hour of the day, the number of transit query in Panel A, and the average speed by trip class in Panel B.

Road network data


Our measures of road network characteristics come from OpenStreetMap (OSM), a collaborative worldwide mapping project. We downloaded OSM data within the light-based boundary of each city through Geofabrik in September 2016 (<http://download.geofabrik.de/asia/india.html>, last accessed, 6 September 2018). We then used *osmnx*, a python package created

³⁷See <https://data.gov.in/catalog/number-buses-owned-public-and-private-sectors-india>, last accessed, 6 September 2018. Note also that Google Maps only officially lists transit authorities spanning 12 Indian cities, corresponding to 10 of our cities, and four multi-region services that share their transit schedules (<http://maps.google.com/landing/transit/cities/>, last accessed 6 September 2018), but queries in an additional 130 cities returned transit components.

Figure A.4: Transit data



Panel A: Distribution of query times across times of the day



Panel B: Travel speeds across time of day by trip class

by Geoff Boeing, to process the OpenStreetMap network as a directed graph of edges and nodes.

Road length

Each edge in the OSM network receives a tag which characterizes its road type. We measure total road length in kilometers for three types of roads:

1. Motorways: The highest capacity roads in a country, equivalent to freeways in the United States. Motorways generally consist of restricted access dual carriage ways with 2 or more lanes in each direction plus emergency hard shoulder.³⁸
2. Primary Roads: The next most important road in a country's transportation system, after motorways and trunks. Generally not dual carriage ways.
3. Total Road Length: aggregation of all road types driveable by motor vehicles and public for everyone to use.³⁹

We note that certain cities have incomplete street networks on OSM. Using satellite data, we visually identified a set of cities for which the road network appears incomplete (Jhansi, on the left-hand panel of Figure A.5, is one such city.) The results are robust to limiting the sample to the subset of cities for which we have a more complete road network.

³⁸We also include the less frequent OSM type "trunks" in the motorways category. Trunks are the next most important types of roads after motorway, and often but not always consist of dual carriage ways.

³⁹In the OSM network, both carriage ways of a motorway count as separate edges (in each direction). We experimented with counting dual carriage ways only once when measuring length, and also with measuring lane-kilometers, instead of just edge kilometers. These adjustments generate measures of length by road type that are very highly correlated with that without adjustments that we show in the paper.

Characterizing the road network

osmnx calculates the compass bearing ("bearing" for short) from each directed edge's origin node to its destination node. The bearing captures the orientation of the edge with respect to true north. We use the distribution of edge bearings in a city to characterize how 'grid-like' its road network is. We measure how grid-like a network is in two separate ways: 'orientation' which captures the share of edges conforming to the network's main grid orientation, and 'Gini' which captures the dispersion in the distribution of edge bearings. We now describe both measures of how grid-like a road network is in more detail.

Orientation. A grid is a series of roads intersecting at perpendicular angles. If a city were a perfect grid network, then all bearings for would be either perpendicular or parallel to each other. The orientation grid metric measures the proportion of edges in a city's road network that conform to the dominant grid orientation in that they are perpendicular or parallel to the modal edge bearing.

Let g index each edge in the road network of city c , and let x_{cg} be the edge bearing rounded to the nearest degree, and x_c^{modal} be the modal edge bearing modulo 90 of city c . For example, if a city's grid were oriented N-E-S-W, then x_c^{modal} would equal 0. Let $\delta_{gc, x_c^{modal}, \nu}$ be an indicator for whether edge g in city c conforms to grid orientation x_c^{modal} within a bandwidth error of ν :

$$\delta_{gc, x_c^{modal}, \nu} = \begin{cases} 1 & \text{if } (x_g - x_{0ct}) \bmod 90 \leq \nu \\ 1 & \text{if } (x_g - x_{0ct}) \bmod 90 \geq (90 - \nu) \\ 0 & \text{else.} \end{cases} \quad (\text{A1})$$

We then compute our grid-like measure as:

$$\text{Orientation}_c = \frac{\sum_{g \in I_c} \delta_{gc, x_c^{modal}, \nu}}{Q_c}, \quad (\text{A2})$$

where I_c is the set of all edges in city c , and Q_c is the number of edges in I_c .

In the paper, we report results using a narrow error bandwidth of $\nu = 2^\circ$. We experimented with a wider bandwidth of 5° . We also experimented with allowing for more than one dominant grid orientation, because for instance larger cities can have smaller sub-grids whose orientation differs from that of the main grid.⁴⁰ These variations produce highly correlated rankings of cities, and we therefore prefer the simplest version above. Visual inspection suggests that our methodology performs well at ranking road networks by how grid-like

⁴⁰We also experimented with weighting edges by length, but such measures appear to overestimate how grid-like small cities with few very long roads are.

Figure A.5: Most and least grid-like city road network using orientation grid metric



Panel A: Chandigarh - Grid Score = 0.54



Panel B: Jhansi - Grid Score = 0.08

they are. Figure A.5 shows the most and least grid-like cities according to the orientation metric, side-by-side.⁴¹

Gini. We modify the definition of the Gini index for income inequality to measure the normalized dispersion of edge bearings. For each city c , we define 360 different possible bearings, indexed by k , and ranked by their frequency such that $k = 1$ is the least frequent bearing and $k = 360$ is the most frequent bearing. In a perfectly gridded city, the four most frequent bearings, spaced 90 degrees apart, would account for 100% of edge bearings. Therefore, we can interpret high values of the following Gini index as corresponding to cities with a more grid-like network:

$$Gini_c = \frac{Q_c \times 360 - 2 \sum_{k=1}^{360} \sum_{l=1}^k \theta_{cl}}{Q_c \times 360}, \quad (\text{A3})$$

where θ_{cl} is the number of edges in city c with bearing l . The Gini and orientation metric have a correlation of 0.53.

The assumption of 360 possible distinct bearings is arbitrary, and we also computed Gini indices after rounding up each bearing to the nearest even degree (i.e., by assuming 180 possible bearings.) We also experimented with defining modulo 90 bearings (instead of modulo 360 as above).⁴² These variations produce Gini indices that are highly correlated with the index defined above that we use in the paper.

⁴¹It is also possible to compute measures of how grid-like the road network is separately for different types of road defined above, instead of only for the total road network. However, visual inspection suggests that these measures do not perform well at capturing overall how grid-like cities are, and for instance motorways are often curved and outside of the main grid.

⁴²For some smaller cities with sparser road networks, the number of distinct edge bearings is less than 360. In these cases, we adjust the calculation to consider only the total set of bearings present in that city, which may be less than 360.

Weather data

Hourly and daily historical weather data (rain, thunderstorm, temperature, humidity, and wind speed) are from the Weather Underground website. (<https://www.wunderground.com/history>, last accessed, 6 September 2018.) Weather Underground (WU) links each city to a station nearby (if there is one) and reads the weather reported by the station at the time it was reported.

We recovered weather data for 112 cities during the trips collection period. The median city-day has 8 weather readings, with a range from 1 to 144. On an average day, 25 of the cities report weather at least once every hour and 13 of them (mostly cities with international airports) report every half hour or more. The number of readings per day for a given city varies little across days.

The remaining 42 cities are missing data for one or more of the following three reasons. First, WU does not recognize the city name (4 cities). Second, WU recognizes the city name, but has no data on it (i.e., not linked to any weather stations – 31 cities). Third, WU re-directs to a different city name, either because: (a) WU recognizes our entry as an alternative name to the returned city, or (b) WU treats the city as a suburb or extension of a larger city nearby (20 cities). In this case, we accepted the returned city as a proxy as long as it was within 50 kilometers of the queried city (8 of 20 cities). Over the two months when we collected weather data, it rained 4.5% of the time and there were thunderstorms 2% of the time.

Comparison with Uber Movement data

We use data from Uber Movement (<https://movement.uber.com>) based on actual trips taken by Uber riders for four Indian cities, Mumbai, Delhi, Hyderabad, and Bangalore. Uber provides data by multi-hour period by day, or alternatively, by individual hour by calendar quarter. Given the importance of within-day variation in mobility, we use the latter for the last quarter of 2016 when we collected our main data. Uber Movement divides each city into small zones and provides a travel time between zone pairs for each hour of the day. For example, in Bangalore, there are 198 zones, averaging less than 4 square kilometers, and Uber Movement reports travel times for 87.2% of all $198 \times 197 \times 24 = 936,342$ zone pair-hours.⁴³

If these were average speeds of actual trips between these zone pairs, we could compare them directly to our GM trips. However, they differ from real trips in a critical way. Uber Movement computes zone-pair travel times using all trips that pass through a pair of zones. Because these zones are fairly small, we expect a typical Uber trip to generate information for many pairs. For example, a real trip that passes through ten zones will provide travel time

⁴³We exclude the travel time information that Uber Movement reports for each ward to itself, which is provided only for the last hour of the day.

information for $10 \times (10-1)/2 = 45$ different pairs. We call these pairs trip segments. These segments will contain a different distribution of zones than the actual trip. For example, the beginning and end zones represent 20% of the zones in an actual ten-zone trip. But in the 45 trip segments, they represent only 9%. This is a problem because we expect that the beginning and end of a trip are on average slower than their middle, for the same reason that long trips are faster than short trips – drivers must use slow local roads at the beginning and end of many trips.

A second complication of comparing our Google Maps data with the Uber Movement data is that there is no way of knowing the actual road distance traveled between two zones in an Uber Movement trip segment. Therefore, we measure the haversine distance between the centroids of the zones of the endpoints of either each trip instance (Google Maps) or each trip segment (Uber Movement) and compute the corresponding effective speed.


To make the samples comparable, we restrict attention to Google Maps trip instances matched to a corresponding Uber Movement observation by zone of origin, zone of departure, and hour of departure. Uber Movement’s effective definition of each city is generally smaller than ours. In Bangalore, 81.3% of our 495,708 Google Maps trip instances have both their origin and their destination in a zone reported by Uber Movement. Uber Movement provides matching travel times data (origin, destination, and hour) for 92.1% of these, or 371,215. Conversely, only 13.1% of zone-pair-hours reported by Uber Movement have a corresponding Google Maps trip in our Bangalore sample.⁴⁴

Consistent with our conjecture above, the effective travel speeds computed from Uber Movement travel times are much faster than the corresponding speeds computed from counterfactual trip instances obtained from Google Maps. We plot these speeds for all hours of the day in figure A.6. For instance, at 6 a.m., we observe an average effective speed of nearly 25 kilometers per hour for Uber Movement instead of about 16 kilometers per hour for Google Maps. Note that we pool data for all four cities together but the patterns are the same for each city taken individually. Note also that we do not condition out trip characteristics as we do in the rest of the paper because some characteristics like trip type are missing from the Uber data while others like trip length cannot be directly compared across both groups. Note finally that for Uber Movement trip segments in figure A.6, we only use trip segments for which the matching Google Map trip instance has a road length of less than 15 kilometers. This is because we compare these Uber Movement trip segments to Google Maps trip instances with a length of less than 15 kilometers.⁴⁵

⁴⁴Using instead all Uber Movement observations only makes minor differences. On the other hand, Google Maps trip instances within the Uber-defined city limits is further from the center on average.

⁴⁵We nonetheless keep in mind that these length ranges are not immediately comparable since trip segments of a given length can be part of much longer trip instances.

Figure A.6: Comparing Uber Movement with different parts of Google Maps trips



Pooled data for four cities (Mumbai, Delhi, Hyderabad, and Bangalore). Series computed as described in the text.

Since Uber Movement oversamples the fast middle parts of trips as noted above, it is more appropriate to compare speeds computed from Uber Movement data with speeds computed for the middle part of Google Maps trips instead of the entire trips. To compute speed for the middle part of trips, we proceed as follows. Using our supplementary data collection from August 2018, Google Maps provides length and duration data for each step of a trip instance (that is, each portion of a trip between changes of direction) in meters and seconds. Unfortunately, step travel speed is not updated in real-time. It appears to reflect an average travel speed for that step.

We thus create step-specific speeds as follows. Let $S_{i\tau}$ be the (real-time) speed of whole trip instance i at time of day τ in 5-minute increments, and $Z_{in\tau}$ be the (average) speed of segment n that is part of trip i at time of day τ , where n could be the first kilometer, or the middle (defined as all but the first and last 2 kilometers). We calculate estimates of $S_{in\tau}$ for all Google Maps trip instances i with an Uber Movement counterpart as $S_{in\tau} = S_{i\tau} \times \frac{Z_{in\tau}}{Z_{i\tau}}$, where $\frac{Z_{in\tau}}{Z_{i\tau}}$ is the speed of segment n relative to the speed of the whole trip. We then average this quantity across all trips in each hour $\bar{\tau}$ of the day to get $S_{n\bar{\tau}}$ for two segments n : the first kilometer and middle, in turn separating out the middle for trips of 5–10 kilometers and 10–15 kilometers.

The resulting speed profiles by time of the day are plotted in figure A.6. While the first

kilometer of trips is sizeably slower than their average speed, the middle part of trips is much faster, especially for longer trips that are more likely to use faster roads.⁴⁶ During most of the day, we find that the speeds we infer from Uber Movement travel times are close to the speeds we estimate for the middle part of 10–15 kilometer trip instances and slightly above the speeds of the middle part of 5–10 kilometer trips.

The pattern of remaining deviations between the Uber Movement and Google Maps data are potentially interesting. The largest deviation is in the uncongested early morning, where Uber movement trips are faster than our Google trips. Recall that speeds for the middle parts of trips are estimated under the assumption that congestion delays affect all parts of trips equally (since step durations reported by Google Maps are invariant during the day). If, consistent with the results of Akbar and Duranton (2018), the middle part of trips on major roads is subject to worse congestion delays than the beginning and end parts on local roads, our calculation of the speed for the middle part of trips should be an underestimate at the least congested hours of the day and an overestimate at the most congested hours of the day. This appears true in at least a relative sense in figure A.6.

Public holidays

In total, 11 public holidays took place during our data collection period in 2016:

1. 15-16 September: Onam in Kerala.
2. 23 September: Haryana's Heroes' Martyrdom Day in Haryana.
3. 28 September: Birthday of Bhagat Singh Ji in Punjab.
4. 30 September: Mahalaya in Karnataka, Orissa, and West Bengal.
5. 10 October: Ayudha Puja in Karnataka and Tamil Nadu, Durga Puja in West Bengal, and Mahanavami in Bihar, Kerala, and Orissa.
6. 12 October: Durga Puja in Orissa.
7. 17 October: Parkash GURPURAB of Sri Guru Ram Das Ji in Punjab.
8. 31 October: Vikram Samvat New Year in Gujarat, Deepavali in Karnataka and Maharashtra, and Vishvakarma Day in Punjab.
9. 1 November: Bhai Bij in Gujarat and Haryana Day in Haryana.
10. 7 November: Chhath Puja in Bihar.
11. 14 November: Rassa Purmina in Orissa.

⁴⁶Although we do not plot them here, we find that the last kilometer of trips is about as slow as the first. The second and second to last kilometers are also slower than the trip average but faster than the first and last kilometers.

On three of these days (31 October, 1 November, and 14 November), we did not happen to be sampling cities in the affected states. For each of the other public holidays, we estimate a variant of our preferred specification (column 5 of table 4) adding an indicator for the holiday, limiting our sample to the cities of the state(s) where the holiday occurs. Given our concern that data quality may be lower in small cities, we alternatively limit samples to cities of less than 0.5 million residents; in some cases this means a single city.

Excluding the exceptions discussed below, mobility on public holidays was significantly faster, even in the small-city samples, consistent with less traffic. Improvements were generally moderate, but occasionally sizable. For instance, we estimate 5 to 6% faster mobility in Bangalore and Kolkata, two highly congested cities, during Mahalaya. For small cities, these improvements in mobility are small, between 0.5 and 2% or less, which is reasonable since congestion delays are limited to start with.

The main exceptions to these patterns are the following. We estimate insignificant coefficients for Ayudha Puja (10 October) in the three small cities of Tamil Nadu and for Mahanavami (also 10 October) in the small cities of Kerala and Orissa.

We also estimate negative and significant coefficients for Durga Puja (10 October) in the small cities of West Bengal, for Parkash Gurburab of Sri Guru Ram Das Ji (17 October) in Amritsar, and Chhath Puja (7 November) in the small cities of Bihar. In all three cases, these are particularly important festivities that draw lots of visitors and involve outdoor celebrations. Negative effects on mobility should be expected. For instance, Parkash Gurburab of Sri Guru Ram Das Ji is a celebration of the founder of Amritsar, the holiest city of Sikhism. We find worse mobility in Amritsar that day while mobility is better in other cities of Punjab, which also celebrate the public holiday but do not draw large crowds.

Other web sources

300 million smartphones in India for the fourth quarter of 2016: <https://www.counterpointresearch.com/india-smartphone-user-base-tops-300-million-as-smartphone-market-grew-18-yoy-in-2016/> (accessed 7 September 2018).

97% share of Android smartphones among all smartphones: <http://indianexpress.com/article/technology/googles-android-captured-97-indian-smartphone-market-share-in-q2-2016-report-2957566/> (accessed 6 September 2018).

71% share of urban mobile internet users in 2015: <http://indianexpress.com/article/technology/tech-news-technology/mobile-internet-users-in-india-to-reach-371-mn-by-june-2016/> (accessed 6 September 2018).

Source for public holidays in India: <https://www.officeholidays.com/countries/india/2016.php>

Appendix B. Derivation and computation of the logit/CES mobility index.

We define the utility from visiting the destination of trip i in city c as:

$$u_{ci} = \log(b_{ci}) + (1 - \sigma) \log(t_{ci}) + \epsilon_{ci}, \quad (\text{B1})$$

where $t_{ci} = \gamma T_{ci}$ is the time cost of a trip to destination i in city c that takes T_{ci} units of time at value of time γ per unit, and ϵ_{ci} , the random component of utility, has a Type I extreme value distribution.⁴⁷ The parameter $\sigma > 1$ is an elasticity of substitution across destinations, and b_{ci} is a trip-specific quality parameter capturing all factors other than time costs making some destinations more desirable than others.⁴⁸

As shown by Anderson *et al.* (1992, pp. 60–61), the expected utility of a traveler in city c is equal to the expected value of u_{ci} 's maximum across the N_c travel destinations available in city c :

$$\mathbf{E} \left(\max_{i \in N_c} \{u_{ci}\} \right) = \log \left(\sum_{i=1}^{N_c} \exp [\log(b_{ci}) + (1 - \sigma) \log(t_{ci})] \right) = \log \left(\sum_{i=1}^{N_c} b_{ci} t_{ci}^{1-\sigma} \right). \quad (\text{B2})$$

Now consider two cities, c and c' . Define a relative price index $G_{c,c'}$ as the factor by which travel costs in city c would have to change in order to equalize expected utility in the two cities:

$$\log \left(\sum_{i=1}^{N_c} b_{ci} (G_{c,c'} t_{ci})^{1-\sigma} \right) = \log \left(\sum_{i=1}^{N_c} b_{c'i} t_{c'i}^{1-\sigma} \right). \quad (\text{B3})$$

It is easy to show that

$$G_{c,c'} = \left(\frac{\sum_i^{N_{c'}} b_{c'i} t_{c'i}^{1-\sigma}}{\sum_i^{N_c} b_{ci} t_{ci}^{1-\sigma}} \right)^{1/(1-\sigma)} = \left(\frac{\sum_i^{N_{c'}} b_{c'i} T_{c'i}^{1-\sigma}}{\sum_i^{N_c} b_{ci} T_{ci}^{1-\sigma}} \right)^{1/(1-\sigma)}, \quad (\text{B4})$$

where the second equality uses $t_{ci} = \gamma T_{ci}$. The relative price index $G_{c,c'}$ is best characterized as a relative travel accessibility index. It is low when comparing cities that have many destinations to those with few (gains from variety), and when comparing cities where travel to those destinations is short-distance and fast to those where it is long-distance and slow.

We now develop a simple non-parametric procedure to isolate a pure mobility index determined only by speed differences across cities. To do this, we replace the denominator of $G_{c,c'}$ with a 'national index' that has exactly the same distribution of trip length as in city

⁴⁷Ben-Akiva and Lerman (1985) are the first to show how to derive a travel accessibility index from a logit model of travel demand. Anderson, de Palma, and Thisse (1992) are the first to show the correspondence between the logit and CES models.

⁴⁸In Table 8, we present an index computed at $\sigma = 0$. Technically, values of $\sigma < 1$ are inconsistent with utility maximization. In practice, the index at $\sigma = 0$ simply weights all trips equally and intuitively corresponds to a perfect complement case.

c , and the same number of trips. This leads to equation (6) in the main text. Note that we inverted the index to ensure that G_c increases with faster speed (the index derived above is a price index increasing with time costs.) We compute \bar{T}_{ci} as the average travel time of all trips in the national sample with length within 1% of that of trip i in city c . We drop any trip with fewer than 10 corresponding trips within 1% of its length in the national sample (less than 0.01% of trips).

We investigate robustness to the parametrization of the quality parameters b_{ci} . For this investigation, we restrict the sample to amenity trips. We do not observe the quality of destinations, but we sampled amenity trips to match the trip shares in the US NHTS, so assuming that $b_{ci} = 1$ for all amenity trips is a reasonable starting point to compute G_c . We then compute variations of this index using random draws of $b_{ci} \in U[1,100]$, thus randomly allowing certain destinations to be more desirable and to carry a higher weight in the index. Indices obtained from these draws are highly correlated with one another and with our benchmark index. This exercise corroborates other findings from Table 8, showing that slow cities are slow for all types of trips, and that weighting certain trips more than others has little impact on our mobility indices.

Finally, we divide trips into M groups and compute the following nested CES/logit mobility index (Sheu, 2014):

$$G_c^{nest} = \frac{\left(\sum_{m=1}^M G_{mc}^{1-\mu}\right)^{\frac{1}{1-\mu}}}{\left(\sum_{m=1}^M \bar{G}_{mc}^{1-\mu}\right)^{\frac{1}{1-\mu}}}, \quad (\text{B5})$$

and

$$G_{mc} = \left(\sum_{i=1}^{N_{mc}} b_{ci} T_{ci}^{1-\sigma}\right)^{\frac{1}{1-\sigma}}, \quad \bar{G}_{mc} = \left(\sum_{i=1}^{N_{mc}} b_{ci} \bar{T}_i^{1-\sigma}\right)^{\frac{1}{1-\sigma}}, \quad (\text{B6})$$

where $\mu > 1$ is the elasticity of substitution across groups, $\sigma > \mu$ is the elasticity of substitution within groups, and N_{mc} is the number of trips in group m in city c . As an example, we can define eight groups, one for each amenity type recorded in Appendix A. In this case, the nested index G_c^{nest} puts less weight on destination types that are relatively slower in city c ; travelers substitute away from them because they are costlier. We compute these indices using exactly the same methodology as before. Setting $\mu = 1.5$ and $\sigma = 2.5$, we experiment with various nesting structures defined by time (e.g., off peak, low peak, high peak), area (e.g., rings), and type of destination (e.g., amenity types), and find a high correlation with our benchmark index in all cases.

Appendix C. A simple model of supply and demand for travel in a city

In this section, we derive the key parameters to estimate the value of increasing uncongested speed versus reducing congestion. Consider the setup in Figure C.1, panel b. Demand is given by $P = a - bQ$. Average costs are $AC = c + fQ$, where we observe c , the intercept, directly as the uncongested inverse speed. The equilibrium inverse speed is:

$$P^{eq} = a - b \frac{a - c}{b + f} = a \frac{f}{b + f} + c \frac{b}{b + f}. \quad (C1)$$

The part of an increase in uncongested speed that is not crowded out by congestion (the shaded rectangle in Figure C.1, panel b) is given by:

$$P_H^{eq} - P_L^{eq} = a - b \frac{a - c_H}{b + f} - \left(a - b \frac{a - c_L}{b + f} \right) = \frac{b}{b + f} (c_H - c_L) \quad (C2)$$

The share of the uncongested mobility improvement ($c_H - c_L$) that persists in equilibrium is $\frac{b}{b+f}$. This share is small if the average cost curve is steep (congestible network) and if demand is flat (drivers are price-sensitive).

To estimate the structural equation (C1), we run the regression:

$$p = \alpha + \beta c + \epsilon \quad (C3)$$

β will be an unbiased estimate of $\frac{b}{b+f}$ unless c is correlated with $a \frac{f}{b+f}$. The key concern here is correlated supply and demand shocks. For example, higher incomes could decrease c and increase a . This would however create a downward bias in β . We thus view our estimate as conservative. In our data, uncongested mobility (c) is correlated with congestion (which rises with a) at only 0.2, suggesting that this is not a severe bias.

The area of the deadweight loss triangle in Figure C.1, panel c is (height*base)/2:

$$DWL = \frac{1}{2} (Q^{eq} - Q^{opt}) * (MC(Q^{eq}) - P^{eq}) = \left(\frac{a - c}{b + f} \right)^2 \frac{f^2}{2(b + 2f)} \quad (C4)$$

We define Δ as the share of travel cost that is a deadweight loss:

$$\Delta = \frac{DWL}{Q^{eq} P^{eq}} = \frac{f^2}{2(b + 2f)} \frac{Q^{eq}}{P^{eq}}. \quad (C5)$$

Our regression estimate $\frac{b}{b+f} = 0.85$ implies that $f = 0.176b$, suggesting that either supply is very flat or demand is very steep. From above, both of these imply that the DWL is small. Plugging $f = 0.176b$ into equation (C5), we get $\Delta = 0.01b \frac{Q^{eq}}{P^{eq}}$. With linear demand, the price elasticity of demand is just $-\frac{dQ}{dP} \frac{P}{Q} = -\sigma = \frac{1}{b} \frac{P}{Q}$, so $\Delta = 0.01/\sigma$. If we are willing to assume unit elasticity locally at equilibrium, then $\Delta = 0.01$: the DWL is one percent of travel costs.

Unreliability

As noted in the main text, we make two assumptions about unreliability: that it increases proportionally with the quantity of travel, and that it is valued as much as travel delay. Incorporating unreliability (u), the average cost curve thus becomes $AC_u = c + (f + u)Q$. Our estimates above assume $u = 0$, so we must first determine the relationship between f and u . The per km cost of an uncongested trip is c . The cost of an equilibrium peak trip including only travel delay is $1.2c$ (based on 20% average travel delay), so in equilibrium $AC = c + fQ = 1.2c$. Accounting for unreliability as well, the equilibrium cost of a trip becomes $(1 + 0.2 + 0.08)c = 1.28c$ (based on 8% average unreliability). Putting this together, the average cost at equilibrium including unreliability is $AC = c + (f + u)Q = 1.28c$, so that $u = 0.4f$.

Next we estimate the downward bias, due to ignoring unreliability, in our regression estimate of how much of uncongested speed differences are crowded by additional traffic. Denote by P_u^{eq} the equilibrium price inclusive of the cost of unreliability, and by P_{obs}^{eq} the price that we observe and use in our regression (our peak mobility index), which excludes unreliability. Their difference is just the difference between the two analogous average cost curves at the same equilibrium:

$$P_u^{eq} - P_{obs}^{eq} = uQ^{eq}$$

This means that the regression that we would like to run is:

$$P_u^{eq} = a - b \frac{a - c}{b + f + u} = a \frac{f + u}{b + f + u} + c \frac{b}{b + f + u}, \quad (c6)$$

but the regression that we actually run is:

$$P_{obs}^{eq} = P_u^{eq} - uQ_{eq} = a \frac{f}{b + f + u} + c \frac{b + u}{b + f + u}. \quad (c7)$$

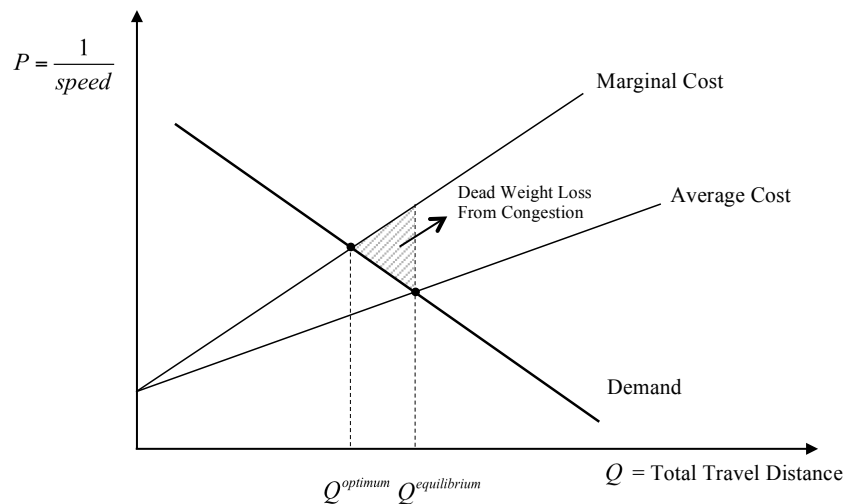
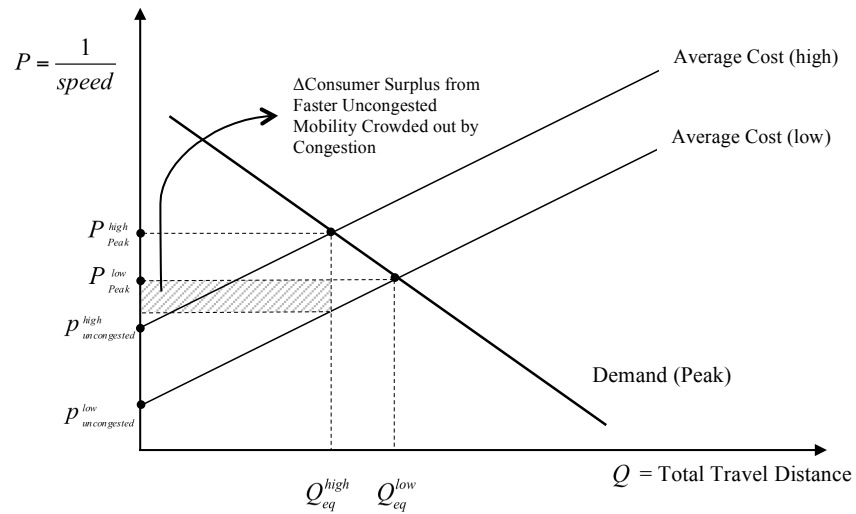

In other words, by ignoring unreliability we obtained a biased estimate $\frac{b+u}{b+f+u}$ of the structural parameter $\frac{b}{b+f+u}$. Our estimate of $\frac{b+u}{b+f+u} = 0.85$, combined with $u = 0.4f$, implies that $\frac{b}{b+f+u} = 0.79$. Similarly, we can obtain Δ (again assuming unit demand elasticity) by replacing f with $f + u$ in (c5):

$$\Delta = \frac{DWL}{Q^{eq} P^{eq}} = \frac{(f + u)^2}{2(b + 2(f + u))} \frac{Q^{eq}}{P^{eq}} = 2.3\%.$$

Appendix D. Additional tables

The four panels of table D.1 duplicate the results of table 4 for each type of trip separately. Table D.2 duplicates table 10 but uses as dependent variable a fixed effect from a trip regression where trips are weighted by how slow they are relative to their speed in absence of traffic

Figure C.1: Theoretical framework for valuing faster uncongested mobility



($\lambda = 0.2$). Table D.3 duplicates the specification of column 6 in table 10 but uses as dependent variables further alternative mobility indices.

Appendix E. A ring analysis of mobility in Indian cities

Although our main findings of city-level correlations in Section 6 are generally stable across a wide variety of specifications, they may be subject to bias due to omitted city-level variables. We now use within-city variation in population, area, and roads to avoid this problem and gain further insights about variation in mobility.

Specifically, we divide each city in our sample into concentric rings. Among other advantages, nearly all radial trips will pass through the same rings, regardless of route. We apply the following transformation of equation (2) which uses the location of trips within cities to estimate a mobility index for each ring within each city:

$$\log S_i = \alpha X_i' + \sum_r R_{rc(i)} \text{share}_{rc(i)}(i) + \epsilon_i, \quad (\text{E1})$$

where $\text{share}_{rc(i)}(i)$ is the share of trip i which takes places within ring r of city c and R_{rc} is a mobility index for ring r of city c . We consider (up to) 5 rings around each city center: 0 to 2 kilometers, 2 to 5, 5 to 10, 10 to 15, and 15 and beyond. We compute each trip's share in each ring using information about the origin and destination. For instance, a radial trip that starts 9 kilometers from the center and finishes one kilometer from the center on the same side will receive a share of 12.5% ($=1/8$) for the first ring of 0 to 2 kilometers, 37.5% for the second ring, 50% for the third ring and 0% for the fourth and fifth ring. We estimate equation (E1) using as controls log trip length, time of day and day of week indicators in manner that is consistent with our baseline index.

In a second step, we can estimate the following regression

$$\hat{R}_{rc} = \kappa_r + \beta_c + \alpha X_{rc}' + \epsilon_i, \quad (\text{E2})$$

where κ_r is a ring fixed effect, β_c is a city fixed effect, and X_{rc} is a vector of explanatory variables at the level of the city-ring. In our dataset, only land area, population, and roads are available separately by city-ring. Two caveats must be kept in mind. First, we winsorize the top and bottom 5% of city-ring effects before estimating equation (E2). This is because some cities barely enter an outer ring and therefore these city-rings have a tiny number of trips. Second, we also expect some equilibrium effects across rings as, for instance, population in

Table D.1: Correlates of log trip speed for specific trip classes

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A. Radial trips							
log trip length	0.28 ^a (0.0063)	0.073 ^b (0.033)	0.074 ^b (0.033)	0.28 ^a (0.0063)	0.069 ^b (0.033)	0.069 ^b (0.033)	0.049 (0.059)
log trip length ²		0.055 ^a (0.011)	0.055 ^a (0.011)		0.057 ^a (0.011)	0.057 ^a (0.011)	0.063 ^a (0.018)
log distance to center		0.15 ^a (0.048)	0.15 ^a (0.048)		0.16 ^a (0.049)	0.16 ^a (0.049)	0.15 ^b (0.076)
log distance to center ²		-0.11 ^b (0.042)	-0.11 ^b (0.042)		-0.11 ^b (0.042)	-0.11 ^b (0.042)	-0.13 ^b (0.061)
Observations	5,102,925	-	-	4,347,207	-	-	2,313,862
R-squared	0.53	0.54	0.54	0.53	0.54	0.54	0.57
Panel B. Circumferential trips							
log trip length	0.26 ^a (0.0083)	0.056 ^c (0.030)	0.056 ^c (0.030)	0.26 ^a (0.0082)	0.053 ^c (0.030)	0.053 ^c (0.030)	0.037 (0.058)
log trip length ²		0.060 ^a (0.010)	0.059 ^a (0.010)		0.060 ^a (0.010)	0.060 ^a (0.010)	0.066 ^a (0.018)
log distance to center		0.15 ^b (0.060)	0.15 ^b (0.060)		0.15 ^b (0.062)	0.15 ^b (0.062)	0.14 (0.11)
log distance to center ²		-0.13 ^b (0.051)	-0.13 ^b (0.051)		-0.13 ^b (0.052)	-0.13 ^b (0.052)	-0.14 ^c (0.082)
Observations	2,261,556	-	-	1,934,692	-	-	1,018,394
R-squared	0.45	0.46	0.47	0.45	0.46	0.47	0.51
Panel C. Gravity trips							
log trip length	0.21 ^a (0.0032)	0.13 ^a (0.012)	0.13 ^a (0.012)	0.21 ^a (0.0032)	0.13 ^a (0.012)	0.13 ^a (0.012)	0.14 ^a (0.014)
log trip length ²		0.016 ^a (0.0031)	0.016 ^a (0.0031)		0.015 ^a (0.0032)	0.015 ^a (0.0032)	0.013 ^a (0.0035)
log distance to center		0.18 ^a (0.051)	0.18 ^a (0.051)		0.18 ^a (0.050)	0.18 ^a (0.050)	0.13 ^c (0.077)
log distance to center ²		0.031 (0.026)	0.031 (0.026)		0.036 (0.025)	0.036 (0.025)	0.047 (0.038)
Observations	7,672,821	-	-	6,539,528	-	-	3,495,291
R-squared	0.38	0.45	0.45	0.37	0.45	0.45	0.46
Panel D. Amenity trips							
log trip length	0.25 ^a (0.0045)	0.17 ^a (0.011)	0.17 ^a (0.011)	0.25 ^a (0.0045)	0.17 ^a (0.010)	0.17 ^a (0.010)	0.16 ^a (0.014)
log trip length ²		0.0064 ^c (0.0034)	0.0064 ^c (0.0034)		0.0059 ^c (0.0033)	0.0059 ^c (0.0033)	0.0081 ^c (0.0044)
log distance to center		0.21 ^a (0.037)	0.21 ^a (0.037)		0.21 ^a (0.036)	0.21 ^a (0.036)	0.17 ^a (0.052)
log distance to center ²		0.0052 (0.019)	0.0052 (0.019)		0.0097 (0.019)	0.0097 (0.018)	0.019 (0.027)
Observations	7,706,854	-	-	6,564,229	-	-	3,492,392
R-squared	0.55	0.60	0.60	0.55	0.60	0.60	0.54
City effect	Y	Y	Y	Y	Y	Y	Y
Day effect	Y	Y	Y	weekd.	weekd.	weekd.	Y
Time effect	Y	Y	Y	Y	Y	Y	Y
Weather	N	N	Y	N	N	Y	only

Notes: OLS regressions with city, day, and time of day (for each 30-minute period) indicators. Log speed is the dependent variable in all columns. Robust standard errors in parentheses. *a*, *b*, *c*: significant at 1%, 5%, 10%. 154 cities in columns 1-7 and 107 in column 8. All trip instances in columns 1-3. Only weekday trip instances in columns 4-6. Only weekday trip instances for which we have weather information in column 7. Weather in column 3 and 6 consists of indicators for rain (yes, no, missing), thunderstorms (yes, no, missing), wind speed (13 indicator variables), humidity (12 indicator variables), and temperature (8 indicator variables). These variables are introduced as continuous indicator variables in column 7. Sample sizes for columns 1 and 4 apply to columns 1-3 and 4-6, respectively.

Table D.2: Correlates of city mobility indices, mobility index for which trips are weighted by powered congestion factor

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log population	-0.19 ^a (0.021)	-0.19 ^a (0.021)	-0.18 ^a (0.021)	-0.18 ^a (0.021)	-0.18 ^a (0.020)	-0.17 ^a (0.023)	-0.18 ^a (0.021)	-0.17 ^a (0.021)
log area	0.13 ^a (0.020)	0.11 ^a (0.022)	0.11 ^a (0.022)	0.10 ^a (0.022)	0.11 ^a (0.021)	0.10 ^a (0.025)	0.11 ^a (0.022)	0.085 ^a (0.024)
log roads		0.015 ^a (0.0057)	0.013 ^b (0.0060)	0.016 ^a (0.0052)	0.011 ^c (0.0058)	0.014 ^b (0.0060)	0.014 ^b (0.0060)	0.014 ^b (0.0056)
log income			0.46 ^a (0.17)	0.48 ^a (0.18)	0.42 ^b (0.18)	0.50 ^b (0.21)	0.46 ^a (0.18)	
log ² income			-0.15 ^a (0.055)	-0.15 ^a (0.055)	-0.14 ^b (0.055)	-0.16 ^b (0.062)	-0.15 ^a (0.055)	
Network / shape				0.36 ^a (0.11)	0.18 ^a (0.062)	0.097 ^b (0.037)		
Pop. growth 90-10							0.020 (0.037)	
share w. car								-0.18 (0.19)
share w. motorcycle								0.31 ^a (0.081)
Observations	153	153	153	153	153	142	153	152
R-squared	0.51	0.52	0.56	0.58	0.58	0.59	0.56	0.58

Notes: OLS regressions with a constant in all columns. The dependent variable is the city fixed effect estimated in the specification reported in column 5 of table 4 where trips are weighted by how slow they are relative to their speed in absence of traffic ($\lambda = 0.2$). Robust standard errors in parentheses. *a, b, c*: significant at 1%, 5%, 10%. See the footnote of table 10 for further details about the explanatory variables.

nearby rings may affect mobility locally. Given the limited precision of our population data, detecting such effects may be out of reach here.

We report results in table E.1. The coefficient on population is -0.084 in our baseline specification, and similar in the rest of the table.⁴⁹ We note that the population coefficients estimated in table E.1 are only about half those estimated in table 10. This may be because our measures of ring population are less precise. We also expect mobility within ring to be determined by population in neighboring rings.⁵⁰ Consistent with table 10, table E.1 also

⁴⁹It is only when we do not control for trip characteristics in the first step in column 2, that we estimate a slightly larger coefficient in absolute value. This is likely because longer trips are faster and predominantly take place in outer rings where population is less dense.

⁵⁰We experimented with specifications that also included population in neighboring rings. Estimated coefficients are generally small and insignificant.

Table D.3: Correlates of city mobility indices with alternative mobility indices

Dep. var.	(1) Effect. sp.	(2) Peak hrs.	(3) Mean	(4) Simp. FE	(5) Amenity	(6) cent.<5 km	(7) Laspeyres	(8) Paasche
log population	-0.15 ^a (0.019)	-0.18 ^a (0.017)	-0.17 ^a (0.029)	-0.17 ^a (0.015)	-0.17 ^a (0.017)	-0.17 ^a (0.016)	-0.17 ^a (0.021)	-0.17 ^a (0.018)
log area	0.072 ^a (0.020)	0.12 ^a (0.018)	0.21 ^a (0.032)	0.15 ^a (0.016)	0.12 ^a (0.019)	0.13 ^a (0.017)	0.12 ^a (0.024)	0.15 ^a (0.021)
log roads	0.014 ^a (0.0048)	0.011 ^b (0.0044)	-0.0066 (0.0097)	0.012 ^a (0.0044)	0.013 ^a (0.0047)	0.013 ^a (0.0043)	0.027 ^a (0.0062)	0.011 ^b (0.0051)
log income	0.25 ^b (0.10)	0.27 ^b (0.11)	-0.032 (0.23)	0.20 ^b (0.094)	0.17 (0.12)	0.19 ^c (0.10)	0.25 (0.18)	0.15 (0.12)
log ² income	-0.073 ^b (0.032)	-0.084 ^b (0.034)	0.0083 (0.064)	-0.054 ^c (0.029)	-0.049 (0.036)	-0.048 (0.032)	-0.088 ^c (0.052)	-0.046 (0.037)
Observations	153	153	153	153	153	153	153	153
R-squared	0.56	0.60	0.28	0.55	0.56	0.54	0.42	0.47

Notes: OLS regressions with a constant in all columns. The dependent variable is the city fixed effect estimated using effective speed in column 1, only peak hour observations in column 2, a simpler speed regression in column 4, only amenity trips in column 5, only trips taking place within 5 kilometers from the center in column 6, our benchmark Laspeyres index in column 7, and a benchmark Paasche index in column 8. The dependent variable in column 3 is the log of a simple mean speed (length-weighted). Robust standard errors in parentheses. *a, b, c*: significant at 1%, 5%, 10%. Log population is constructed from the town population from the 2011 census. Log roads is log kilometers of primary roads within the city-light.

reports small positive coefficients for area. On the other hand, the coefficient on roads is generally negative, though it is only significantly different from zero when we focus on the city centers. Although we do not report the details here, this negative coefficient is driven mainly by the central ring when roads effects are allowed to vary by ring in columns 5 and 6. Finally, table E.1 also reports that mobility is generally faster in outer rings, which confirms earlier results from section 4.

Table E.1: Correlates of city mobility indices, rings analysis

	(1) Base	(2) No Step 1 Control	(3) < 5 km	(4) < 3 km	(5) Base	(6) <5m	(7) Peak	(8) Peak <5 km
log ring population	-0.084 ^a (0.013)	-0.13 ^a (0.018)	-0.086 ^a (0.010)	-0.089 ^a (0.010)	-0.085 ^a (0.013)	-0.088 ^a (0.011)	-0.090 ^a (0.013)	-0.086 ^a (0.010)
log ring area	0.038 ^b (0.018)	0.039 (0.024)	0.054 ^a (0.014)	0.058 ^a (0.014)	0.028 (0.019)	0.050 ^a (0.015)	0.048 ^a (0.018)	0.059 ^a (0.014)
log roads	-0.010 (0.0095)	-0.0017 (0.013)	-0.017 ^b (0.0076)	-0.018 ^b (0.0076)			-0.011 (0.0097)	-0.018 ^b (0.0076)
ring 2	0.14 ^a (0.019)	0.30 ^a (0.026)	0.082 ^a (0.015)	0.062 ^a (0.015)	0.10 ^b (0.043)	0.048 (0.035)	0.13 ^a (0.019)	0.073 ^a (0.015)
ring 3	0.20 ^a (0.025)	0.44 ^a (0.033)	0.092 ^a (0.020)	0.058 ^a (0.020)	0.19 ^a (0.036)	0.078 ^a (0.029)	0.20 ^a (0.025)	0.079 ^a (0.020)
ring 4	0.19 ^a (0.031)	0.46 ^a (0.042)	0.038 (0.024)	0.00043 (0.025)	0.15 ^a (0.045)	-0.014 (0.036)	0.19 ^a (0.031)	0.025 (0.025)
ring 5	0.20 ^a (0.042)	0.53 ^a (0.057)	0.036 (0.034)	-0.0067 (0.034)	0.15 ^a (0.055)	0.048 (0.044)	0.20 ^a (0.043)	0.018 (0.034)
roads per ring	N	N	N	N	Y	Y	N	N
Observations	467	467	466	465	467	466	467	466
R-squared	0.56	0.72	0.47	0.42	0.57	0.48	0.57	0.44

Notes: OLS regressions with a city fixed effect and a ring fixed effect in all columns (145 cities in all regressions). The dependent variable is the city-ring fixed effect estimated as per equation (E2). Robust standard errors in parentheses. *a*, *b*, *c*: significant at 1%, 5%, 10%. Column 1 is our baseline estimation for which city-ring effects are estimated as described in the text. Column 2 considers city ring effects estimated with out trip controls in the first step. Columns 3 and 4 only consider trips with a length of less than 5 and 3 kilometers respectively. Columns 5 and 6 estimate separate roads effects for each ring. Columns 7 and 8 duplicate columns 1 and 3 but only consider peak-hour trips.