Cities in Bad Shape: Urban Geometry in India*

Mariaflavia Harari†

September 2019

Abstract

The spatial layout of cities is an important feature of urban form, highlighted by urban planners but overlooked by economists. This paper investigates the causal economic implications of city shape in India. I measure cities' geometric properties over time using satellite imagery and historical maps. I develop an instrument for urban shape based on geographic obstacles encountered by expanding cities. Compact city shape is associated with faster population growth and households display positive willingness to pay for more compact layouts. Transit accessibility is an important channel. Land use regulations can contribute to deteriorating city shape.

JEL: R10, R30

*I am grateful to two anonymous referees, Jan Brueckner, Nathaniel Baum-Snow, Alain Bertaud, Dave Donaldson, Denise Di Pasquale, Esther Duflo, Gilles Duranton, Michael Greenstone, Melanie Morten, Daniel Murphy, Paul Novosad, Bimal Patel, Ben Olken, Champaka Rajagopal, Otis Reid, Albert Saiz, Chris Small, Kala Sridhar, Matthew Turner, Maisy Wong and seminar participants at MIT, NEUDC, UIUC, Columbia SIPA, LSE, Zürich, Wharton, the World Bank, Carey, the Minneapolis FED, the IGC Cities Program, the NBER Summer Institute, the Meeting of the Urban Economics Association, NYU, the Cities, Trade and Regional Development Conference at the University of Toronto, CEMFI, UPF, Stockholm School of Economics, Stockholm University, the IEB Urban Economics Conference, the Barcelona Summer Forum, PSU, CEU, and Stanford for helpful comments and discussions. JoonYup Park, Yuan Pei, and Candice Wang provided excellent research assistance. I am thankful to Prottoy Akbar, Victor Couture, Gilles Duranton, and Adam Storeygard for kindly sharing their data. I declare that I have no relevant or material financial interests that relate to the research described in this paper.

†The Wharton School, University of Pennsylvania. 428 Vance Hall, 3733 Spruce Street, Philadelphia, PA 19104-6301. Phone: (215) 5733503. Email: harari@wharton.upenn.edu.

1 Introduction

The United Nations estimates that cities will add more than 2.5 billion people by 2050, with nearly 90% of this increase occurring in Asia and Africa (UN, 2014). This will likely trigger a massive expansion in urban land (Seto et al., 2011), with India alone predicted to require 18.6 million more hectares by 2030 (McKinsey, 2010). Faced with the challenge of facilitating urban expansion, policy makers are making decisions on urban planning and infrastructural investments that will have persistent effects on the spatial configuration of economic activity within and across cities. Understanding how cities grow and which city forms best promote quality of life and economic growth is thus paramount.

This paper contributes to the debate on how to accommodate urban development by studying the economic implications of a previously overlooked feature of urban form: city shape. While largely ignored by the economics literature, the geometry of a city's footprint has long been emphasized by urban planners as important for transit accessibility and service delivery. All else being equal, a city with a more compact layout is characterized by shorter distances within the city, potentially affecting transit accessibility, public service delivery (such as electricity), and household and firm location choices. This, in turn, could impact firms' productivity and households' quality of life (Bertaud, 2004; Cervero, 2001), particularly in the developing world, where levels of service provision are lower and many city dwellers lack individual means of transportation.

Despite a common perception that developing country cities are expanding rapidly and haphazardly (Suzuki et al., 2010; UN-Habitat, 2016), we have little understanding of how the shape of urban expansion influences households and firms within and between cities. Among the key empirical challenges are the lack of data and the endogeneity of city shape, which is in itself an equilibrium outcome. By leveraging satellite-derived data and plausibly exogenous variation driven by topography, I provide the first causal estimates of the impacts of city shape on economic outcomes, in the context of Indian cities. I find that city geometry affects household location choices across urban areas: compact cities are associated with faster population growth and a negative compensating real wage differential, which suggests that they offer higher quality of life (Rosen, 1979, Roback, 1982).

The first contribution of this paper lies in the data and the measurement of the spatial properties of Indian cities over time. With over 470 urban agglomerations in rapid expansion (Census of India, 2011) and the world's second-largest urban population (UN, 2015), India is a particularly important context for studying urban form. An allegedly chaotic urban growth has been associated with sprawl and potentially distortive land use regulations (McKinsey, 2010), which makes urban form an important item in the Indian policy debate. However, systematic data on Indian cities and their spatial structures is not readily available. I assemble a novel panel dataset that covers over 350 Indian cities between 1950 and 2011 and includes detailed information on each city's spatial properties, microgeography, and city-level outcomes. I trace the dynamic evolution of urban footprints by combining newly geo-referenced historical maps (1950s) with satellite imagery of night-time lights (1992-2010).

The second contribution is to quantify city compactness through a shape metric that I then embed in a standard urban economic model. I employ a shape index used in urban planning, based on the average distance between any two points in a polygon. Higher values of this index indicate a less compact urban footprint, and longer within-city distances. As an example, consider the cities of Kolkata and Bangalore (Figure 1): Kolkata has a distinctive elongated shape, stretching along the North-South axis, whereas Bangalore, roughly shaped like a pentagon, has a more compact layout. Controlling for city area, the average linear distance between any two points in the city is 27% longer in Kolkata than it is in Bangalore.[1] This is likely to become more pronounced over time, as I find that large cities have a tendency to deteriorate in shape as they grow.

The third contribution of this paper concerns the identification strategy. Estimating the causal impact of city shape on economic outcomes is challenging, as the spatial structure of a city at any point in time is in itself an equilibrium outcome. Urban shape is determined by the interactions of city population growth, natural constraints, and policy choices, such as land use regulations and infrastructural investments. I propose a novel instrumental variable for urban geometry that combines geography with a mechanical model for city expansion. The underlying idea is that, as cities expand in space and over time, they face geographic constraints - steep terrain or water bodies - leading to departures from an ideal circular expansion path. The relative position in space of such constraints allows for a more or less compact development pattern, and the instrument captures this variation.

The construction of my instrument requires two steps. First, I employ a mechanical model for city expansion to predict the area that a city should occupy in a given year, based on its projected historical population growth. Using a predicted expansion path is important since the city's actual growth path would be endogenous. Second, I consider the largest contiguous set of developable land pixels within this predicted radius; these pixels together form a polygon that I denote as "potential footprint". I then instrument the shape of the actual city footprint with the shape of the potential footprint.

The resulting instrument varies at the city-year level, allowing me to focus on long differences in shape between 1950 and 2010 and abstract from time-invariant city characteristics. The identification effectively relies on changes in shape that a city undergoes over time, as a result of hitting geographic obstacles. Importantly, my instrument captures variation in the relative position of topographic obstacles, rather than the presence or the extent of particular geographic features, and its explanatory power is not limited to cities with extremely constrained topographies (e.g., coastal or high-altitude cities).

To frame the empirical question of the economic impacts of city shape, I turn to a simple framework of spatial equilibrium across cities (Rosen, 1979, Roback, 1982). As households and firms optimally choose where to locate, I hypothesize that they account for city shape when evaluating the trade-offs associated with different locations. Compact cities may offer advantages associated with better service delivery or with greater accessibility, stemming from the fact that all locations within the city are closer to one another. If compact shape makes a city operate more efficiently, population will flow to that city, bidding housing rents up and wages down, until utility is equalized everywhere. This argument suggests that compact cities should have, in equilibrium, larger populations and lower real wages.

With my instrument in hand, I take these reduced-form predictions to the data. I begin by demonstrating that more compact cities experience faster population growth. A one standard deviation im-

[1] This example is discussed in greater detail in Section 3.

provement in city compactness, corresponding to a reduction in the average within-city distance of 360 meters, is associated with a 3% population increase. Naïve OLS estimates have the opposite sign, as they are confounded by the fact that larger and faster-growing cities also tend to have more disconnected shapes in equilibrium.

These results are robust to using different methods to delineate urban footprints and alternative shape indicators, and survive many falsification checks. One of the main threats to the identification is that the instrument may capture direct effects of local geography on city-level outcomes - for example, water bodies can have an inherent amenity value or provide productivity advantages. Reassuringly, the instrument is not correlated with geographic characteristics such as elevation, distance from the coast, or ruggedness. To further strengthen the identification, I allow for differential responses to city shape in cities with different geographies or different soil characteristics (such as bedrock depth, presence of minerals, or crop suitability) and find very similar results. My results are also stable if I exclude from the sample cities with particular characteristics, including coastal and high-altitude cities, as well as fast- or slow-growing ones. Another concern for the identification stems from potential pre-existing trends that may be correlated with historical population growth rates. I show that the results are consistent using an alternative version of the instrument, that employs a completely mechanical model for city expansion and does not rely on projected historical population.

Next, I turn to rent and wage differentials across cities. In the spatial equilibirum framework, poor accessibility and worse service delivery in non-compact cities may require cross-city compensating differentials, to the extent that households and firms cannot fully optimize against poor shape at the within-city level. For example, in a non-compact, low-accessibility city, households may be forced to live or shop in less preferable locations, if their first-best ones require excessively long trips. Consistent with this hypothesis, I find that compact cities are characterized by lower real wages. I further provide a back-of-the-envelope calculation of households' implied willingness to pay for compact shape, equivalent to 5% of their income for a one standard deviation improvement in compactness. Along the same lines, I calculate the implied impact of city shape on firm productivity through the lens of the model, finding negligible effects. This suggests that firms may be able to offset the negative impacts of poor shape through margins other than their cross-city location choices.

Turning to mechanisms, I consider the two main channels emphasized by urban planners: service delivery and transit accessibility. I find no meaningful impacts of city shape on the share of households connected to tap water or electricity, suggesting that disconnected shape is not standing in the way of the delivery of utilities. In contrast, several pieces of evidence point to the importance of accessibility. First, the negative impact of non-compact shape on population is mitigated in cities with a denser and better-functioning road network or a larger share of households with car access. Second, cities with worse shapes have a less dense road network, suggesting higher costs of providing infrastructure in more disconnected cities.

Furthermore, work commuting patterns may be affected by city shape. The impacts of city shape on realized commutes are *a priori* ambiguous: households may respond to longer potential distances by incurring longer commutes but also by giving up certain trips entirely. This is difficult to investigate empirically due to a lack of commuting surveys, but I indirectly shed light on work commuting patterns by examining the location of firms within cities. Using street addresses to geo-locate establishments, I find that firms located in non-compact cities tend to cluster in few employment sub-centers. This suggests that firms may be able to neutralize the effects of poor city shape and still take advantage of agglomeration by locating near one another, leaving it to workers to bear the costs of longer commutes. This is consistent with the interpretation of the compensating differentials discussed above, where city shape is associated with positive household willingness to pay but with no differences in firm productivity.

Finally, I consider the role of policy, specifically land use regulations, as one of the determinants of city shape. I show that more permissive vertical building limits, in the form of higher Floor Area Ratios (FARs),[2] result in less spread-out and more compact cities. For given geography, increasing FARs by one improves compactness by one standard deviation. This provides new evidence on the potentially distortive effects of land use regulations in India (Brueckner and Sridhar, 2012), highlighting a new margin: making cities less compact.

Taken together, these results indicate that the spatial configuration of cities has real consequences for the quality of life of urban dwellers, for their location patterns across cities, and, potentially, for their welfare. This has important implications for policy makers taking decisions related to urban planning or infrastructure, particularly in rapidly growing cities, and suggests that the impact of urban policies on city shape should be accounted for in cost-benefit analyses.

The rest of the paper is organized as follows. Section 2 provides some background on urbanization in India and reviews the related literature. Section 3 discusses the dataset and descriptive patterns. Section 4 outlines the conceptual framework. Section 5 details the empirical strategy and the instrument. Section 6 presents my main empirical results. Section 7 addresses identification threats. Section 8 interprets the results on wages and rents in terms of compensating differentials and provides willingness-to-pay estimates. Section 9 presents results on mechanisms and heterogeneous effects, discussing the interactions between city shape and transit, utilities, and land use regulations. Section 10 concludes.

2 Background and Previous Literature

India represents a relevant and promising setting to study urban spatial structures. With a current urban population of 460 million (World Bank, 2018) growing at a 2.3% yearly rate, India has the second-largest urban population in the world, after China, and is projected to host 250 million new urban dwellers by 2030 (McKinsey, 2010). Importantly for my empirical strategy, India has a large number of cities of meaningful size, with over 50 urban agglomerations having more than 1 million inhabitants.

During the period studied in this paper (1950 through 2011), India experienced a massive urban transition, with the urban population growing from 62 to 377 million according to the Census. This has been accompanied by a significant physical expansion of urban footprints, at an estimated rate of 4.84% yearly between 1970 and 2000 (Seto et al., 2011). Urban expansion has typically occurred

[2] FARs are defined as the maximum allowed ratio between a building's floor area and the area of the plot on which it sits. Higher values are associated with taller buildings. The average FAR in the cities in my sample is 2.3.

beyond urban administrative boundaries (IIHS, 2013; World Bank, 2013), making it difficult to track in space using official administrative sources.

The more-than-proportional expansion in urban land has been associated with haphazard development and poor urban planning. Sprawl, lengthy commutes, and limited urban mobility are often cited among the perceived harms of rapid urbanization (World Bank, 2013; McKinsey, 2010). There is also a concern that existing land use regulations might contribute to distorting urban form (Sridhar, 2010; Glaeser, 2011). In particular, sprawl has been linked to vertical limits in the form of restrictive Floor Area Ratios (Bertaud, 2002 and 2004; Brueckner and Sridhar, 2012; Glaeser, 2011; Sridhar, 2010).[3]

Literature directly related to the geometric layout of cities is scant, but a number of strands are tangentially connected to this theme. The economics literature on urban spatial structures has mostly focused on the determinants of city size and of the population density gradient, often assuming that cities are circular or radially symmetric (see Anas et al., 1998, for a review). The implications of city geometry are left mostly unexplored.[4] A large body of empirical literature investigates urban sprawl (see Glaeser and Kahn, 2004), typically in the U.S. context, suggesting longer commutes as one of its potential costs. Although some studies identify sprawl with non-contiguous development (for instance, Burchfield et al., 2006, in the U.S; Baruah et al., 2017, in Africa), which is related to the notion of "compactness" that I investigate, in most analyses the focus is on decentralization and density, neglecting differences in geometry. I focus on a different set of spatial properties of urban footprints: conditional on the overall amount of land used, I consider geometric properties capturing compactness, and view population density as an outcome variable.

The geometry of cities has attracted the attention of the quantitative geography and urban planning literature, from which I borrow indicators of city shape (Angel et al., 2009). Descriptive analyses of the morphology of cities and their dynamics have been carried out in the urban geography literature (see Batty, 2008, for a review), which emphasizes the scaling properties of cities. Urban planners emphasize the link between city shape, intra-urban trip length, and accessibility, claiming that contiguous, compact, and monocentric urban morphologies are more favorable to transit (Bertaud, 2004; Cervero, 2001).

In terms of methodology, my work is related to that of Burchfield et al. (2006), who employ remotely-sensed data to track the extent of sprawl in US cities over time. The data I employ comes mostly from night-time, as opposed to day-time, imagery, and covers a longer time span.[5] Furthermore, Saiz (2010) examines geographic constraints to city expansion and relates it to the elasticity of housing supply. I use the same definition of geographic constraints, but I employ them in a novel way to construct a time-varying instrument for city shape.

Finally, by highlighting the implications of city shape for accessibility, this paper also complements

[3] Another example is the Urban Land Ceiling and Regulation Act, which has been claimed to hinder intra-urban land consolidation and restrict the supply of land available for development within cities (Sridhar, 2010).

[4] One exception is Bento et al. (2005), who incorporate a measure of city shape in their investigation of the link between urban form and travel demand in US cities. Differently from their approach, I incorporate time variation in urban form and I treat city shape as endogenous.

[5] Recently, night-time lights have been employed to detect urban markets in India by Baragwanath-Vogel et al. (forthcoming).

a growing literature on infrastructure, transit, and urban expansion in developing countries (Baum-Snow et al., 2017; Storeygard, 2016) and India in particular (Akbar et al., 2018, 2019; Kreindler, 2018).

3 Data

3.1 Sources and Dataset Construction

I assemble a city-year level dataset covering over 350 cities in the main estimation sample, for a period ranging from 1950 to 2010. I include data on the geometric properties of urban footprints, topography, and various city-level economic outcomes - in particular, population, wages and housing rents. This section discusses my primary data sources. A detailed description of data sources and methods can be found in Section A in the Appendix.

Urban footprints

The first step in constructing my dataset is to trace the footprints of Indian cities at different points in time. I retrieve the boundaries of urban footprints from two sources. The first is a set of historical maps of India from the U.S. Army Map Service (ca. 1950), that I georeferenced and used to trace the boundaries of urban areas as of 1950. A fragment of one such map, showing the city of Mumbai, is shown in Figure 2. From these maps I am able to trace the footprints of 351 cities.

Second, I employ the DMSP/OLS Night-time Lights dataset, a series of night-time satellite imagery recording the intensity of Earth-based lights for every year between 1992 and 2010, with a resolution of approximately 1 square kilometer.[6] I use night-time lights imagery to delineate urban areas by considering spatially contiguous lighted pixels surrounding a city's coordinates, with luminosity above a pre-defined threshold of 35. This approach is illustrated in Figure 3. Employing higher (lower) thresholds results in more (less) restrictive definitions of urban areas and fewer (more) detected footprints overall, but does not affect the main results, as I show in Section 6.[7] In general, the resulting definition of urban areas is broad, extending beyond administrative boundaries. Through this procedure I retrieve up to 450 footprints per year.

Although this approach is not immune from measurement error, this is not a major concern in this setting since both area and shape of urban footprints will be instrumented throughout my analysis. Among other things, this addresses non-classical measurement error in the extents of urban footprints - for instance, due to a correlation between income and luminosity. Moreover, the goal is not to provide absolute estimates of urban land cover, but rather to explain changes. The long difference or fixed effects panel specifications employed throughout the paper account for differences in the definition of urban areas in different years (particularly between the U.S. Army maps and the night-time lights).

Combining these two sources, I retrieve footprints for a total of 6,172 city-years. The main estimation sample focuses on 2010-1950 long differences and includes 351 cities.

[6] These data have been widely employed in the economics literature, mainly for purposes other than urban mapping, starting with the seminal work of Henderson et al. (2012).

[7] See Baragwanath-Vogel et al. (forthcoming) for a discussion of the extents of Indian urban areas delineated employing night-time lights.

Shape metrics

Next, I quantify the compactness of urban footprints in each city-year. There are many possible indexes that measure compactness, defined as the extent to which a polygon's shape departs from that of a circle. I employ the disconnection index, an indicator borrowed from the urban planning literature (Angel et al., 2009). The index is defined as the average Euclidean distance, in kilometers, between any two points within a polygon, as illustrated in Figure A.1 in the Appendix.[8] For a given footprint area, higher values of the disconnection index are associated with larger distances between points in the city and a less compact shape. Figure A.2 in the Appendix provides examples of polygons with varying degrees of disconnection: elongated shapes and polygons with recesses and gaps (similar to urban areas growing around topographic obstacles) are all associated with greater disconnection relative to circular polygons with similar areas. For robustness, in Section 6 I also consider alternative indexes of compactness, which tend to be highly correlated with one another.

Importantly, any compactness index based on distances within a polygon is mechanically correlated with polygon area. In order to disentangle the effect of geometry *per se* from that of city size, in all of my specifications I control for the area of the footprint (which in the instrumental variables specification will be separately instrumented for, as discussed in Section 5). Alternatively, the index can be normalized, computing a version that is invariant to the area of the polygon. My results are robust to this alternative approach (discussed in Section 7).

To illustrate how the index maps onto urban shape, Figure 1 displays the footprints of Bangalore and Kolkata in 2005, where Bangalore's footprint has been rescaled so that they have the same area. Among India's best-known cities, Bangalore and Kolkata have among the most and the least compact geometries.[9] The difference in the rescaled disconnection index indicates that, if Kolkata had the same compact shape as Bangalore, the average potential distance within the city would be shorter by 4.4 kilometers.[10]

Outcomes

The outcome data I consider include city population, wages, and rents. Population data at the city level for the period 1871-2011 is obtained from the Census, available at 10-year intervals. Urban footprints, as retrieved from the night-time lights dataset, do not always have an immediate Census counterpart. The calculation of footprint-level population totals requires intermediate steps, detailed in Section A3 in the Appendix.

Wages and rents data are not systematically available at the city level for India. I thus employ

[8]Section A2 in the Appendix provides the mathematical formula. The index is calculated numerically, by sampling pairs of interior points from a polygon and averaging their distances. The shortest connecting paths used to define distance do not need to lie within the polygon.

[9]For the interested reader, Table A1 in the Appendix shows a list of the top most and least compact cities among those with over a million inhabitants. Cities are ranked by their normalized shape index, so that the ranking is not confounded by city size.

[10]This comparison is based purely on shape, holding city area constant. Even if Bangalore has a relatively efficient geometry, the overall spatial extent of the city may well be inefficiently large, as highlighted by Bertaud and Brueckner (2005).

coarser, district-level data and use district urban averages as proxies for city-level averages (as in Chauvin et al., 2017). The matching between cities and districts is not one to one. I thus provide results obtained with different matching approaches (including dropping districts that include more than one city or considering only the top city in each district).

Data on wages are drawn from the Annual Survey of Industries (ASI), waves 1990 and 2010.[11] The ASI consists of a series of repeated cross-sections covering manufacturing plants in the formal sector.[12]

Data on rents are drawn from the National Sample Survey (rounds 2005-2006 and 2007-2008), in which households are asked about the amount spent on rent and about the floor area of their dwelling.[13] Average rents calculated from NSS data are likely to underestimate the market rental rate, due to rent control provisions in most major cities of India (Dev, 2006). In the Appendix I thus show that my results are similar if I exclude the bottom 25% of reported rents for each city, where it is *a priori* more likely to find observations from rent controlled units.

For robustness, I also employ an alternative source of data on rents: the India Human Development Survey (2005 and 2012), used amongst others by Chauvin et al. (2017). An advantage over the NSS data is that respondents report not only rent but also a number of dwelling characteristics, allowing me to consider the residuals of a hedonic rents regression as an outcome.

Other data

To construct my city shape instrument, I employ high-resolution data on each city's microgeography. I consider land pixels as "undevelopable" when they correspond to a water body or have a slope above 15% (as in Saiz, 2010). Data on water bodies and slope are drawn respectively from the MODIS Raster Water Mask (with a resolution of 250 meters) and the ASTER dataset (30 meters). Figure 4 illustrates this classification for the Mumbai area.

For my robustness checks, I collect data on many additional city characteristics including topography and geology controls. Furthermore, for my analyses of mechanisms and heterogeneous effects, I assemble data on infrastructure (including current road length from Openstreetmap), firm location within cities (from street addresses in the Economic Census), availability of public services and slum population (from the Census), and land use regulations (from Sridhar, 2010). All these data sources are discussed in Section A in the Appendix.

Assembling city-year level data for Indian cities is not a straightforward exercise due to a lack of city-level sources and poor matching across different datasets. It should further be noted that data at a more disaggregated level is typically not available for India. In particular, a limitation is that I cannot systematically observe household location and commuting patterns within cities.

[11] Using intermediate waves in a panel specification yields similar results.

[12] This selective coverage may affect my results, to the extent that manufacturing is systematically over- or underrepresented in cities with worse shapes. However, I examined the relationship between shape and the industry mix of cities, employing data from the Economic Census, and found no obvious patterns. The share of manufacturing appears to be slightly lower in non-compact cities, but this figure is not significantly different from zero, which alleviates the selection concern discussed above.

[13] For those who own, an imputed figure is provided. Results are similar when excluding owners from the sample.

3.2 Descriptive Statistics

As a preliminary step to the causal investigation of the impacts of city shape, I provide descriptive evidence on the spatial properties of Indian city footprints. Summary statistics for city area and shape are provided in Table 1. Recall that shape is measured as the average within-city distance, in kilometers, with higher values of the index denoting less compact shapes. Panel A reports statistics for the full panel of years in which the footprints of those cities are observed. Panel B shows averages for years 1950 and 2010 and for the long difference on which most of my analyses are based.

The average city in my sample is relatively large, with a population of over 600 thousand inhabitants, a night-light-based footprint area of 118 square kilometers (about twice the land area of the borough of Manhattan), and an average within-city distance of 4.7 kilometers as of 2010. As expected, there is considerable variation in city shape across cities (partly driven by the variation in city area) and less variation for a given city over time. In 2010, the standard deviation in shape across cities was 4.2 kilometers, while the average within-city standard deviation across the years in the panel is only 1 kilometer. This is not surprising, given the path-dependence of urban form. City shape also appears to have a skewed distribution, consistent with similar patterns in the distribution of city size and area. Finally, the average within-city distance in 2010 is nearly five times larger than in 1950. While this may indicate a deterioration in urban shape over time, it is confounded by the massive expansion in the land area of urban footprints over the sample period and by differences in the methodologies used to trace urban areas in different years. To gain further insight, rather than focusing on the absolute values, in Table A2 I examine correlates of city shape, in levels and changes. While these correlations are purely descriptive, they help establish some key stylized facts in the data and motivate my causal estimation strategy. Each row in Table A2 corresponds to a city attribute and reports OLS coefficients from two independent cross-sectional regressions: in column 1 I regress city shape in 2010 on the city attribute, controlling for 2010 city area; column 2 is analogous but the dependent variable is the 2010-1950 long difference in city shape and I control for 1950 city area. A description of the variables and a more detailed discussion of the results is provided in Section A4 in the Appendix. Summary statistics for the correlate variables are in Table A3. Below I highlight the key patterns.

First, there is a clear positive equilibrium correlation between city size and non-compact shape, which affects the interpretation of naïve OLS regressions of shape on city-level outcomes. Panel A shows that large cities have worse shapes and tend to become less compact over time. Ranking cities by their 1951 population, cities in the lower quartile are associated with lower values of the disconnection index, whereas those in the top quartile are associated with higher values, both in levels and in changes. In Section 5 I discuss potential explanations for these patterns: the selection of topographically privileged locations when a city is originally founded, sprawl as a response to expanding infrastructure, or difficulties in enforcing urban planning regulations in rapidly-growing cities.

In line with the results in panel A, panel B shows that cities with less compact shapes also have better access to electricity, tap water, and cars, likely reflecting the fact that larger settlements also have higher incomes and better public services. At the same time, disconnected cities do not appear to have a more developed road network, which could stem from the difficulty of providing road infrastructure in cities with disconnected layouts. Interestingly, the disconnection index is positively correlated with the average distance to workplace in 2011 (drawn from the Census), also pointing to more difficult transit in non-compact cities. Section 9 is devoted to further causal exploration of these patterns.

Panel C highlights that city-level geographic features associated with natural advantage (such as a city's distance from the coast or crop suitability) are not predictors of city shape. An identification strategy based on instrumenting city shape with local topography may raise concerns associated with confounding effects of local geography through natural advantage. The lack of correlation in panel C is reassuring and these concerns are further assuaged by the robustness checks discussed in Section 7.

Finally, panel D shows the correlation between city shape and a number of non-predetermined characteristics capturing initial conditions, such as distance from other cities or colonial origin. State capitals appear to have experienced a large deterioration in shape, consistent with the correlation between shape and city size. Moreover, initial shape appears to be a strong predictor of current shape, suggesting that changes in shape are more informative than levels.

4 Conceptual Framework

In this Section I present a simple economic framework to motivate my empirical analysis, connecting the economic value of city shape with population, wages, and housing rents. According to the urban planning view, compact urban layouts offer households and firms advantages in terms of accessibility and public services, stemming from shorter within-city distances (Bertaud, 2004; Cervero, 2001). I embed this idea in a model of spatial equilibrium across cities. In this framework, households and firms optimally choose in which city to locate and, in equilibrium, they are indifferent across cities with different attributes. I hypothesize that they may value the "compactness" of a city as they evaluate the trade-offs associated with locating in different cities. Spatial equilibrium implies that, if compact shape makes a city operate more efficiently, population will flow into that city, bidding up housing costs and bidding down labor costs, until utility is equalized everywhere (Henderson, 1974). Along the same lines, if compact cities offer productivity advantages, they will attract firms, which will bid up labor costs, until profits are equalized. The city-level responses of population and factor prices to changes in city shape thus allow me to shed light on the value of compact urban layouts for households and firms by revealed preferences (Rosen, 1979, Roback, 1982). Below I provide a brief description of the model (adapted from Glaeser, 2008) to highlight the key reduced-form implications. The full model is provided in Section B in the Appendix, along with a discussion of caveats and extensions.

Assume identical and perfectly mobile households choosing optimally where to locate among a menu of cities. Their utility depends on the consumption of a (numéraire) tradeable good C and housing H, and on a vector of city characteristics θ. Households' maximization problem reads:

$$\max_{C,H} U(C,H,\theta) \; s.t. \; C = W - p_h H \tag{1}$$

where W is labor income and p_h is the rental price of housing (both of which are city-specific). Broadly, θ captures all utility costs and gains of living in a city. For the purposes of my discussion, it is useful to conceptualize θ as consisting of three components: "public services" θ_P, "transit accessibility" θ_T, and "consumption amenities" θ_A. All else being equal, better public services (such as electricity or water), greater accessibility, and better amenities (such as good climate) improve household utility. Denoting city shape with S, I assume that S can affect θ_P and θ_T, in line with the conjectures of urban planners.

The role of "accessibility" in this context deserves some discussion. Subsumed in this cross-city framework is a complex within-city location and travel problem that households face once they have chosen a city. This involves simultaneously choosing where to reside, work, shop, and consume leisure within the city, what mode of transportation to use, and how many trips to make out of a large menu of potential trips (Small and Verhoef, 2007). A city with poor shape can be thought of as offering a worse menu of choices than one with good shape, as some of the potential trips offered are longer. Once the within-city problem is solved, commuting trips are realized in equilibrium. Note that, as a city's shape deteriorates and potential distances increase, realized trips may become longer or shorter, as one of the possible household responses is to give up certain trips entirely or substitute them with shorter ones (for example, shopping in the neighborhood instead of traveling to a far away mall). Although the within-city rent gradient may partially offset direct commuting costs (as suggested by spatial equilibrium within cities á la Alonso, 1964), to the extent that households cannot fully optimize against bad shape through within-city margins, poor city shape will affect city choice and require cross-city compensating differentials. These stem from potential non-pecuniary costs associated with living in a poorly-shaped city, including the disutility from living in less preferable locations so as to avoid long commutes, the sheer displeasure of sitting in traffic, or the disutility of renouncing a trip to avoid this displeasure. These "quality of life" costs are parsimoniously captured in the model by allowing S to reduce θ_T.[14]

Along the same lines, S may affect θ_P, capturing the fact that utilities can be more efficiently delivered through spatial networks in more compact cities. Note that in India public services are primarily funded by states and local taxes have an extremely limited role (Jaitley, 2018), hence θ_P plausibly does not appear in the budget constraint.

Spatial equilibrium requires that indirect utility V be equalized across cities, implying:

$$V(W - p_h H, H, \theta) = \overline{\upsilon}. \tag{2}$$

Embedded in (2) is the intuition that, in equilibrium, wages and rents equalize utility differences. Households implicitly pay for a better θ bundle, including better accessibility or amenities, through a combination of higher rents p_h and lower wages W.

In the production sector, competitive firms optimally choose where to locate and produce the tradeable good according to production function $C = AF(N, K, \overline{Z})$, where A represents a city-specific productivity parameter, N is labor, K is traded capital and $\overline{Z}$ is a fixed supply of non-traded capital. Similar

[14] Modeling the within-city location and travel problem solved by households is notoriously challenging both theoretically and in terms of data requirements (Small and Verhoef, 2007). In Section C in the Appendix I present a drastically simplified version of this problem by embedding a monocentric city (Alonso, 1964) in a setting with constraints to urban shape. The reduced-form predictions of the within-city model are consistent with those of the cross-city framework.

Empirically, pinning down the within-city responses to city shape would require more disaggregated data than what is available for India. However, in Section 9 I offer suggestive evidence on some of these responses, including firm location and work commutes.

to households, firms may benefit from compact city shape through better access to services or because of greater accessibility, which I capture by allowing S to affect A via two components, A_P and A_T. Firms also face a within-city problem where they optimize along various margins, including choosing where to locate within a city. To the extent that they cannot fully optimize against bad shape, S will affect A and their choice of city. Normalizing the price of traded capital to 1, firms' profit maximization problem yields the following zero-profit condition:

$$\pi(W,A,\overline{Z}) = 0. \tag{3}$$

This embeds the intuition that more productive cities must pay higher wages in equilibrium, as they attract more firms.

Finally, the model features developers competitively producing housing in each city, building over a fixed supply of land. Combining the indifference condition of households, firms, and developers, the model delivers equilibrium population N, wages W, and housing rents p_h in a given city as a function of θ and A:[15]

$$Y = f(\theta,A), Y \in \{N,W,p_h\}. \tag{4}$$

To illustrate the reduced-form predictions of the model, assume that non-compact shape S is negatively affecting households ($\frac{\partial\theta}{\partial S} < 0$) but not firms ($\frac{\partial A}{\partial S} = 0$). This would be the case if, for example, households located in non-compact cities faced longer commutes, or were forced to live in a less preferable location so as to avoid long commutes, while firms were unaffected - because of better access to transportation technology, or because of being centrally located within a city. All else being equal, a city with less compact shape should then have a smaller population, higher wages, and lower housing rents. Intuitively, households prefer cities with good shapes, which drives rents up and bids wages down in those locations.

Suppose, instead, that poor city geometry negatively affects both the utility of households and the productivity of firms ($\frac{\partial\theta}{\partial S} < 0, \frac{\partial A}{\partial S} < 0$). This would be the case if firms have worse access to utilities or are prevented from taking full advantage of agglomeration spillovers in non-compact cities. The model's predictions are similar, except that the effect on wages will be ambiguous, given that now both firms and households prefer to locate in compact cities.

Motivated by the model, in Section 6 I begin by examining the reduced-form impacts of city shape on population, wages, and rents. Furthermore, in Section 8 I discuss compensating differentials and provide estimates of the implied willingness to pay for compact shape. The evidence suggests that S affects households' quality of life (via θ) but does not have a meaningful impact on firm productivity (A) in equilibrium. In Section 9 I investigate mechanisms, considering accessibility (θ_T) and service delivery (θ_P), and I find evidence in support of the former channel.

[15] See equarions (B.9), (B.10), and (B.11) in the Appendix for closed-form solutions.

5 Empirical Strategy

In this Section I propose an empirical strategy to take to the data the reduced-form predictions outlined above. To fix ideas, consider city population N as an outcome. Denote the shape of city c in year t as $S_{c,t}$, where higher values denote less compact shapes, and let $area_{c,t}$ be the area of the urban footprint. The equation of interest is:[16]

$$\log(N_{c,t}) = a \cdot S_{c,t} + b \cdot \log(area_{c,t}) + \eta_{c,t}. \tag{5}$$

The main concern in estimating the above relationship is the endogeneity of urban geometry. The observed spatial structure of a city at a given point in time is the result of the interaction of exogenous factors, such as geography, and factors endogenous to population, such as the city's growth rate and policy choices. Examples of policies affecting city shape include master plans, land use regulations, that can promote more or less compact patterns, and investments in road infrastructure, that can generate distinctive patterns of urban growth along transport corridors. This induces a simultaneous correlation between city shape and city size. In general, the sign of the OLS bias will be ambiguous, as the selection effects induced by the endogenous determinants of city shape operate in different directions. Below I provide a qualitative discussion, and in Section D in the Appendix I provide an analytical derivation.

One endogenous determinant of city shape is local institutional capacity. Areas with stronger state capacity tend to have better urban planning and enforcement of master plans, and may be more compact, all else being equal. At the same time, cities with stronger institutional capacity and well-functioning local governments also tend to be more successful and faster-growing cities. This may result in fast-growing cities having better shapes, for reasons unrelated to the value of compactness. This selection effect (denote it as A) would thus tend to generate a negative correlation between non-compact shape and population.

Another kind of selection effect (denote it as B) is due to the fact that population growth may make cities less compact. Mechanically, as cities grow, they tend to deteriorate in shape: intuitively, a city is originally founded in a topographically privileged location, and as it expands over time it will typically extend into terrain that is less preferable. Furthermore, a city experiencing faster population growth may be harder to manage from an urban planning perspective, resulting in more chaotic development. There could also be effects mediated by infrastructural investment: more highways connecting into large and fast-growing cities could lead to sprawl and non-compact development (echoing Baum-Snow, 2007). Finally, large cities may have more fragmented governance as they stretch over multiple administrative units (as Delhi's urban agglomeration, which covers multiple states). This may result in uncoordinated urban planning and more difficult enforcement, all leading to less compact development. All of these effects would tend to generate a positive correlation between non-compact shape and population. The OLS estimate for the impact of bad shape on city population will thus be a combination of the causal impact (via utility-equalizing population flows), gross of selection effects of type A and B discussed above.

[16] This is the empirical counterpart of reduced-form equation (B.9) derived in the Appendix.

5.1 Instrumental Variable Construction

In order to address these concerns, I employ an instrumental variables approach that exploits both temporal and cross-sectional variation in city shape. Intuitively, my identification relies on plausibly exogenous changes in shape that a city undergoes over time, as a result of encountering topographic obstacles along its expansion path. More specifically, I construct an instrument that isolates the variation in urban shape driven by topography and mechanically predicted urban growth. Such instrument varies at the city-year level, incorporating the fact that cities hit different sets of topographic obstacles at different stages of their growth.

To operationalize this identification strategy, I instrument the *actual* shape of the observed footprint at a given point in time with the *potential* shape the city can have, given the geographic constraints it faces at that stage of its predicted growth. Specifically, I consider the largest contiguous patch of developable land, i.e. land not occupied by a water body nor by steep terrain, within a given predicted radius around each city. I denote this contiguous patch of developable land as the city's "potential footprint". I compute the shape indicator of the *potential* footprint and use it as an instrument for the shape of the *actual* urban footprint. What gives time variation to this instrument is the fact that the predicted radius is time-varying, and expands over time based on a mechanical model for city expansion. Using predicted growth is important as actual growth would be endogenous.

The procedure for constructing the instrument is illustrated in Figure 5 for the city of Mumbai. Recall that I observe the footprint of a city c in year 1950 (from the U.S. Army maps) and then yearly between 1992 and 2010 (from the night-time lights dataset). I take as a starting point the minimum bounding circle of the 1950 city footprint (Figure 5a). To construct the instrument for city shape in 1950, I consider the portion of land that lies within this bounding circle and is developable, i.e., not occupied by water bodies nor steep terrain. The largest contiguous patch of developable land within this radius is colored in green in Figure 5b and represents what I define as the "potential footprint" of the city of Mumbai in 1950. In subsequent years $t \in \{1992, 1993..., 2010\}$ I consider concentrically larger radii $\widehat{r}_{c,t}$ around the historical footprint, and construct corresponding potential footprints lying within these predicted radii (Figures 5c and 5d).

The projected radius $\widehat{r}_{c,t}$ is obtained by postulating a mechanical model for city expansion in space, that is based on a projection of the city's historical (1871-1951) population growth rates. In particular, $\widehat{r}_{c,t}$ answers the following question: if the city's population continued to grow as it did between 1871 and 1951 and population density remained constant at its 1951 level,[17] what would be the area occupied by the city in year t? More formally, the steps involved are the following:

(i) I project log-linearly the 1871-1951 population of city c (from the Census) in all subsequent years, obtaining the projected population $\widehat{pop_{c,t}}$, for $t \in \{1992, 1993..., 2010\}$.

(ii) Denoting the actual population of city c in year t as $pop_{c,t}$, I pool together the 1950-2010 panel of cities and estimate the following regression:

$$\log(area_{c,t}) = \alpha \cdot \log(\widehat{pop_{c,t}}) + \beta \cdot \log\left(\frac{pop_{c,1951}}{area_{c,1950}}\right) + \gamma_t + \varepsilon_{c,t}. \qquad (6)$$

[17] Area is observed in 1950 and matched to Census population data from 1951 to calculate density in 1951.

From the regression above, I obtain $\widehat{area_{c,t}}$, the *predicted* area of city c in year t.[18]

(iii) I compute $\widehat{r_{c,t}}$ as the radius of a circle with area $\widehat{area_{c,t}}$:

$$\widehat{r_{c,t}} = \sqrt{\frac{\widehat{area_{c,t}}}{\pi}}. \tag{7}$$

The circle with radius $\widehat{r_{c,t}}$ from Figures 5c and 5d thus represents the area the city would occupy if it continued to grow as in 1871-1951, with unchanged density, and if the city expanded freely and symmetrically in all directions.

The variation in city shape captured by this time-varying instrument is induced by geography interacted with mechanically predicted city growth. This excludes, by construction, the variation resulting from policy choices. The instrument is also arguably orthogonal to most time-varying confounding factors - such as rule of law or local politics - that may be correlated with both city shape and the outcomes of interest. Focusing on variation induced by topography avoids the selection effects of type A discussed above (more successful cities attracting better planners). Using variation from the mechanical model for city expansion, instead of the city's actual growth, helps avoid selection effects of type B discussed above (faster growing cities deteriorating in shape).

Note that city area has to be included in the estimating equation (5) to account for the fact that, mechanically, larger cities are characterized by longer distances. However, including actual area as a control is problematic: a city's expansion in land area will reflect population growth, part of which will be a response to changes in shape. To avoid this simultaneity, I employ projected historical population as an instrument for city area, mirroring the approach I follow in the construction of the shape instrument. By predicting city area using historical population 1871-1951, I isolate the variation in city area that is driven by a city's fundamentals, and exclude the variation induced by recent responses to city shape.

5.2 Estimating Equations

With this instrument in hand, I proceed to estimate the following specification:

$$\log(Y_{c,t}) = a \cdot S_{c,t} + b \cdot \log(area_{c,t}) + \mu_c + \rho_t + \eta_{c,t} \tag{8}$$

where $Y_{c,t}$ is the outcome of interest, $S_{c,t}$ is the shape of the *actual* footprint, $area_{c,t}$ is the area of the urban footprint, and μ_c and ρ_t are city and year fixed effects.

Two regressors are endogenous: the regressor of interest $S_{c,t}$ and the control variable $\log(area_{c,t})$. The corresponding instruments are $\widetilde{S_{c,t}}$, the shape of the potential footprint, and $\log(\widehat{pop_{c,t}})$, the same projected historical population used in the model for urban expansion (described above).

This leads to the following two first-stage equations:

$$S_{c,t} = \sigma \cdot \widetilde{S_{c,t}} + \delta \cdot \log(\widehat{pop_{c,t}}) + \omega_c + \varphi_t + \theta_{c,t} \tag{9}$$

[18]As a robustness check, I also consider an alternative implementation of the instrumental variables approach, that postulates a common rate of expansion for all cities, equivalent to the average rate of expansion across all cities in the sample. This alternative approach is detailed in the Appendix, Section E, and discussed in Section 7 among the other robustness checks.

and

$$\log(area_{c,t}) = \alpha \cdot \widetilde{S_{c,t}} + \beta \cdot \log(\widehat{pop_{c,t}}) + \lambda_c + \gamma_t + \varepsilon_{c,t}. \tag{10}$$

Since many of my outcomes are not available on a yearly basis and the year-to-year variation in city shape is limited, throughout the paper I present most results as long differences, yielding the following estimating equation:

$$\Delta \log(Y_c) = a \cdot \Delta S_c + b \cdot \Delta \log(area_c) + \eta_c \tag{11}$$

where the long differences (denoted by Δ) are taken over 2010-1950 unless otherwise indicated. The corresponding first-stage estimating equations are long difference versions of equations (9) and (10) above. This approach differences out time-invariant city characteristics.

Finally, a small subset of outcomes are available for a single cross-section, in which case cross-sectional versions of (8), (9), and (10) are estimated.

6 Main Results

In this Section I discuss first-stage estimates of the relationship between predicted and actual city shape and the impact of city shape on population, wages, and rents.

6.1 First Stage

Table 2 presents results from estimating the two first-stage equations, relating city shape (in odd columns) and area (in even columns) to the geography-based instrument described above and to projected historical population. Potential shape, as determined by topographic obstacles, is indeed predictive of actual city shape, with F statistics above conventional levels. Columns 1 and 2 present results from the baseline long difference specification employed throughout the paper, where I consider changes in city shape and area between 2010 and 1950. Columns 3 and 4 report the equivalent specification, but in panel format, using all of the data from intermediate years as well, which yields very similar results and a slightly stronger first stage.

This exercise is of inherent interest as it sheds light on the land consumption patterns of Indian cities as a function of their geography. Interestingly, the area of the *actual* footprint appears to be positively affected by the shape of the *potential* footprint (columns 2 and 4). While this partly reflects the mechanical correlation between shape and footprint area, it also suggests that the topographic configurations that make cities less compact may also make them expand more in space. This could be because topographic constraints induce a leapfrog, more land-consuming development pattern, or could reflect an inherent difficulty in planning for parsimonious land use in constrained contexts. This also clarifies that "constrained" cities in this context should not be thought of as land-scarce in absolute terms, but rather cities where growth has to occur around topographic obstacles.

6.2 Population

Table 3 reports estimates of the impact of city shape on population, the main outcome of interest. I estimate long-difference equation (11) by IV in column 1 and OLS in column 2. The results are similar in the panel version (equation (8)), reported in the Appendix.[19]

The IV estimates show that, as cities become less compact, conditional on area, their population growth declines. The magnitudes are best understood in standardized terms. Recall that higher values of the shape index imply longer distances and less compact geometry. A one-standard deviation increase in normalized shape for the average-sized city in the panel (which has radius 4.8 kilometers) corresponds to roughly 360 meters. Holding constant city area, this increase in the average distance between points in the city is associated with a 3.5% decline in population. Through the lens of the model, this is consistent with households valuing compact city layouts as they choose across cities. To the extent that households value compact city shape, spatial equilibrium forces coupled with national population growth will result in population flowing into compact cities at a faster rate.

Conversely, the OLS results in column 2 indicate a positive correlation between city shape and population growth. This confirms the descriptive patterns highlighted in Section 3: in equilibrium, faster growing cities are cities that grow into more disconnected shapes. Specifically, the 0.022 OLS coefficient implies a deterioration in shape of 450 meters for a one percent increase in population, . The discussion in Section 5 suggests potential channels through which this positive selection effect may operate: more difficult urban planning or governance, urban growth occurring along transit corridors, and the mechanical tendency of cities to expand into less favorable terrain.

These results are robust to employing more or less restrictive definitions of urban areas. As discussed in Section 3, delineating urban areas using night-time lights requires setting a luminosity threshold above which a pixel is considered "urban". In Table A5 I provide results using a less restrictive threshold of 30 (columns 1 through 4) and a more restrictive one of 40 (columns 5 through 8). As expected, the lower threshold detects more urban areas, which end up having larger footprints; conversely, the higher threshold detects fewer urban areas and delineates smaller footprints. Despite differences in sample size and in the areas of cities, the estimated impacts of shape on population are very similar to the baseline ones.

These results are also robust to employing alternative shape indicators. In Tables A6 and A7 I consider different shape metrics, detailed in Section A2 in the Appendix. Again, I find similar results, with less compact cities associated with slower population growth.

6.3 Wages and Rents

Next, I examine the impact of city shape on wages and rents, which in the model provide compensating differentials to households and firms as they allocate across cities. A caveat to the empirical

[19] In Table A4 I show that the first-stage and population results also hold in the full panel of city-years. Columns 1 and 2 show the first-stage using all city-years detected in the night-time lights and not just those in the long-difference sample of cities present in the U.S. Army maps. Columns 3 through 6 show that the IV and OLS impacts of city shape on population are very similar using a panel specification, both in the long-difference sample (columns 5 and 6) and in the full sample (columns 3 and 4).

analyses below is that wages and rents are measured more noisily than population, as discussed in Section 3 (with further details provided in Section A in the Appendix).

In Table 4, I report the IV and OLS relationship between average wages and city shape, providing suggestive evidence that non-compact cities are associated with higher wages. The dependent variable is the 2010-1990 long difference in the log urban average of individual yearly wages in the city's district, from the Annual Survey of Industries. The average yearly wage in 2010 was 187 thousand Rupees, at 2015 purchasing power. As discussed in Section 3, the ASI data are available at the district level and the matching between districts and cities is not one to one. I thus provide results for three samples: one including any city that can be matched (columns 1 and 2); one that only includes cities for which there is a one-to-one mapping with a district (columns 3 and 4); and finally a sample including only the top city in each district (columns 5 and 6). Since not all districts can be matched, the sample size is smaller than in the population sample and the first stage is also weaker. The IV estimates tend to be imprecise, with the shape coefficient being only borderline significant in column 3 and significant at the 10% level in column 5, but the qualitative pattern suggests a positive impact of city shape on wages, both in the OLS and in the IV.

Table 5 reports the same set of specifications for housing rents, providing suggestive evidence of lower rents in less compact cities. The dependent variable is the 2008-2006 difference of the log yearly housing rent per square meter, averaged throughout all urban households in the district, from National Sample Survey data. The average yearly rent per square meter in 2006 was 703 Rupees, at 2015 prices. The estimates appear only borderline significant in column 3, with a p-value between 0.10 and 0.15. Again, the lack of precision in the results can partly be attributed to data limitations: measurement error, an imperfect match between cities and districts, loss of power from smaller sample size (which also weakens the first stage), and the limited time variation in the data (drawn from two consecutive rounds of NSS data). However, subject to these caveats, the qualitative pattern that emerges is quite consistent: the impact of disconnected shape on rents is negative in the IV and close to zero in the OLS. These patterns are similar if I exclude from the calculation of average rents the bottom 25% of the rents distribution in a district, which may be more likely to belong to rent-controlled units (see Table A8 in the Appendix).

In Table A9 in the Appendix I show similar qualitative results using an alternative source of rents data, the Indian Human Development Survey (IHDS). The correlation between the IHDS and the NSS data is positive but weak (0.3), reflecting measurement error in both sources. Nevertheless, the sign of the impact of city shape on rents is still negative in the IV and positive in the OLS. In column 3 and 4 of Table A9 the dependent variable is a rent residual, obtained from a hedonic regression of rents on housing attributes provided in the IHDS survey (details are provided in the Appendix). The qualitative pattern of lower rents in less compact cities is preserved.

Taken together, the finding of higher wages and lower rents in non-compact cities is consistent with a compensating differential interpretation. In the model, if compact city shape provides advantages in terms of quality of life or productivity, compact cities will be characterized by higher rents and wages that may be higher or lower depending on whether households or firms value compact shape the most.

To the extent that households value city shape more than firms, they will bid wages up in compact cities. I discuss the magnitudes and interpretation of this compensating differential through the lens of the model in Section 8.

7 Threats to Identification

In this Section I address the main threats to identification, focusing on population as an outcome variable. I begin by discussing concerns related to direct effects of geography, followed by confounding by initial conditions or diverging trends. Finally, I discuss an alternative estimation strategy that employs a single instrument and does not rely on controlling for projected historical population.

7.1 Direct effects of geography

The exclusion restriction for the shape instrument requires that potential shape only affects the outcomes of interest though the constraints that it posits to urban form. One of the major identification threats is that the instrument may be correlated with geographic characteristics that have direct time-varying impacts on the outcomes of interest. For example, the topographic constraints that affect city shape, such as coasts and slopes, may also make cities intrinsically more or less attractive for households and/or firms. Indeed, the literature documents many channels through which physical geography affects local development: amongst others, see Combes et al. (2010), Rosenthal and Strange (2004), and Barr et al. (2011) on geology, density, and agglomeration; Burchfield et al. (2006) on local geography and density; Bleakley and Lin (2012) on coastal configurations and ports; Nunn and Qian (2011) on potato suitability and urbanization; and Nunn and Puga (2012) on ruggedness and local economic development. In particular, the reader may worry about geographic features that have inherent consumption amenity value (e.g. coasts or lakes), production amenity value (e.g. the presence of mineral deposits, or fertile land), or that may impact construction costs (e.g. terrain ruggedness or bedrock depth).

These direct effects of geography could bias the IV results in different directions. For example, if the instrument picked up the effect of coasts and the latter were landscape amenities, the estimated effects of bad shape on population would be biased towards positive values. Conversely, if potential shape were less compact in areas with particularly deep bedrock, the IV impacts of shape on population could be biased towards more negative values, as they would be mediated by higher construction costs in those cities (Barr et al., 2011).

Below, I show that the IV results are unlikely to be driven by confounding effects of these geographic characteristics. As a preliminary step, in Table A10, I show that the instrument is uncorrelated with most of these geographic variables. Each row reports the coefficient from a separate OLS regression. In column 1 I report pairwise correlations between changes in potential shape 2010-1950 and a number of predetermined city characteristics, including elevation, distance from the coast, distance from the nearest river or lake, distance from mineral deposits, terrain ruggedness (capturing slope), bedrock depth (which the literature has linked to high-rise construction costs, population density and ultimately agglomeration), and crop suitability (that may be higher near cities with water bodies). For completeness, in column 2 I provide the same correlations for the projected population instrument (used to control for city area). A description of the controls is provided in Section A4 and summary statistics are reported in Table A3 in the Appendix. Reassuringly, changes in potential shape are uncorrelated with most of these characteristics.[20]

In Table 6, I show that the IV estimates for the impact of city shape on population are robust to controlling for all of the characteristics listed above. I report the same IV specification as in Table 3, column 1, augmented with time-invariant geography controls. This amounts to allowing for differential changes across cities with different geographic characteristics. All point estimates are very similar to the baseline one of -.096, assuaging concerns of confounding.

In Table A11 I provide IV results for different sample cuts. My results are minimally affected by excluding from the sample coastal and mountainous cities, high-ruggedness cities, cities near rivers or lakes, cities with minerals, cities with high bedrock depth and cities with high crop suitability. This is reassuring that the results are not driven by a very peculiar set of compliers: most cities are affected in their shape by the position of topographic constraints, and not only those with particular topographies.

The lack of correlation between the shape instrument and geographic variables such as elevation or distance to the coast may appear surprising, since potential shape is calculated based on constraints stemming from steep slopes and water bodies. Importantly, the instrument's variation does not stem from the generic presence of water bodies or steep slopes, nor from the presence of particularly large obstacles (e.g. a mountain or lake), but rather from the relative position in space of these constraints. In fact, a city could be very "constrained" in terms of share of land lost to bad topography, but may still be able to expand in a compact way. For example, suppose all topographic obstacles are concentrated East of the center of a city. This will not prevent the city from expanding in a relatively compact way on the West side. On the other hand, if a city is surrounded by obstacles in multiple directions, it will have to grow around those obstacles generating a less compact pattern. This is the variation that the instrument is capturing.

7.2 Initial conditions and pre-existing trends

The estimation of the population response to changes in city shape may be confounded by underlying city-specific trends, potentially driven by fundamentals or initial conditions. Controlling for projected historical population growth 1871-1951 through the city area instrument partially addresses these concerns, as it allows changes in city shape to only affect deviations from the city's long-run path. However, there is still a concern of changes in cities' fundamental trends that are not captured by past projected growth. I address these concerns by showing that my baseline results are very similar when controlling for a battery of city characteristics and across various sample cuts. I also present falsification tests using instrument leads and future changes.

In Table A12 I consider potential diverging trends by initial conditions, extending the tests of Table 6 to include non-predetermined characteristics as controls. The first-stage and IV estimates

[20] There is a weak correlation between changes in potential shape and distance from mineral deposits. This is addressed in Table A11 by showing that the results are robust to excluding cities near mineral deposits.

are qualitatively similar to the baseline ones. In columns 1 through 3 I control for initial shape at the beginning of the sample, allowing cities that start out with different constraints to evolve along different paths. Not surprisingly, in this more demanding specification instruments are weaker, but the qualitative impacts of city shape are preserved. In columns 4 through 6 I control for direct British rule, which may be associated with particular city management or urban planning approaches (Baruah et al., 2017), and in columns 7 through 9 I include a capital city dummy, motivated by the strong correlation between capital cities and non-compact shape highlighted in Table A2. Again, results are very similar.

In Table A13 I present additional sample cuts based on non-predetermined characteristics. In columns 1 through 9 I consider cities with different population growth patterns and show that the first-stage and IV results are very similar when excluding particular sets of cities. In columns 1 though 3 I exclude cities that at any point during the years in the panel have experienced negative population growth from one year to the other. In columns 4 though 6 I exclude fast-growing cities, defined as cities whose 2011-1951 growth rate was in the top 10th percentile. In columns 5 and 6 I exclude the slow-growing cities, defined as the bottom 10% growers. Again, this is reassuring that the compliers are not a peculiar set of cities. Along the same lines, in columns 10 through 12 I exclude the top 10% most constrained cities. The definition of "constrained" refers to the share of land within the 2010 city radius that is lost to topographic constraints.[21]

To further assuage the concern of underlying pre-trends, in Table 7 I provide a falsification test regressing changes in outcomes on instrument leads and future instrument changes. Reassuringly, past changes in population are not predicted by future values of the instrument. Specifically, in column 1 I regress 2001-1951 population changes on 2005 and 2010 instruments; in column 2 I consider 2001-1991 population changes as a dependent variable instead. In column 3 I regress 1991-1951 population changes on instruments measured in 1995 and 2000. Finally, in column 4 I regress 1991-1951 changes in shape on 2010-2001 changes in the instruments. None of the shape coefficients are statistically different from 0, and many of the point estimates are positive (the opposite sign of the main effects). Table A14 in the Appendix includes a similar test for rents and wages, which shows no pattern that would be suggestive of pre-trends correlating with the shape instrument.

7.3 Single-instrument specification

Employing projected historical population in the construction of the shape instrument and as an instrument for area may raise identification concerns. The identifying assumption is that projected historical population predicts actual city area and shape, but does not affect current population and other outcomes other than through "deep fundamentals" uncorrelated with shape. A violation of the exclusion restriction may arise if historical population growth 1871-1951 not only predicted current population and city expansion through fundamentals, but also affected it through shape itself. This could be the case if city shape also followed a long-run trend and the 1871-1951 population growth

[21] This is similar to the measure that Saiz (2010) relates to housing supply elasticity. This robustness check suggests that the results are not driven by particularly constrained (and thus potentially more supply-inelastic) cities. In future work, a richer model could characterize the impacts of topographic constraints and shape on housing supply elasticity, to fully disentangle them from those driven purely by geometry. I discuss this in Section B in the Appendix.

partially responded to it.

These concerns can be assuaged by an alternative identification strategy that does not rely on using projected historical population at any stage. This alternative approach, detailed in Section E in the Appendix, involves normalizing both sides of equation (4) by city area. The ensuing estimating equation has population density as an outcome and normalized city shape as the only explanatory variable, which is treated as endogenous. The corresponding instrument is a normalized version of potential shape, based on topographic obstacles encountered along a city's predicted expansion path. However, in this alternative version, the city's predicted expansion path is not based on historical population growth, but is completely mechanical, based on the average rate of city expansion in the panel.

The results of this estimation are reported in Table 8. Column 1 reports the first stage, showing that potential normalized shape is a strong predictor of actual normalized shape. Column 2 reports IV estimates for the impact of normalized shape on population density, showing that population density declines as normalized shape deteriorates. The magnitudes are consistent with the estimates from the baseline specification: as normalized shape deteriorates by one standard deviation, population density declines by approximately one standard deviation.

8 Compensating Differentials and Willingness to Pay

In this Section I provide an interpretation of the reduced-form results on wages and rents from Section 6 through the lens of the model. In a Rosen-Roback framework, higher real wages in disconnected cities can be interpreted as the implicit premium that households pay in order to live in cities with more compact shapes. To calculate households' willingness to pay for city shape, I begin by expressing households' indifference condition (2) as a log-separable function, which can be derived from a Cobb-Douglas utility function:

$$\log(W) - \alpha \log(p_h) + \log(\theta) = \log(\overline{\upsilon}) \tag{12}$$

where α is the share of housing in consumption.

Differentiating (12) with respect to S provides a way to quantify the extent to which S affects indirect utility via θ:

$$\frac{\partial \log(\theta)}{\partial S} = \alpha \frac{\partial \log(p_h)}{\partial S} - \frac{\partial \log(W)}{\partial S}. \tag{13}$$

The marginal willingness to pay for a unit improvement in S equals the difference between the semi-elasticity of housing prices to S, weighted by the share of housing in consumption α, and the semi-elasticity of wages. As an empirical counterpart of (13), I estimate the following:

$$\widehat{\lambda_\theta} = \alpha \widehat{B_P} - \widehat{B_W} \tag{14}$$

where $\widehat{B_P}$ and $\widehat{B_W}$ are estimates of the reduced-form impact of city shape S on, respectively, log rents and wages. To calibrate α, I compute the share of household expenditure devoted to housing for urban households, according to the NSS Household Consumer Expenditure Survey data in my sample. This figure amounts to 0.16.[22]

[22]While this figure may seem low, it is consistent with the evidence from other developing countries (Chauvin et al.,

Estimates of $\widehat{\lambda}_\theta$, obtained from pooling the IV regressions of Tables 4 and 5, are reported at the bottom of Table 5. The willingness to pay for a one kilometer improvement in city shape ranges between 0.13 and 0.17 log points, depending on the specification, with p-values between 0.04 and 0.12. In standardized terms, this implies a willingness to pay between 4.7 and 6% for a one standard deviation improvement in city compactness, corresponding to an increase in the average within-city distance of approximately 360 meters.[23] Note that relying on OLS, as opposed to IV estimates of $\widehat{B_P}$ and $\widehat{B_W}$ would yield smaller willingness-to-pay estimates, ranging between 1 and 2%.

8.1 Discussion

The positive estimated willingness to pay for good shape $\widehat{\lambda}_\theta$ can be interpreted as evidence that households view compact shape as affecting their quality of life as they evaluate the trade-offs associated with different cities. Through a more structural lens, this estimate can also be useful to sign and bound potential welfare effects of city shape.

First, λ_θ can be viewed as an upper bound for welfare effects of deteriorating shape, to the extent that reality is somewhere in between a scenario with infinitely elastic or infinitely inelastic supply of urban dwellers (Donaldson and Hornbeck, 2016). With a fixed total urban population at the country level, equilibrium indirect utility $\overline{\upsilon}$ will increase everywhere if θ increases in one city. In the Cobb-Douglas case, λ_θ coincides with the welfare impact of a one unit improvement in shape in all cities. The assumption of a fixed total population is extreme, as many migrants into cities are coming from the countryside rather than reallocating across cities. The alternative extreme assumption is that of a perfectly elastic supply of migrants to cities, with indirect utility being pinned down by a reservation utility in the countryside. In this scenario, any improvement in θ will result in larger city populations but no welfare change, which provides a lower bound of zero for the welfare effect of city shape.

Furthermore, in a richer model with heterogeneous households, the Rosen-Roback indifference conditions will hold for the marginal household, but there will be welfare impacts on inframarginal households. Intuitively, the latter will not be perfectly compensated for bad shape through higher real wages and their utility will be affected by changes in θ. The welfare impacts on those households will depend on the relative elasticity of labor and of housing supply, with a lower local elasticity of labor implying a larger household incidence (Moretti, 2011). Analyzing distributional impacts of deteriorating city shape requires a richer model incorporating landlords and tenants as well as heterogeneous incomes and/or migration costs and is left for future research.

The calculation of λ_θ is subject to a number of caveats, that could be addressed in future work. While the notion of compensating differentials based on rents and wages is very general, the calculation above relies on Cobb-Douglas functional form, implying homothetic preferences and a constant housing expenditure share. Which functional form best describes housing expenditure in a developing country setting is an open question.

2017). Employing the IHDS data as an alternative source I find a similar number.

[23]In order to evaluate this magnitude, this figure could be compared to estimates of the value of other amenities. However, no such estimates are available for India. As a reference, covering 360 additional meters on foot twice a day takes about 9 minutes, or 2% of an 8-hour working day.

Second, the model does not allow for heterogeneous agents to sort into locations based on their preferences or skills. This particularly affects the interpretation of the estimated impact of shape on wages. The latter may reflect sorting and differences in the skill composition of the workforce (Combes et al., 2008), which I cannot control for given the information in the ASI data.[24] The estimated compensating differentials should be thus thought of as an underestimate of true equalizing differences for those with a strong preference for compact layouts, and an overestimate for those with weak preferences.

Third, the model assumes that the housing supply elasticity is the same across cities. Allowing for heterogeneity across cities would affect the magnitude of the response of rents and wages: to the extent that good shape positively affects household utility, in more inelastic cities the impacts on population and wages would be attenuated and the impact on rents would be amplified (in absolute terms).

Furthermore, the model implicitly assumes that S only affects θ. A richer model, providing a micro-foundation for how city shape affects households and firms, may also allow for city shape to affect other objects, including the elasticity of housing supply.

Finally, the model does not allow for externalities. With congestion, $\widehat{\lambda_\theta}$ would understate the true willingness to pay for compact shape, as it would be estimated gross of equilibrium congestion effects.

8.2 Implied Productivity Impacts

Next, I consider the implied productivity impacts of city shape on firms, $\frac{dA}{dS}$. Similar to the calculation for households, I begin by expressing firms' indifference condition (3) under the assumption of a Cobb-Douglas production function:

$$(1-\gamma)\log(W) = (1-\beta-\gamma)(\log(\overline{Z}) - \log(N)) + \log(A) + \kappa_1 \tag{15}$$

where parameters β and γ represent the shares of labor and tradeable capital in a Cobb-Douglas production function.

Totally differentiating (15) with respect to S allows to pin down the effect of S on productivity as:

$$\frac{\partial \log(A)}{\partial S} = (1-\beta-\gamma)\frac{\partial \log(N)}{\partial S} + (1-\gamma)\frac{\partial \log(W)}{\partial S}. \tag{16}$$

I estimate the empirical counterpart of the above as

$$\widehat{\lambda_A} = (1-\beta-\gamma)\widehat{B_N} + (1-\gamma)\widehat{B_W}, \tag{17}$$

where $\widehat{B_N}$ is the estimated reduced-form impact of shape on population.

Setting β to 0.4 and γ to 0.3 (as in Glaeser, 2008), IV-based estimates of $\widehat{\lambda_A}$ range from -0.12% (under the most conservative point estimates of $\widehat{B_W}$) to 0.8% for a one standard deviation deterioration in city shape. In the pooled specification, none of the estimates are statistically different from zero (with p-values ranging from 0.5 to 0.9). [25]

[24] In Section 9 I discuss evidence that compact cities have a larger share of slum dwellers, which may suggest sorting of lower-skill workers into compact cities. However, the wages of slum dwellers are unlikely to be driving my results, as the ASI data that I employ only covers the formal sector.

[25] Utilizing OLS, as opposed to IV estimates, yields positive and statistically significant impacts of bad shape on productivity in the 1.1-1.4% range, in line with the OLS pattern of more disconnected cities being the larger and plausibly most

These estimates appear very small, suggesting that city shape does not affect firms in the cross-city equilibrium. This does not indicate that a city's layout is *ex ante* irrelevant for firms. Rather, the interpretation is that firms do not require a compensation for poor city geometry through factor prices, whereas households do. Put differently, in equilibrium, firms may be able to optimize against "bad" shape, in a way that households cannot. This may be related to the relative location of households and firms within cities. This hypothesis is explored in Section 9, where I investigate how firms respond to city shape in their location choices within cities, by looking at the spatial distribution of employment.

9 Mechanisms and Heterogeneous Effects

The urban planning literature emphasizes two main channels through which the compactness of city layouts may affect households and firms: transit (θ_T and A_T in the model) and public service delivery (θ_P and A_P). Shorter distances improve accessibility and may facilitate the provision of infrastructure, as well as ease the delivery of services provided through spatial networks, such as water and electricity (Bertaud, 2004; Cervero, 2001). In what follows I provide evidence on both channels, showing that accessibility is plausibly the main mechanism. I also discuss heterogeneous effects of city shape shedding further light on the way in which disconnected cities operate.

9.1 Accessibility and Transit

In this Section I begin by discussing the heterogeneous impacts of city shape as a function of a city's infrastructure and ease of transit, taking infrastructure as given. I then discuss the equilibrium relationship between city shape and infrastructural provision. Finally, I provide suggestive evidence on the location of firms and commutes to work.

Heterogeneity by infrastructure

Recall that my shape indicator is based on Euclidean distances between points in a city, abstracting from the road network and transport technology. All else being equal, a well-functioning road network should mitigate the impacts of bad shape: for example, a disconnected city with a fast highway may ultimately be very accessible.

In Table 9 I show that the negative impacts of shape on population are indeed mitigated in cities with better-functioning within-city transit. I augment my baseline IV specification with interactions between shape and a number of transit-related variables that capture how easy commutes are in the city: road length (columns 1 through 3), indices capturing the functionality of the road network from Akbar et al. (2018, 2019) (columns 4 and 5),[26] and availability of cars (columns 6 through 8). The sources and construction of each variable are detailed in Section A5 in the Appendix. All interaction terms yield positive coefficients: while bad shape tends to reduce population growth, the effects are attenuated for

productive cities.

[26]The grid conformity index (column 4) measures the extent to which a city's current road network is laid out as a regular grid; it correlates with better vehicular mobility. The proximity index (column 5) is a city-level measure of distance accessibility capturing how easy it is to reach shopping centers, train stations, restaurants, and other amenities within a city.

cities where commutes are plausibly easier due to better road infrastructure and motorized means of transportation.

While these results are highly suggestive, I caution that they can only be interpreted causally if one takes infrastructure as given. In reality, infrastructure provision is simultaneously determined with urban shape (as discussed in detail below) and affected by city income, which may confound the estimation of the interaction terms in Table 9. I mitigate endogeneity concerns through various approaches: I consider lagged interaction variables (in columns 2 and 7) and employ state-level, instead of city-level variables (columns 3 and 8). One specific concern is that cities with better infrastructure tend to be higher-income, more successful cities. To assuage this source of confounding, in Table A15 I replicate the estimation of Table 9 additionally controlling for the number of banks in 1981 as a proxy for city income. While this is admittedly an endogenous control, including it does not change the main estimates, suggesting that the interactions with infrastructure are not solely capturing income differences.

City shape and provision of infrastructure

Below I elaborate on the equilibrium relationship between city shape and infrastructure. Urban infrastructure is jointly determined with city shape and the sign of the reduced-form relationship between the two is *a priori* ambiguous. On the one hand, infrastructural provision is an endogenous response to changes in a city's layout: as built-up areas expand, the road network also tends to expand to service these areas. At the same time, infrastructure is a co-determinant of a city's layout: new built-up areas often arise around transit corridors. This two-way relationship tends to generate a *positive* correlation between bad shape and urban road length. On the other hand, topographic obstacles that lead to disconnected shape may also increase the costs of providing infrastructure, resulting in a *negative* correlation between bad shape and urban road length. Similar arguments can be made regarding road quality.[27]

In Table A16, panel A, I empirically examine the equilibrium relationship between city shape and road length. I find that disconnected cities tend to have a shorter road network, conditional on city area, both in absolute and in per capita terms. In columns 1 through 8 I consider roads in 2019 as measured from Openstreetmap. I report estimates from a cross-sectional version of my benchmark IV and OLS specification, where regressors are defined in 2010 levels. Bad shape is associated with shorter total length of roads (columns 1 and 2) and motorways (columns 3 and 4). A one standard deviation in normalized shape (approximately 360 additional meters) is associated with a 6% shorter road network (column 1). The pattern is similar when considering road length per 2011 population (columns 5 through 8). Although precision varies, these results are qualitatively similar in the OLS (even columns) and in the IV specifications (odd columns). In columns 9 through 12 I consider a specification in changes, where the dependent variable is the difference between 2019 Openstreetmap road length and 1981 urban road length from the Census. While weaker, the negative IV estimates

[27] Bad topography may increase the cost of maintaining or upgrading roads, resulting in lower road quality in cities with poor shapes. At the same time, planners may choose to compensate for poor accessibility by investing in road quality, for example increasing the number of lanes of the main city's artery. Thus, the relationship between city shape and infrastructure quality is *a priori* ambiguous.

confirm the pattern highlighted thus far. The OLS estimates are positive, in line with the spurious correlation between city growth and deteriorating shape. Taken together, these results suggest that as cities expand to become more disconnected, the road network is not keeping up, plausibly due to higher costs of providing infrastructure in topographically constrained settings. Thus, poor city shape may hurt accessibility not only directly, but also by making infrastructural provision more costly.

For the interested reader, in panel B of Table A16 I consider indexes related to the internal functioning of city transit from Akbar et al. (2019). Some of the evidence points to moderately worse mobility in disconnected cities, as measured by within-city transit speeds. In Section A5 in the Appendix I provide details on the construction of these indexes and caveats to the interpretation of the corresponding regression results.

Commuting and within-city responses to city shape

The evidence provided thus far points to transit accessibility as one of the key channels through which city shape affects households. The question may then arise on how city shape maps to commuting behavior. The theoretical prediction is ambiguous: all else being equal, disconnected city shape is associated with lower accessibility and higher *potential* costs of travel within the city. However, *realized* commuting costs may be higher or lower, depending on households' elasticity of demand for trips and on the endogenous location of employers and retailers within the city. As a city becomes more disconnected, households may respond through different margins: one is to incur longer commutes, but others include locating closer to one's job and giving up some trips entirely. For example, they may choose to shop in their neighborhood instead of taking a lengthy trip to their preferred mall. This would result in shorter, rather than longer realized commutes in disconnected cities. Moreover, firms may respond to deteriorating shape by dispersing throughout the city or forming new business districts to be closer to their workers or clients: if a city becomes more polycentric as it becomes more disconnected, commutes should also shorten.[28]

The lack of systematic data on Indian households' location patterns and travel behavior limits the scope for investigating the within-city responses to poor accessibility in a conclusive way. With this in mind, I provide two pieces of evidence on the endogenous responses of firms and workers to city shape, examining the location of firms within cities and households' distance to workplace.

In Table 10, columns 1 and 2, I consider the clustering of firms within cities and show that disconnected cities do not have more dispersed employment. Specifically, I use street addresses and reported employment of productive establishments from the 2005 Economic Census to detect employment subcenters using the approach developed by McMillen (2001).[29] Employment subcenters are identified as locations that have significantly larger employment density than nearby ones and that have a significant impact on the overall employment density function in a city. Details on the data and the procedure are

[28]Models of endogenous subcenter formation emphasize firms' trade-off between a centripetal agglomeration force and the lower wages that accompany shorter commutes in peripheral locations. See Anas et al. (1998) for a review of the literature on polycentricity.

[29]I geo-code the street addresses of productive establishments covered in the 5th Economic Census using Google Maps. I retrieve consistent coordinates for approximately 240 thousand establishments in about 190 cities.

provided in the Appendix, in Sections A6 and F, respectively. I estimate a cross-sectional version of the benchmark OLS and IV specification with the log number of employment sub-centers as a dependent variable, for year 2005. Subject to the limitations of cross-sectional inference and small sample size, less compact cities have, if anything, fewer subcenters, a pattern found both in the IV and the OLS.

This is consistent with the interpretation that, as cities grow into more disconnected shapes, firms continue to cluster in a few locations within a city, plausibly so as to take advantage of agglomeration, and they leave it to workers to bear the costs of longer commutes.[30] This is also in line with the with the findings discussed in Section 8 that poor shape has meaningful impacts for households, but has negligible impacts on firms in equilibrium.

This interpretation also suggests that disconnected cities should be characterized by longer trips to work in equilibrium. Absent commuting data at the city level, in columns 3 through 6 of Table 10 I consider a noisy district-level proxy for work commute length derived from the 2011 Census. The latter provides a breakdown of workers in a district by travel mode and reported distance to work, by coarse bins (0-1, 2-5, 6-10, 11-20, 21-30, 31-50, and above 50 kilometers). I calculate weighted average distance to work separately for workers commuting by car and on foot. Additional details on the construction of this variable are provided in Section A4 in the Appendix. I report a cross-sectional version of the benchmark IV and OLS specifications, for year 2010. In the OLS (columns 4 and 6), poor shapes are associated with longer commutes by car and shorter commutes on foot. This could point to two heterogeneous kinds of responses to deteriorating shape: those with cars endure longer commutes, whereas those without cars choose locations of work and/or residence that are closer to one another. Plausibly, many of the workers who report walking to work are employed in local informal jobs, which may be an alternative to formal employment in inaccessible parts of the city. I caution that this result is not robust as the corresponding IV estimates (columns 3 and 5) are small and insignificant (albeit with the same sign). Noisy results are to be expected given the inherent limitations in the data: the unit of observation is the district (larger than the city) and the distance bins are probably too coarse to capture differences in commuting length in medium and smaller cities.[31]

9.2 Public Services

The second channel through which city shape may affect households and firms is public service delivery. More compact layouts may reduce the cost of providing services such as water, electricity or sewerage, resulting in higher levels of access. In Table 11 I examine the impact of city shape on households' access to electricity and tap water, but I do not find any meaningful effects. Specifically, I consider the 2011-1991 long difference in the log number and share of households with access to electricity (panel A) and tap water on premises (panel B). When considering the total number of households, OLS estimates indicate a positive correlation between disconnected shape and service access (columns 2 and 6), plausibly due to the fact that larger cities tend to have worse shape. However,

[30]In the model, firms compensate workers for these longer commutes with high wages. Firms are able to pay higher wages in more disconnect cities because poorly-shaped cities have smaller population and returns to labor are decreasing (see Section B in the Appendix).

[31]Recall that the median value of the disconnection index is 2.6 kilometers.

when examining the share of households with access (columns 3, 4 and 7 and 8), both OLS and IV estimates are close to zero.

These results appear to run counter the prediction of the urban planning literature and the findings of Baruah et al. (2017), who find worse service access in African cities that are more sprawled. However, one reason may be that service access in urban India is quite high to begin with: in 1991, the shares of households with access to electricity and tap water were 82% and 70% respectively. Furthermore, the urban planning argument may be more relevant for services delivered along a centralized grid, while electricity and water access in urban India is also granted through decentralized means such as small-scale private service providers (Kariuki and Schwartz, 2005). These decentralized solutions may provide a way around the difficulty of servicing the more disconnected parts of a city.

9.3 Slum Population

Complementary to the investigation of the impact of city shape is the question of which types of households bear the costs of poor city shape. On the one hand, compact cities may be more favorable to the poor because they may offer better connectivity to jobs and services. However, lower real wages in compact cities may reduce the housing floor space that the poor can afford and price them out of the formal market (Bertaud, 2004).

While I cannot systematically observe household income, I examine the share of slum dwellers from the Census. In Table 12 I show that cities with less compact shapes have overall fewer slum dwellers, both in absolute terms (columns 1 and 2) and relative to total population (columns 3 and 4). The dependent variables are 2011-1981 long differences in the log number and share of slum households as identified by the Census and the regressors are defined as 2010-1950 long differences. Results are similar in the IV and OLS specifications. Two interpretations are possible. The first is that higher equilibrium rents in compact cities are forcing more households into sub-standard housing. The second relates to sorting of poorer migrants into cities with more compact shapes, possibly because of their lack of individual means of transport and consequent higher sensitivity to commute lengths.[32]

9.4 Land Use Regulations and City Shape

Taken together, the evidence presented in this paper suggests that poor city shape affects household location choices and potentially their quality of life. Given that most cities cannot expand radially due to their topographies, an important question arises on the role of policy and on what kind of land use regulations best accommodate city growth. Below I provide evidence on the interactions between land use regulations, urban growth, and city shape by focusing on a controversial regulatory tool: Floor Area Ratios (FARs).

FARs are restrictions on building height expressed in terms of the maximum allowed ratio of a

[32]These results may raise concerns related to the interpretation the wages results from Section 6.3: lower wages in more compact cities may be driven by low-productivity workers disproportionately locating in these cities, consistent with my findings on slum dwellers. Recall, however, that my wage sample covers the formal sector only and is therefore unlikely to include a large share of slum workers.

building's floor area over the area of the plot on which it sits. Higher values allow for taller buildings. The average value of FARs in my sample is 2.3, a very restrictive figure compared to international standards. Previous work has linked conservative FARs in Indian cities to suburbanization and sprawl (Sridhar, 2010; Bertaud and Brueckner, 2005).

In Table 13, I show that restrictive FARs lead to less compact city shapes. I employ data on FARs in 55 Indian cities as of 2005, from Sridhar (2010).[33] I report the two first-stage equations, linking potential shape and projected historical population to city shape and area, augmented with interactions between each of the two instruments and FARs. Given the small number of cities, in order to leverage more time variation in the data, I present a panel version of the two first-stage equations, similar to columns 3 and 4 in Table 2. The interaction between projected population and FARs has a negative impact on city shape (column 1) and city area (column 2). This suggests that cities with laxer FARs may expand less in space (consistent with Sridhar, 2010), and may expand in a more compact fashion than their projected growth would imply. A one standard deviation increase in FARs (0.6) is associated with an absolute reduction in the shape index of roughly 1 kilometer for each percent increase in projected population.

In other words, higher FARs may slow down the deterioration in city shape that fast city growth entails: if growing and topographically constrained cities are allowed to grow vertically, they will not expand horizontally as much and will plausibly remain more compact. These findings are particularly important for the larger cities in India, as they are those with the most pronounced natural tendency to deteriorate in shape and also those with the most restrictive FARs (Sridhar, 2010).

10 Conclusion

In this paper I examine the causal economic implications of city shape in the context of India, exploiting variation in urban form driven by topography. Embedding city shape in a classic urban economics model, I connect the notion of geometric city compactness with that of spatial equilibrium across cities, and provide novel causal evidence that city compactness affects urbanization patterns. Less compact urban layouts, conducive to longer within-city distances, are associated with lower quality of life and potential welfare costs for households, primarily driven by worse accessibility. This is particularly important for the largest cities in India, that have a tendency to become less compact over time.

As India prepares to accommodate an unprecedented urban growth in the next decades, the challenges posed by urban expansion are gaining increasing importance. On the one hand, policy makers are concerned about the perceived harms of haphazard urban expansion, including sprawl and limited urban mobility (World Bank, 2013). On the other hand, existing policies, especially land use regulations, are viewed as potentially distortive of urban form (Glaeser, 2011; Sridhar, 2010). This paper contributes to informing the policy debate on both fronts. Although this study focuses on geographic obstacles (which are mostly given) in order to gain identification, there is a range of policy options to prevent the deterioration in connectivity that fast city growth entails and to mitigate the negative impacts on bad shape, for example by improving urban mobility. My findings also suggest that urban

[33] Given that FARs are updated infrequently, these mid-2000s data are a reasonable proxy for FARs in place throughout the sample period.

connectivity can be indirectly improved by promoting more compact development through land use regulations: restrictive FARs, among the most controversial of such regulations, result in less compact footprints. This suggests that any distortive effects on urban morphology should be accounted for when evaluating the costs of land use regulations. A number of other urban planning practices and regulations currently in place in Indian cities have been viewed as conducive to sprawl (Bertaud, 2002) and could be explored in future work.[34]

In future research, it would be interesting to provide a richer theoretical microfoundation for the effects of city shape on the behavior of households and firms, shedding further light on the channels through which city shape matters. A richer model could also be used to investigate heterogeneous distributional effects and to pin down welfare consequences. More disaggregated data at the sub-city level will be required to empirically investigate these ramifications.

References

[1] Akbar, P., V. Couture, G. Duranton, and A. Storeygard (2018), "Accessibility in Urban India", working paper.

[2] Akbar, P., V. Couture, G. Duranton, and A. Storeygard (2019), "Mobility and Congestion in Urban India", working paper.

[3] Alonso, W. (1964), *Location and land use, Cambridge: Harvard University Press.*

[4] Anas, A., R. Arnott, and K. A. Small (1998), "Urban Spatial Structure", Journal of Economic Literature, 36 (3), 1426-1464.

[5] Angel, S., J. Parent, and D. L. Civco (2009), "Ten Compactness Properties of Circles: A Unified Theoretical Foundation for the Practical Measurement of Compactness", The Canadian Geographer, 54 (4), 441-461.

[6] Baragwanath-Vogel, K., G. Hanson, R. Goldblatt, and A. K. Khandelwal (forthcoming), "Detecting Urban Markets with Satellite Imagery: An Application to India", Journal of Urban Economics.

[7] Barr, J., T. Tassier and R. Trendafilov (2011), "Bedrock Depth and the Formation of the Manhattan Skyline, 1890-1915", Journal of Economic History, 71 (4), 1060-1077.

[8] Baruah, N., V. J. Henderson, and C. Peng (2017), "Colonial Legacies: Shaping African Cities", SERC Urban and Spatial Programme Discussion Paper.

[9] Batty, M. (2008), "The Size, Scale, and Shape of Cities", Science, 319 (5864), 769-771.

[10] Baum-Snow, N. (2007), "Did Highways Cause Suburbanization?", Quarterly Journal of Economics 122 (2), 775-805.

[11] Baum-Snow, N., L. Brandt, J. V. Henderson, M. A. Turner and Q. Zhang (2017), "Roads, Railroads and Decentralization of Chinese Cities", Review of Economics and Statistics, 99(3), 435-448.

[12] Bento, A., M. L. Cropper, A. M. Mobarak, and K. Vinha (2005), "The Effects of Urban Spatial Structure on Travel Demand in the United States", Review of Economics and Statistics, 87 (3), 466-478.

[34]Examples include: the Urban Land Ceiling Act, which has been claimed to hinder intra-urban land consolidation; rent control provisions, which prevent redevelopment and renovation of older buildings; regulations hindering the conversion of land from one use to another; and, more generally, complex regulations and restrictions in central cities, as opposed to relative freedom outside the administrative boundaries of cities.

[13] Bertaud, A. (2002), "The Economic Impact of Land and Urban Planning Regulations in India", working paper.

[14] Bertaud, A. (2004), "The Spatial Organization of Cities: Deliberate Outcome or Unforeseen Consequence?", working paper.

[15] Bertaud, A. and J. K. Brueckner (2005), "Analyzing Building-Height Restrictions: Predicted Impacts and Welfare Costs", Regional Science and Urban Economics, 35 (2), 109-125.

[16] Bleakley, H. and J. Lin (2012), "Portage and Path Dependence", The Quarterly Journal of Economics, 127 (2), 587-644.

[17] Brueckner, J. and K.S. Sridhar (2012), "Measuring Welfare Gains from Relaxation of Land-use Restrictions: The Case of India's Building-Height Limits", Regional Science and Urban Economics, 42 (6), 1061-67.

[18] Burchfield, M., H. G. Overman, D. Puga and M. A. Turner (2006), "Causes of Sprawl: A Portrait from Space", The Quarterly Journal of Economics, 121 (2), 587-633.

[19] Census of India (2011). Provisional Population Totals, Office of the Registrar General and Census Commissioner, New Delhi.

[20] Cervero, R. (2001), "Efficient Urbanisation: Economic Performance and the Shape of the Metropolis", Urban Studies, 38 (10), 1651-1671.

[21] Chauvin, J. P., E. Glaeser, Y. Ma and K. Tobio (2017), "What is Different About Urbanization in Rich and Poor Countries? Cities in Brazil, China, India and the United States", Journal of Urban Economics, 98, 17-49.

[22] Combes, P-P., Duranton, G. and L. Gobillon (2008), "Spatial Wage Disparities: Sorting Matters!", Journal of Urban Economics, 63 (2), 723-742.

[23] Combes, P-P., G. Duranton, L. Gobillon and S. Roux (2010), "Estimating Agglomeration Economies with History, Geology, and Worker Effects", in E. Glaeser (ed.), *The Economics of Agglomeration*, Cambridge, MA: NBER and University of Chicago Press.

[24] Dev, S. (2006), "Rent Control Laws in India: A Critical Analysis", CCS Working Paper No. 158, Centre for Civil Society, New Delhi.

[25] Donaldson, D. and R. Hornbeck (2016), "Railroads and American Economic Growth: A Market Access Approach", The Quarterly Journal of Economics, 131 (2), 799-858.

[26] Glaeser, E. (2008), *Cities, Agglomeration and Spatial Equilibrium*, Oxford: Oxford University Press.

[27] Glaeser, E. (2011), *Triumph of the City: How Our Greatest Invention Makes Us Richer, Smarter, Greener, Healthier, and Happier.* New York: Penguin Press.

[28] Glaeser, E. and M. Kahn (2004), "Sprawl and Urban Growth" in V. Henderson. and J. Thisse (eds.), *The Handbook of Regional and Urban Economics*, Amsterdam: North Holland Press.

[29] Henderson, V. (1974), "The Sizes and Types of Cities", American Economic Review, 64 (4),640-656.

[30] Henderson, V., A. Storeygard and D. N. Weil (2012), "Measuring Economic Growth from Outer Space", American Economic Review, 102 (2), 994-1028.

[31] Indian Institute for Human Settlements (2013), "Urban India 2011: Evidence", working paper.

[32] Jaitley, A. (2018), "Reconciling Fiscal Federalism and Accountability: Is there a Low Equilibrium Trap?", in *Economic Survey of India 2017-2018*, New Delhi: Government of India, Ministry of Finance.

[33] Kariuki, R. and J. Schwartz (2005), “Small-Scale Private Service Providers of Water Supply and Electricity: A Review of Incidence, Structure, Pricing, and Operating Characteristics”, The World Bank, Policy Research Working Paper Series 3727.

[34] Kreindler, G. (2018), “The Welfare Effect of Road Congestion Pricing: Experimental Evidence and Equilibrium Implications”, working paper.

[35] McKinsey Global Institute (2010), “India’s Urban Awakening: Building Inclusive Cities, Sustaining Economic Growth”.

[36] McMillen, D. P. (2001), “Nonparametric Employment Subcenter Identification”, Journal of Urban Economics, 50 (3), 448-473.

[37] Moretti, E. (2011), “Local Labor Markets”, in *Handbook of Labor Economics*, eds. O. Ashenfelter and D. Card, Amsterdam: North-Holland.

[38] Nunn, N. and D. Puga (2012), “Ruggedness: The Blessing of Bad Geography in Africa”, Review of Economics and Statistics, 94 (1), 20-36.

[39] Nunn, N. and N. Qian (2011), “The Potato’s Contribution to Population and Urbanization: Evidence from an Historical Experiment”, The Quarterly Journal of Economics 126 (2), 593-650.

[40] Roback, J. (1982), “Wages, Rents and the Quality of Life”, Journal of Political Economy, 90 (6), 1257-1278.

[41] Rosen, S. (1979), “Wage-Based Indexes of Urban Quality of Life”, in P. Mieszkowski and M. Strazheim (eds.), *Current Issues in Urban Economics*, Baltimore: Johns Hopkins University Press.

[42] Rosenthal, S. S. and W. C. Strange (2004), “Evidence on the Nature and Sources of Agglomeration Economies”, in V. Henderson. and J. Thisse (eds.), *The Handbook of Regional and Urban Economics*, Amsterdam: North Holland Press.

[43] Saiz, A. (2010), “The Geographic Determinants of Housing Supply”, Quarterly Journal of Economics, 125 (3), 1253-1296.

[44] Seto, K. C., Fragkias, M., Güneralp, B. and M. K. Reilly (2011), “A Meta-Analysis of Global Urban Land Expansion”, PLoS ONE, 8 (6), 1-9.

[45] Small, K. and E. Verhoef (2007), *The Economics of Urban Transportation*, New York, Routledge.

[46] Sridhar, K. S. (2010), “Impact of Land Use Regulations: Evidence From India’s Cities”, Urban Studies, 47 (7), 1541-1569.

[47] Suzuki, H., A. Dastur, S. Moffatt, N. Yabuki, H. Maruyama (2010), “Ecological Cities as Economic Cities”, The World Bank.

[48] Storeygard, A. (2016), “Farther on Down the Road: Transport Costs, Trade and Urban Growth in Sub-Saharan Africa”, Review of Economic Studies, 83 (4), 1263-1295.

[49] United Nations Department of Economic and Social Affairs, Population Division (2014), “World Urbanization Prospects: The 2014 Revision”.

[50] UN-Habitat (2016), “World Cities Report 2016”.

[51] U.S. Army Map Service (ca. 1950), *India and Pakistan Topographic Maps, Series U502, 1:250,000,* U.S. Army Map Service Topographic Map Series.

[52] World Bank (2013), *Urbanization Beyond Municipal Boundaries: Nurturing Metropolitan Economies and Connecting Peri-urban Areas in India*, Washington D.C: The World Bank.

[53] World Bank (2018), World Bank Open Data, *https://data.worldbank.org/*.

Table 1: Descriptive statistics

A. Panel 1950, 1992-2010

	Obs.	*Mean*	*Median*	*St. Dev.*	*Min*	*Max*
Area, km^2	5028	73.15	24.14	191.93	0.43	3,986
Shape, km	5028	3.58	2.60	3.29	0.35	38.21
Potential shape, km	5028	3.16	2.65	1.88	0.42	20.03
City Population	1135	480,626	133,229	1,546,857	5,822	22,085,130

Notes: this table reports descriptive statistics from the 351 cities in the main estimation sample, in all years for which data is available. City area, shape, and potential shape are observable in years 1950 and 1992-2010. City population is available for Census years 1951, 1991, 2001, and 2011.

B. Long difference

	1950	*2010*	*2010-1950*
Area, km^2	3.743	118.0	114.2
	(7.22)	(304.15)	(298.64)
Shape, km	1.012	4.714	3.703
	(0.71)	(4.24)	(3.81)
Potential shape, km	1.376	3.985	2.608
	(0.99)	(2.30)	(1.78)
City Population	106,807	657,420	550614
	(325,337)	(1,968,436)	(1,703,455)

Notes: this table reports variable averages for the 351 cities in the main estimation sample for years 1950, 2010, and for the long difference 2010-1950. For city population, 1950 and 2010 correspond to Census years 1951 and 2011, respectively. Standard deviations in parentheses.

Table 2: First stage

	Long difference, 2010-1950		Panel 1950, 1992-2010	
	(1)	(2)	(3)	(4)
Dependent variable	*Δ Shape, km*	*Δ Log area*	*Shape, km*	*Log area*
Δ Potential shape, km	1.941***	0.232***		
	(0.249)	(0.0488)		
Δ Log projected population	-2.226***	0.0467		
	(0.484)	(0.131)		
Potential shape, km			1.503***	0.185***
			(0.241)	(0.0495)
Log projected population			-1.354***	0.213*
			(0.278)	(0.122)
Observations	351	351	5,028	5,028
AP F stat shape	27.51	27.51	78.36	78.36
AP F stat area	9.14	9.14	13.60	13.60
KP F stat	12.86	12.86	15.97	15.97
City FE			Y	Y
Year FE			Y	Y

Notes: this table reports OLS estimates of the first-stage relationship between city shape and area, and the instruments discussed in Section 5. Each observation is a city in cols. 1 and 2 and a city-year in cols. 3 and 4. The dependent variables are the 2010-1950 long differences in city shape, in km, in col.1, and city area, in squared km, in col. 2; and levels of city shape and city area in cols. 3 and 4. The regressors are the shape of the potential footprint and log projected historic population, in long differences in cols. 1 and 2, and in levels in cols. 3 and 4. Shape is defined as the average distance between two points in the city. Cols. 3 and 4 include city and year fixed effects. Angrist-Pischke and Kleibergen-Paap F statistics are reported. Robust standard errors in parentheses (clustered at the city level in cols. 3 and 4). *** p<0.01,** p<0.05,* p<0.1.

Table 3: Impact of city shape on population

Dependent variable: Δ Log population, 2011-1951		
	(1)	(2)
	IV	OLS
Δ Shape, km	-0.0964**	0.0222***
	(0.0439)	(0.00721)
Δ Log area	0.851***	0.213***
	(0.238)	(0.0338)
Observations	351	351
AP F stat shape	27.51	
AP F stat area	9.14	
KP F stat	12.86	

Notes: this table reports estimates of the impact of city shape on population. Each observation is a city. The dependent variable is the 2011-1951 long difference in log city population. The regressors are the 2010-1950 long differences in city shape, in km, and log city area. Estimation is by IV in col. 1 and OLS in col. 2. Angrist-Pischke and Kleibergen-Paap F statistics are reported in col. 1. Robust standard errors in parentheses.*** p<0.01,** p<0.05,* p<0.1.

Table 4: Impact of city shape on wages

	Dependent variable: Δ Log wage, 2010-1990					
	(1)	(2)	(3)	(4)	(5)	(6)
	IV	OLS	IV	OLS	IV	OLS
Sample	*All*		*Only districts with one city*		*Only top city per district*	
Δ Shape, km	0.0364	0.0336**	0.0728	0.0466***	0.0562*	0.0349**
	(0.0354)	(0.0132)	(0.0470)	(0.0154)	(0.0293)	(0.0150)
Δ Log area	-1.057	-0.0787	0.0542	-0.0371	-0.418	-0.0668
	(0.944)	(0.0833)	(0.368)	(0.130)	(0.435)	(0.101)
Observations	183	183	80	80	145	145
AP F stat shape	13.86		10.21		10.4	
AP F stat area	1.94		3.76		3.10	
KP F stat	1.67		1.76		2.28	
Avg. yearly wage, 1992	72	72	72	72	72	72
Avg. yearly wage, 2010	187	187	193	193	187	187

Notes: this table reports estimates of the impact of city shape on wages. Each observation is a city. The dependent variable is the 2010-1990 long difference in the log urban average yearly wage in the city's district. The regressors are the 2010-1992 long differences in city shape, in km, and log city area. Estimation is by IV in odd columns and OLS in even columns. In cols. 3 and 4 the sample is restricted to districts with only one city. In cols. 5 and 6 the sample is restricted to the top cities in their respective districts. Angrist-Pischke and Kleibergen-Paap F statistics are reported. Average yearly wages are in thousand 2018 Rupees. Robust standard errors in parentheses.*** p<0.01,** p<0.05,* p<0.1.

Table 5: Impact of city shape on rents

	Dependent variable: Δ Log rent 2008-2006					
	(1)	(2)	(3)	(4)	(5)	(6)
	IV	OLS	IV	OLS	IV	OLS
Sample	*All*		*Only districts with one city*		*Only top city per district*	
Δ Shape, km	-0.606	0.000310	-0.486	-0.0172	-0.697	0.0145
	(0.521)	(0.0472)	(0.310)	(0.0675)	(0.648)	(0.0476)
Δ Log area	-2.367	-0.0101	-1.245	-0.101	-1.955	-0.0594
	(2.145)	(0.0902)	(1.044)	(0.108)	(2.094)	(0.0970)
Observations	262	262	134	134	215	215
AP F stat shape	9.60		14.77		5.11	
AP F stat area	3.00		6.12		2.80	
KP F stat	1.67		2.93		1.20	
Avg. yearly rent per m^2, 2006	703	703	707	707	700	700
Implied willingness to pay	-.133		-0.151**		-.167	
$0.16 \cdot \beta_{Rents} - \beta_{Wages}$	(0.082)		(.072)		(0.107)	
	[0.104]		[0.037]		[0.115]	

Notes: this table reports estimates of the impact of city shape on housing rents. Each observation is a city. The dependent variable is the 2008-2006 long difference in the log urban average of housing rent per square meter in the city's district. The regressors are the 2008-2006 long differences in city shape, in km, and log city area. Estimation is by IV in odd columns and OLS in even columns. In cols. 3 and 4 the sample is restricted to districts with only one city. In cols. 5 and 6 the sample is restricted to the top cities in their respective districts. Angrist-Pischke and Kleibergen-Paap F statistics are reported. Average yearly rents are in 2018 Rupees. The implied willingness to pay is calculated as discussed in Section 6 and is based on coefficients from this table and Tables 4. Robust standard errors in parentheses, p-values in brackets.*** p<0.01,** p<0.05,* p<0.1.

Table 6: IV impact of city shape on population, robustness to confounding trends

	Dependent variable: Δ Log population, 2011-1951						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Δ Shape, km	-0.0976**	-0.0964**	-0.103**	-0.0923**	-0.0993**	-0.0872**	-0.0888**
	(0.0453)	(0.0426)	(0.0472)	(0.0428)	(0.0445)	(0.0395)	(0.0391)
Δ Log area	0.857***	0.851***	0.887***	0.839***	0.874***	0.799***	0.798***
	(0.244)	(0.229)	(0.258)	(0.232)	(0.244)	(0.215)	(0.211)
Control	-5.56e-05	-4.03e-09	-0.00285	0.000845	-0.000397	-0.0223**	0.190
	(0.000167)	(0.000115)	(0.00237)	(0.000544)	(0.000319)	(0.0103)	(0.123)
Observations	351	351	351	351	351	351	351
AP F stat shape	27.32	29.64	26.71	28	26.62	30.06	32.25
AP F stat area	8.89	9.37	8.41	9.37	8.69	10.34	9.75
KP F stat	12.44	13.63	11.90	12.93	12.11	14.27	15.68
Characteristic	Elevation	Distance from coast	Distance from river/lake	Distance from mineral deposit	Ruggedness	Bedrock depth	Crop suitability

Notes: this table reports the same IV specification as in Table 3, col. 1, augmented with time-invariant controls, described in Section A.4 in the Appendix. Robust standard errors in parentheses.*** p<0.01,** p<0.05,* p<0.1.

Table 7: Falsification test with lagged outcomes, population

Dependent variable:	*Δ Log population, 2001-1951*	*Δ Log population, 2001-1991*	*Δ Log population, 1991-1951*	*Δ Log population, 1991-1951*
	(1)	(2)	(3)	(4)
2005 Potential shape, km	0.0237	0.076		
	(0.129)	(0.0494)		
2010 Potential shape, km	0.0157	-0.0361		
	(0.0910)	(0.0344)		
1995 Potential shape, km			0.0601	
			(0.196)	
2000 Potential shape, km			-0.00273	
			(0.183)	
Δ Potential shape, km, 2010-2001				-0.0509
				(0.0407)
2005 Projected population	-4.524***	-0.561		
	(0.823)	(0.356)		
2010 Projected population	4.387***	0.515		
	(0.806)	(0.347)		
1995 Projected population			-4.087***	
			(0.848)	
2000 Projected population			3.933***	
			(0.831)	
Δ Projected population, 2010-2001				1.996***
				(0.417)
Observations	238	238	267	267

Notes: this table presents a falsification test to show that the shape instrument is not correlated with past population growth rates. Each observation is a city. The dependent variables are long differences of log population and the regressors are levels and long differences of the instruments. Estimation is by OLS. Robust standard errors in parentheses. *** p<0.01,** p<0.05,* p<0.1.

Table 8: First stage and impact of city shape on population, single instrument

	(1)	(2)	(3)
	First Stage	IV	OLS
Dependent variable:	*Δ Norm. Shape*	*Δ Population Density*	
Δ Potential norm. shape	0.0996***		
	(0.0188)		
Δ Norm. shape		-171.8***	-22.19***
		(37.32)	(7.806)
Observations	351	351	351
AP F stat shape	28.05	28.05	
KP F stat	21.07	21.07	
Mean dep var in levels, 2010	0.969	6.614	
Mean dep var in levels, 1950	1.066	29.872	

Notes: this table reports estimates of the impacts of city shape on population, employing the strategy discussed in Section E in the Appendix. Col. 1 reports the first stage of actual normalized shape on potential normalized shape. Cols. 2 and 3 report the IV and OLS estimates of the impact of normalized shape on population density, measured in thousand inhabitants per square km. All variables are expressed as 2010-1950 long differences (2011-1951 for population). The construction of the instrument is based on a purely mechanical model for city expansion. Normalized shape is the area-invariant version of the disconnection index (see Section A.2 in the Appendix). The mean and standard deviation of normalized shape in the panel are respectively 0.96 and 0.08. Angrist-Pischke and Kleibergen-Paap F statistics are reported. Robust standard errors in parentheses. *** p<0.01,** p<0.05,* p<0.1.

Table 9: Heterogeneous effects of infrastructure and transit

	Dependent variable: Δ Log population, 2011-1951							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Shape, km	-0.247**	-0.410**	-0.173**	-0.152***	-0.163***	-0.316**	-0.290**	-0.227**
	(0.115)	(0.205)	(0.0829)	(0.0507)	(0.0611)	(0.125)	(0.114)	(0.0971)
Δ Shape X Roads	7.48e-06**	0.000167**	0.0109**	0.249***	0.386*	0.000356**	0.000772***	0.197**
	(3.44e-06)	(7.75e-05)	(0.00455)	(0.0763)	(0.211)	(0.000142)	(0.000292)	(0.0946)
Δ Log area	1.192***	1.468**	0.997***	0.864***	0.882***	1.482***	1.421***	1.269***
	(0.462)	(0.665)	(0.347)	(0.216)	(0.267)	(0.532)	(0.492)	(0.410)
Observations	336	336	336	123	123	246	246	246
AP F stat interaction	911.41	35.14	328.24	41.95	36.37	504.84	517.16	305.6
AP F stat shape	10.29	4.60	13.07	7.66	8.34	6.31	7.07	7.54
AP F stat area	4.91	3.16	6.57	10.32	12.02	3.38	3.63	4.37
KP F stat	7.95	5.67	9.96	9.79	8.41	6.06	6.39	7.09
Interaction variable	Roads 2019	Roads 1981	State roads 1981	Proximity	Grid conformity	Cars 2011	Cars 2001	State cars 1984
Mean interaction var	695	150	1.517	.001	.132	19	8	105

Notes: this table reports the same IV specification as in Table 3, col.1, augmented with interactions between city shape and the transit-related variables indicated in each column. Roads 2019 is the total length of roads in a city's 2010 lit-up shape, as reported in 2019 in Openstreetmap. Roads 1981 is the length of city urban roads in 1981 from the Census. State roads in 1981 is to the total length of urban roads in a state in 1981. All road length variables are in km. Proximity is an index of distance accessibility, from Akbar et al (2018). Grid conformity is a measure of the regularity of a city's primary road grid from Akbar et al (2019). Cars 2011 indicates the number of households with a car, from the Census, in thousands. Cars 2001 is analogous for 2001. State cars 1984 is the number of motor vehicles registered in a state. State variables are drawn from the Ministry of Road Transport and Highways and normalized by state urban area (in square km). Further details can be found in Section A.5 in the Appendix. Robust standard errors in parentheses.*** p<0.01,** p<0.05,* p<0.1.

Table 10: Employment centers and work trips

Dependent variable:	*Log nr subcenters, 2005*		*Log avg. distance to work , 2011*			
			Car		*Walk*	
	(1)	(2)	(3)	(4)	(5)	(6)
	IV	OLS	IV	OLS	IV	OLS
Shape, km	-0.0639*	-0.0579***	0.00257	0.0143***	-0.00314	-0.0138***
	(0.0379)	(0.0154)	(0.0181)	(0.00470)	(0.0193)	(0.00453)
Log area	0.611***	0.571***	-0.0171	-0.0649***	0.0781	0.115***
	(0.125)	(0.0568)	(0.0758)	(0.0243)	(0.0824)	(0.0230)
Observations	188	188	238	238	238	238
AP F stat shape	6.74		5.02		5.02	
AP F stat area	18.92		16.74		16.74	
KP F stat	4.76		4.44		4.44	

Notes: this table reports cross-sectional estimates of the impact of city shape on the number of employment subcenters in 2005 (cols. 1 and 2) and average reported distance to work in 2011 (cols. 5 through 8). Each observation is a city in cols. 1 and 2 and a district in cols. 5 through 8. In cols. 1 and 2 the dependent variable is the log number of employment subcenters in a city in 2005, detected using the method described in Section F in the Appendix based on establishment addresses from the Economic Census. In cols. 5 through 8 the dependent variable is the log weighted average distance to work of workers in a district, from the 2011 Census. Cols. 5 and 6 consider workers commuting by car and cols. 7 and 8 consider workers commuting on foot. The regressors are city shape, in km, and log city area, measured in 2005 (cols. 1 and 2) and 2010 (cols. 5 through 6). Estimation is by IV in odd columns, and OLS in even columns. Angrist-Pischke and Kleibergen-Paap F statistics are reported. Robust standard errors in parentheses.*** p<0.01,** p<0.05,* p<0.1.

Table 11: Impact of city shape on public services

	A. Electricity				***B. Tap water***			
Dependent variable:	*Δ Log nr households, 2011-1991*		*Δ Log share households, 2011-1991*		*Δ Log nr households, 2011-1991*		*Δ Log share households, 2011-1991*	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	IV	OLS	IV	OLS	IV	OLS	IV	OLS
Δ Shape, km	0.229	0.0741***	0.0268	-0.00279	0.185	0.0815***	-0.0167	0.00459
	(0.152)	(0.0103)	(0.0249)	(0.00227)	(0.113)	(0.0156)	(0.0380)	(0.0122)
Δ Log area	1.685	-0.0948*	0.340	0.0398**	1.427	-0.176**	0.0823	-0.0409
	(1.789)	(0.0544)	(0.311)	(0.0160)	(1.517)	(0.0716)	(0.467)	(0.0527)
Observations	201	201	201	201	201	201	201	201
AP F stat shape	8.88		8.88		8.88		8.88	
AP F stat area	2.49		2.49		2.49		2.49	
KP F stat	1.77		1.77		1.77		1.77	

Notes: this table reports estimates of the impact of city shape on public services. The specifications reported are similar to those in Table 3. The dependent variables are the 2011-1991 long differences in the log number (cols.1, 2, 5, and 6) and share (cols. 3, 4, 7, and 8) of households with service access. Cols. 1 through 4 report results for electricity and cols. 5 through 8 report results for tap water. The regressors are the 2010-1992 long differences in city shape, in km, and log city area. Estimation is by IV in odd columns and OLS in even columns. The average shares of households with electricity and tap water in 1991 are, respectively, 0.82 and 0.7. Angrist-Pischke and Kleibergen-Paap F statistics are reported. Robust standard errors in parentheses.*** p<0.01,** p<0.05,* p<0.1.

Table 12: Impact of city shape on slum population

	(1)	(2)	(3)	(4)
	IV	OLS	IV	OLS
Dependent variable:	*Δ Log slum population, 2011-1981*		*Δ Log slum population share, 2011-1981*	
Δ Shape, km	-0.154*	-0.0387**	-0.167**	-0.0502***
	(0.0798)	(0.0149)	(0.0823)	(0.0143)
Δ Log area	0.651	0.0378	0.702	-0.0614
	(0.760)	(0.107)	(0.757)	(0.103)
Observations	200	200	200	200
AP F stat shape	15.98		15.98	
AP F stat area	4.08		4.08	
KP F stat	6.17		6.17	

Notes: this table reports estimates of the impact of city shape on slum population. The specifications reported are similar to those in Table 3. The dependent variables are 2011-1981 long differences in the log number (cols.1 and 2) and share (cols. 3 and 4) of slum households in a city. The regressors are 2010-1950 long differences in city shape, in km, and log city area. Estimation is by IV in odd columns and OLS in even columns. The average share of slum households in 1981 is 0.2. Angrist-Pischke and Kleibergen-Paap F statistics are reported. Robust standard errors in parentheses.*** p<0.01,** p<0.05,* p<0.1.

Table 13: Impact of FARs on city shape

Dependent variable	*Shape, km*	*Log area, km²*
	(1)	(2)
Log projected population	2.995	1.981**
	(2.755)	(0.788)
Log projected population x FAR	-1.975*	-0.705**
	(1.023)	(0.319)
Potential shape, km	0.158	-0.184
	(1.182)	(0.232)
Potential shape, km x FAR	0.667	0.137
	(0.487)	(0.107)
Observations	1,182	1,182
City FE	Y	Y
Year FE	Y	Y

Notes: this table reports estimates of the relationship between Floor Area Ratios, city shape, and area. Each observation is a city-year. Cols. 1 and 2 are similar to cols. 3 and 4 in Table 2, augmented with interactions between FARs and each of the instruments. Estimation is by OLS. All specifications include city and year fixed effects. Standard errors clustered at the city level in parentheses. *** p<0.01,** p<0.05,* p<0.1.

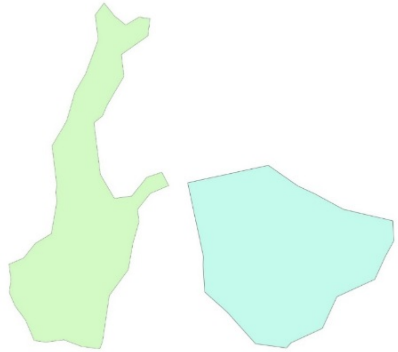

Shape metric	Kolkata	Kolkata Normalized	Bangalore	Bangalore Normalized
Disconnection, km	20.4	1.20	16.0	0.94
Remoteness, km	14.8	0.87	11.8	0.69
Spin, km^2	287.1	0.99	9.4	0.55
Range, km	65.2	3.83	44.5	2.62

Figure 1: Shape metrics: an example

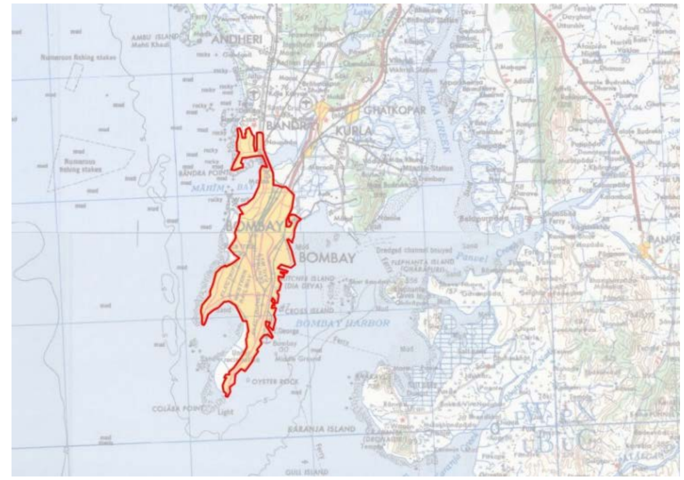

Figure 2: U.S. Army India Topographic Maps

Figure 3: DMS/OLS night-time lights, year 1992

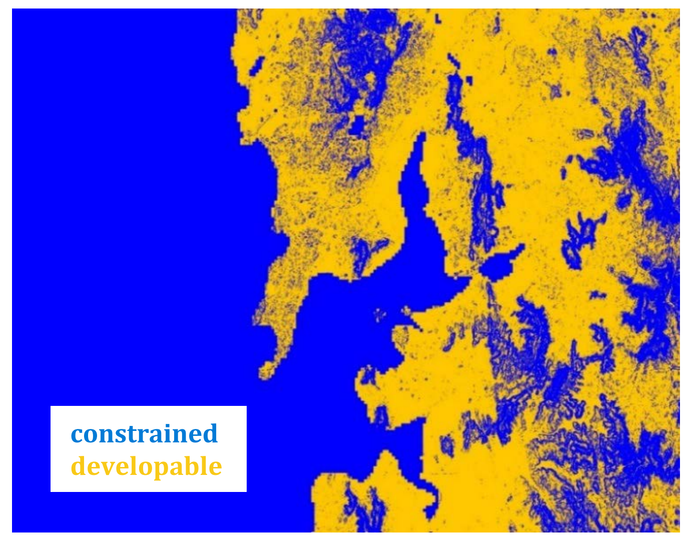

Figure 4: Developable vs. constrained land

5a

5b

5c

5d

Figure 5
Instrument construction

Appendix for Online Publication

This Appendix is organized as follows. Section A discusses in detail the data employed in the paper. Section B presents a formal derivation of the spatial equilibrium model outlined in Section 4 in the paper. Section C outlines an alternative model of within-city spatial equilibrium to shed light on the rent gradient in a city with irregular shape. Section D provides an analytical derivation of the bias in the estimation of the OLS impacts of city shape. Section E discusses an alternative version of the shape instrument, used in one of the robustness checks, that does not rely on projecting historical population. Section F describes the procedure used to detect employment sub-centers to obtain the estimates in Table 10.

A. Data

Below I provide details on the sources and methods employed to assemble my dataset. I start by describing how I retrieve urban footprints (Section A1) and measure their geometric properties (Section A2). Next, I discuss population, wages, and rents (Section A3). In Section A4 I discuss the control variables employed in Table A2 and in the robustness checks in Tables 6, A10, A11, and A12. In Section A5 I discuss variables related to infrastructure and transit. In Section A6 I discuss the remaining data sources - geographic constraints used for the construction of the instrument and variables employed for heterogeneous effects analyses, including firms' addresses, slum population, and floor area ratios.

A1. Urban footprints

The first step in constructing the dataset is to trace the footprints of Indian cities at different points in time and measure their geometric properties. The boundaries of urban footprints are retrieved from two sources. The first is the U.S. Army India and Pakistan Topographic Maps (U.S. Army Map Service, ca. 1950), a series of detailed maps covering the entire Indian subcontinent at a 1:250,000 scale. These maps consist of individual topographic sheets. I geo-referenced each of these sheets and manually traced the reported perimeter of urban areas, which are clearly demarcated. These maps are from the mid-50s, but no specific year of publication is provided. For the purposes of constructing the city-year panel, I label these data as 1950 and match them with Census data from 1951.

The second source is the DMSP/OLS Night-time Lights dataset. This consists of night-time imagery recorded by satellites from the U.S. Air Force Defense Meteorological Satellite Program (DMSP) and reports the recorded intensity of Earth-based lights, measured by a six-bit number (ranging from 0 to 63). This data is reported for every year between 1992 and 2010, with a resolution of 30 arc-seconds (approximately 1 square kilometer). The use of the DMSP-OLS dataset for delineating urban areas is quite common in urban remote sensing (Henderson et al., 2003; Small et al., 2005; Small et al., 2013). The methodology is the following: first, I overlap the night-time lights imagery with a point shapefile with the coordinates of Indian settlement points, taken from the Global Rural-Urban Mapping Project (GRUMP) Settlement Points dataset (Balk et al., 2006; CIESIN et al., 2011). I then set a luminosity threshold (35 in my baseline approach, as explained below) and consider spatially contiguous lighted areas surrounding the city coordinates with luminosity above that threshold. This approach can be replicated for every year covered by the DMSP/OLS dataset.

The choice of luminosity threshold results in a more or less restrictive definition of urban areas, which will appear larger for lower thresholds.[1] To choose luminosity thresholds appropriate for India, I overlap the 2010 night-time lights imagery with available Google Earth imagery. I find that a luminosity threshold of 35 generates the most plausible mapping for those cities covered by both sources.[2] In my main estimation sample (that includes cities covered in 1950 and 2010), the average city footprint occupies an area of approximately 73 square kilometers. In Table A5, I show that the first-stage results and main population IV results are robust to using alternative luminosity thresholds 30 and 40.

Using night-time lights as opposed to alternative satellite-based products, in particular day-time imagery, is motivated by a number of advantages. Unlike products such as aerial photographs or high-resolution imagery, night-time lights cover systematically the entire Indian subcontinent, and not only a selected number of cities. Moreover, they are one of the few sources allowing to detect changes in urban areas over time, due to their yearly temporal frequency. Finally, unlike multi-spectral satellite imagery, night-time lights do not require any sophisticated manual pre-processing and cross-validation using alternative sources.[3]

It is well known that urban maps based on night-time lights tend to inflate urban boundaries, due to "blooming" effects (Small et al., 2005).[4] This can only partially be limited by setting high luminosity thresholds. In my panel, urban footprints as reported for years 1992-2010 thus reflect a broad definition of urban agglomerations, which typically goes beyond the current administrative boundaries. This contrasts with urban boundaries reported in the US Army maps, which seem to reflect a more restrictive definition of urban areas (although no specific documentation is available). Throughout my analysis, I focus on long differences or include year fixed effects, which amongst other things control for these differences in data sources, as well as for different calibrations of the night-time lights satellites.

The resulting panel dataset of urban footprints is unbalanced for two reasons: first, some settlements become large enough to be detectable only later in the panel; second, some settlements appear as individual cities for some years in the panel, and then become part of larger urban agglomerations in

[1] Determining where to place the boundary between urban and rural areas always entails some degree of arbitrariness, and in the urban remote sensing literature there is no clear consensus on how to set such threshold. It is nevertheless recommended to validate the chosen threshold by comparing the DMSP/OLS-based urban mapping with alternative sources, such as high-resolution day-time imagery, which in the case of India is available only for a small subset of city-years.

[2] For years covered by both sources (1990, 1995, 2000), my maps also appear consistent with those from the GRUMP - Urban Extents Grid dataset, which combines night-time lights with administrative and Census data to produce global urban maps (CIESIN et al., 2011; Balk et al., 2006).

[3] An extensive portion of the urban remote sensing literature compares the accuracy of this approach in mapping urban areas with that attainable with alternative satellite-based products, in particular day-time imagery (e.g. Henderson et al., 2003; Small et al., 2005). This cross-validation exercise has been carried out also specifically in the context of India by Joshi et al. (2011) and Roychowdhury et al. (2009). The conclusion of these studies is that none of these sources is error-free, and that there is no strong case for preferring day-time over night-time satellite imagery if aerial photographs are not systematically available for the area to be mapped.

[4] DMSP-OLS night-time imagery overestimates the actual extent of lit area on the ground, due to a combination of coarse spatial resolution, overlap between pixels, and minor geolocation errors (Small et al., 2005).

later years. The number of cities in the panel ranges from 351 to 457, depending on the year considered. As a result, the "long-difference" sample used in the baseline specifications includes 351 observations. In Appendix Table A4, I show that the results continue to hold in the full sample.

The criterion for being included in the analysis is to appear as a contiguous lighted shape in the night-time lights dataset. This appears to leave out only very small settlements.

A2. Shape metrics

The notion of "compactness" of an urban footprint is borrowed from the urban planning and landscape ecology literature (Angel et al., 2009 and Parent et al., 2009). Intuitively, the geometric concept of compactness rests on the idea that the circle is the most compact shape. The extent to which a polygon's shape departs from that of a circle can be measured through many distinct indexes, all based on the distribution of points within a polygon.

My benchmark indicator throughout the paper is the disconnection index (corresponding to the "cohesion" index in Angel et al., 2009). It is defined as the average Euclidean distance between all pairs of interior points within a polygon and can be viewed as a proxy for the length of all potential trips within the city, without restricting one's attention to those to or from the center. Higher values of the index denote longer distances within the city and less compact shape. Specifically, consider n random interior points sampled within a polygon and denote the distance between points j and i as d_{ij}. The index is defined as follows:

$$S = \frac{\sum_{i=1}^{n} \sum_{j}^{n} d_{ij}}{n(n-1)}.$$

This is illustrated in Figure A.1 for four hypothetical sample points.

I compute the index using the Parent et al. (2009) ArcGIS routines.[5] The Shape Metrics tool considers 20,000 interior points, uniformly distributed throughout the polygon in a grid pattern. Then, for computational ease, the index is computed for 30 samples of 1000 randomly selected points within this set and averaged. I compute the index in kilometers.

The disconnection index is mechanically correlated with polygon area. In order to disentangle the effect of geometry *per se* from that of city size, two approaches are possible. One is to explicitly control for the area of the footprint, which I do in the baseline specification throughout the paper. Alternatively, the index can be normalized, computing a version that is invariant to the area of the polygon. I do so by computing first the radius of the "Equivalent Area Circle" (EAC), namely a circle with an area equal to that of the polygon. I then normalize the index of interest by dividing it by the EAC radius.[6]

Figure A2 reports the disconnection index computed for selected shapes, where S represents the non-normalized index and nS the normalized version. These examples illustrate departures from the circular shape that are associated with higher values of the disconnection index. Elongated shapes (such as v in the figure) and polygons with recesses and gaps (such as iii and iv, similar to urban areas growing around topographic obstacles) are all associated with greater disconnection.

---

[5] I am thankful to Vit Paszto for help with the ArcGIS shape metrics tool.

[6] My normalization is slightly different from that proposed by Parent et al. (2009).

Figure A.1: Calculation of the disconnection index, example from Parent et al. (2009).

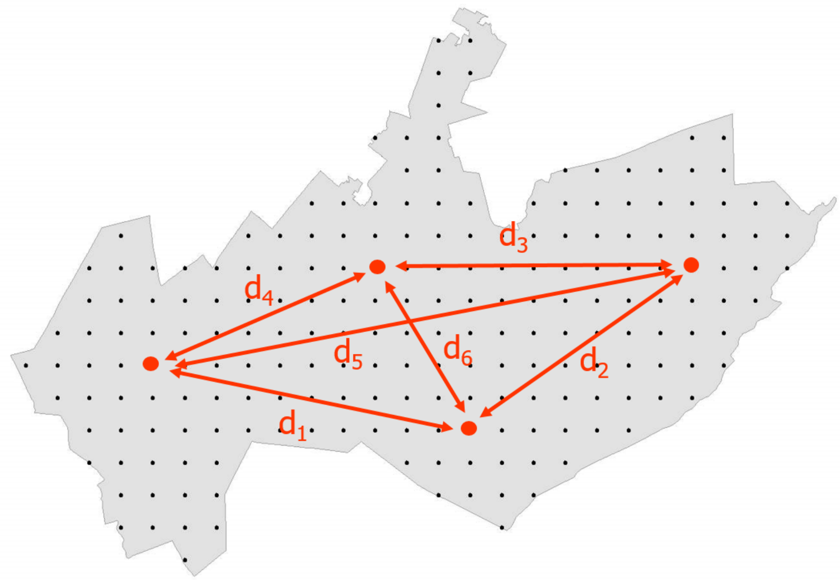

$$S = \frac{d_1 + d_2 + \cdots + d_6}{6}$$

The shortest connecting paths used in the computation of average distances do not need to lie within the polygon. In this regard, the index may underestimate distances that account for the placement of roads. Furthermore, the index is defined for contiguous polygons only: as a result, a built-up area disconnected from the main urban footprint would not contribute to the index. In this respect, the index may underestimate the disconnectedness of non-contiguous development.

For robustness, I also compute three additional shape metrics:

(i) The *remoteness* index ("proximity" in Angel et al., 2009) is the average distance between all interior points and the centroid.[7] It can be viewed as a proxy for the average length of all potential trips to the center.

(ii) The *spin* index is computed as the average of the squared distances between interior points and the centroid. This is similar to the remoteness index, but gives more weight to the polygon's extremities, corresponding to the periphery of the footprint. This index has particularly high values for footprints that have tendril-like projections.

(iii) The *range* index captures the maximum distance between two points on the shape perimeter, representing the longest possible distance between two points within the city.

Similarly to the benchmark indicator, all the indexes above are measured in kilometers, higher values denoting a greater departure from circularity and longer within-city distances, and can be normalized by the EAC radius.

[7] The centroid of a polygon, or center of gravity, is the point that minimizes the sum of squared Euclidean distances between itself and each vertex.

Figure A.2: Disconnection index for a sample of polygons, adapted from Parent et al. (2009).

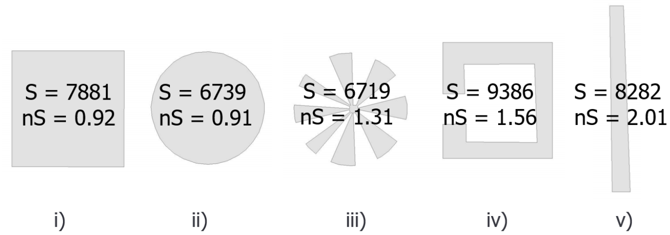

Even though these indexes represent independent properties, in practice they tend to be highly correlated and should be viewed as different proxies for the same broad notion of "compactness". In my dataset, the correlation between any two indexes is between 0.82 and 0.97 (except for the correlation between spin and range that is 0.7). Table A6 and A7 in the Appendix shows that the first-stage and IV results are robust to employing different indexes.

A3. Outcome variables

Population City-level data for India is difficult to obtain. The only systematic source that collects data explicitly at the city level is the Census of India, conducted every 10 years. I employ population data from Census years 1871-2011.[8] As explained in Section 5, historical population (1871-1941) is used to construct the instrument, whereas population drawn from more recent waves (1951, 1991, 2001, and 2011) is used as an outcome variable.

It is worth pointing out that "footprints", as retrieved from the night-time lights dataset, do not always have an immediate Census counterpart in terms of town or urban agglomeration, as they sometimes stretch to include suburbs and towns treated as separate units by the Census. A paradigmatic example is the Delhi conurbation, which as seen from the satellite expands well beyond the administrative boundaries of the New Delhi National Capital Region. When assigning population totals to an urban footprint, I sum the population of all Census settlements that are located within the footprint, thus computing a "footprint" population total. Moreover, in order to assemble a consistent panel of city population totals over the years one also has to account for changes in the definitions of "cities", "urban agglomerations", and "outgrowths" across Census waves. Mitra (1980) provides harmonized figures for all Census waves up to 1971 and I harmonized the rest of the waves.

Wages and rents For wages and rents, I rely on the National Sample Survey and the Annual Survey of Industries, which provide, at most, district identifiers. I thus follow the approach of Greenstone and Hanna (2014) and Chauvin et al. (2017): I match cities to districts and use district urban averages as proxies for city-level averages. It should be noted that the matching is not always perfect, for a number of reasons. First, it is not always possible to match districts as reported in these sources to Census districts, and through these to cities, due to redistricting and inconsistent numbering throughout this

[8]Historical population totals are drawn from Mitra (1980). Census data for years 1991 to 2001 were taken from the Census of India electronic format releases. 2011 Census data were retrieved from the Census official website.

period. Second, there are a few cases of large cities that cut across districts (e.g., Hyderabad). Finally, there are a number of districts which contain more than one city from my sample. I provide results for three samples: one including any city that can be matched; one that only includes cities for which there is a one to one mapping with a district; and finally a sample where I only include the top city in each district. The matching process introduces considerable noise and leads to results that are relatively less precise and less robust than those I obtain with city-level outcomes.

Data on wages are taken from the Annual Survey of Industries (ASI). These are repeated cross-sections of plant-level data collected by the Ministry of Programme Planning and Implementation of the Government of India. The ASI covers all registered manufacturing plants in India with more than fifty workers (one hundred if without power) and a random one-third sample of registered plants with more than ten workers (twenty if without power) but less than fifty (or one hundred) workers. As mentioned by Fernandes and Sharma (2012) amongst others, the ASI data are extremely noisy in some years, which introduces a further source of measurement error. I employ a long difference specification using waves 1990 and 2010. However the results using intermediate waves from 1994, 1995, 1997, 1998, 2009 in the panel are similar.

A drawback of the ASI data is that it covers the formal manufacturing sector only.[9] This may affect the interpretation of my results, to the extent that this sector is systematically over- or underrepresented in cities with worse shapes. I examined the relationship between shape and the industry mix of cities, employing Economic Census data, and found no obvious patterns. The share of manufacturing appears to be slightly lower in non-compact cities, but this figure is not significantly different from zero, which somewhat alleviates the selection concern discussed above.

Unfortunately, there is no systematic source of data for property prices in India across a sufficient number of cities. I construct a proxy for the rental price of housing drawing on the National Sample Survey (Household Consumer Expenditure schedule), which asks households about the amount spent on rent. In the case of owned houses, an imputed figure is provided. Rounds 62 (2005-2006), 63 (2006-2007), and 64 (2007-2008) are the only ones for which the urban data is representative at the district level and which report total dwelling floor area as well. In my baseline specification I focus on rounds 62 and 64 and take a long difference, but results are similar using all three waves in the panel version.

I construct a measure of rent per square meter based on total rent amount and floor area. These figures are likely to be underestimating the market rental rate, due to the presence of rent control provisions in most major cities of India (Dev, 2006). While I cannot observe which figures refer to rent-controlled housing, as an attempt to cope with this problem, I also construct an alternative proxy for housing rents which focuses on the upper half of the distribution of rents per meter. This is *a priori* less likely to include observations from rent-controlled housing.

For robustness, I also consider an alternative source, used by Chauvin et al. (2017): the India

[9] An alternative source of wages data is the National Sample Survey, Employment and Unemployment schedule. This provides individual-level data that cover both formal and informal sector. However, it is problematic to match these data to cities. For most waves, the data are representative at the NSS region level, which typically encompasses multiple districts.

Human Development Survey (IHDS), comprising two rounds (2005 and 2012). The IHDS is a nationally representative household survey including 971 urban neighborhoods across India. It reports monthly expenditures on housing rents as well as housing characteristics (other than dwelling size), including: number of rooms, house type (house with no shared walls, house with shared walls, flat, chawl, slum housing, or other), housing surrounded by sewage, predominant wall type (grass/thatch, mud/unburnt bricks, plastic, wood, burned bricks, GI sheets or other metal, stone, or cement/concrete), predominant roof type(grass/thatch/mud/wood, tile, slate, plastic, GI metal/asbestos, cement, brick, or stone concrete), and predominant floor type (mud, wood/bamboo, brick, stone, cement, tiles/mosaic, or others). This is an advantage over the NSS data as I can run a hedonic regression of rents on the above characteristics and consider the residuals. At most, district identifiers are provided, and not all districts can be matched to Census ones. As a result, only about 260 cities can be matched to an IHDS district.

A4. Correlates of city shape and controls

Below I discuss the results and the sources of the city-level variables employed in Table A2 in this Appendix. Many of these variables are also employed throughout the paper in a number of robustness checks. Summary statistics are provided in Table A3. Unless otherwise specified, all the distance variables are in kilometers and are calculated from the Central Business District (CBD) of each city (based on the centroid of the 1950 footprint).

Table A2 shows the correlation between city shape, in levels and changes, and a number of city-level attributes. Each row presents the OLS coefficients of two distinct regressions: in column 1 I regress city shape in 2010 on the relevant city attribute, controlling for 1950 city area; column 2 is similar, but the dependent variable is the 1950-2010 long difference in city shape and I control for city area in 1950.

In Panel A I rank cities by their 1951 population and assign quartile dummies to each. The bottom three quartiles are all associated with more compact shapes whereas the top quartile is associated with less compact shape. This pattern holds both in levels (column 1) and in changes (column 2). In Section 5 I argue that the correlation between bad shape and city size is spurious and driven by the tendency of cities to deteriorate in shape as they expand.

In Panel B, I consider the channels highlighted by the urban planning literature: access to public services and urban transit. Specifically, I consider the share of households with connections to electricity, with tap water on premises, and with cars, from the 2011 Census (town-level tables).[10] Less compact cities are associated with a larger share of households connected to public services. This runs counter the predictions of the urban planning literature, arguing that compact cities provide better service access. However, any potential causal effects are likely to be confounded by the fact that less compact cities tend to be the largest and highest-income cities in the sample. This can also explain the correlation between non-compact shape and the share of households with cars.

[10] A number of variables employed in the paper are drawn from the Census town-level tables. It should be noted that these tables cover Census towns only, excluding small settlements that may fall within a city's lit-up footprint.

Next, I consider the length of the urban road network. I overlap 2019 road maps from Openstreetmap on 2010 city outlines as defined by the nightlights, following the approach of Akbar et al. (2019). Openstreetmap is a collaborative worldwide mapping project. A caveat in employing these data is that the degree of accuracy and comprehensiveness may vary across cities, raising concerns of measurement error. Despite disconnected cities being larger and more developed, they do not appear to have a denser road network. In levels, higher values of the disconnection index are associated with a shorter urban road network. In changes, the sign becomes negative, but the effect size is small: holding city area constant, as the average within-city trip increases by one kilometer, the road network expands by one meter. This suggests that the provision of infrastructure in non-compact cities may indeed be more difficult, as urban planners suggest.

Interestingly, non-compact shape is positively correlated, both in level and in changes, with the average distance to workplace in 2011. This variable is calculated based on district-level Census data. The 2011 Census reports the number of urban workers in each district residing at different reported distances from their workplaces, by coarse bins (0-1, 2-5, 6-10, 11-20, 21-30, 31-50, or above 50 kilometers). I calculate a district-level average distance to work by averaging the mean distance within each bin,[11] weighted by the share of workers. I then match each city to a district.[12] This should be viewed as a noisy proxy for commuting distance within the city, as the distance bins are coarse and include large distances relative to the average city size. The correlation with city shape (in levels) is positive, suggesting that less compact cities may be associated with longer commutes. Furthermore, the positive correlation with changes in shape indicates that cities with long commutes are cities that became less compact than they were. This is plausible for a city that starts out as monocentric and compact, and grows into a less compact shape over time, with commutes to the center becoming longer over time. Conversely, commutes may not be as long in a city that has always been non-compact and perhaps more polycentric. These patterns are explored further in Section 9 in the paper.

In Panel C, I consider pre-determined city characteristics related to geography and geology:

- Elevation is measured in the CBD, based on data from the Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) Global Digital Elevation Model (NASA and METI, 2011). Mountainous cities are defined as having elevation above 700 meters.

- Distance from the nearest river or lake is measured combining large rivers from the Natural Earth 2.0.0 dataset and lakes from the WWF Global Lakes and Wetlands Database, Level 2. The river/lake dummy is equal to 1 for cities whose CBD lies within 5 kilometers of a river or lake.

- Distance from nearest mineral deposit is calculated based on the location of mineral deposits recorded in the Mineral Resources Data System dataset, assembled by the U.S. Geological Survey. A city is considered with a mineral deposit if there is a deposit within 50 kilometers of the CBD.

[11] I consider 60 kilometers for the "above 50 kilometers" bin.

[12] I show results for all cities, but results are similar when excluding districts with multiple cities or when considering only the main city in each district.

- Ruggedness (in meters) is drawn from the G-Econ gridded dataset (Nordhaus et al., 2006) and is measured at the 1 degree (approximately 100 kilometers) level. It is calculated based on the average absolute change in elevation between adjacent 10 Arc-minutes cells included in each 1 degree cell. Higher values imply more variation in elevation and greater terrain ruggedness. I match each city CBD to the corresponding 100-km grid cell and assign the corresponding value.

- Bedrock depth (in meters) is drawn from the SoilGrids dataset (Hengl et al., 2017), a global gridded dataset at a 250m resolution. I take the average bedrock depth within 100 kilometers of a city's CBD.

- Crop suitability (in tonnes per hectare per year) is calculated based on the potential yields (for low-input, rainfed production) of the top 5 most suitable crops in India (dryland rice, wetland rice, maize, millet, sorghum), drawn from the FAO's Global Agro-Ecological Zones (GAEZ) dataset. The raw data is available at a resolution of 5 arc minutes (approximately 10 kilometers). Yields are averaged within 100 kilometers of the CBD.

Elevation, distance from the coast, distance from water bodies, and terrain ruggedness capture the presence of geographic constraints to city expansion. However, the raw correlation with city shape is insignificant, except for ruggedness which is associated with less compact cities. This suggests that what affects city shape is not the generic presence of particular geographic features and one may have to account for the exact position of geographic constraints - which motivates the way in which I construct my instrument. Similarly, I find no meaningful correlation with bedrock depth, that has been associated in the literature with higher construction costs for high-rises (Barr et al., 2011). Crop suitability and the distance from mineral deposits, which may affect the city's productivity, are also not significant.

Finally, in Panel D I consider other, non-predetermined city characteristics. The British direct rule dummy identifies cities in districts that were formerly part of British India, based on Iyer (2010). Distance from state headquarters, from district headquarters, and from the nearest city with more than 100,000 inhabitants are drawn from the 2011 Census (town-level tables). As expected, shape is persistent in time, as highlighted by the positive coefficient for initial shape. More remote cities, further away from state or district headquarters, tend to be more compact, but the correlation is weak. Cities that were under direct British rule are on average less compact, consistent with the findings of Baruah et al. (2017) on British colonial cities being more sprawled, but this correlation is only borderline significant. Conversely, there is a strong tendency of state capitals to deteriorate in shape, probably because they are also the largest cities.

A5. Infrastructure and transit

Below I discuss the data related to transit and infrastructure employed in Tables 9, A15, and A16.[13]

[13] The data on work commutes employed in Table 10 is discussed in Section A4 above.

Current road network I measure the current length of the road network by overlapping digital roadmaps from 2019 Openstreetmap with 2010 city outlines as defined from the night-time lights. In Table A16 I also consider motorways, defined as the subset of road segments labeled by Openstreetmap as "motorways" or "trunks", corresponding to dual-carriage roads similar to freeways in the US (Akbar et al., 2019).

Indices by Akbar et al. (2018, 2019) In Tables 9, A15, and A16 I consider indices developed for Indian cities by Akbar et al. (2018, 2019).[14] The authors provide estimates of the unit cost of commuting in Indian cities using transit times predicted by Google Maps in 2016, and aggregate these estimates into city-level indices of vehicular mobility. Their methodology is primarily based on feeding into Google Maps origin-destination pairs and collecting information on the duration and length of these artificial trips. In addition, they also create indices related to the spatial properties of the road network in a city, defined overlapping 2016 Openstreetmap with city outlines defined using a combination of night-time lights and other satellite-based products.

The proximity index (Akbar et al., 2018) is a measure of distance-based accessibility. It is based on the road distance between random points in the city and a number of amenities (shopping centers, train stations etc.), selected by Google Maps within a pre-specified radius. Higher values indicate greater accessibility and shorter trips.

The grid conformity index (Akbar et al., 2019) measures the extent to which the city's 2016 road network is laid out as a regular grid. It measures the proportion of edges in a city's road network that conform to the dominant grid orientation, by being perpendicular or parallel to the modal edge bearing. Higher values indicate more regular grids and correlate with better vehicular mobility.

The mobility index (Akbar et al., 2019) is their benchmark index of vehicular mobility and is based on the speed of simulated trips. This index abstracts from city shape as the length of the simulated trips is pre-specified by the authors. Factors affecting this index include road density, road quality, and traffic congestion.

Note that, *a priori*, there is no clear mapping between city shape and the Akbar et al. (2018, 2019) indices, as the latter depend on the internal functioning of the road network and traffic patterns, and not on the city's layout. In Table A16, panel B, I provide IV and OLS estimates for the relationship between city shape and the three indices discussed above, subject to the caveat of very weak instruments. Poor city shape is associated with lower mobility, both in the OLS and the IV. This may stem from disconnected cities having a less functional road network, as highlighted in Table A16, panel A. However, the magnitudes are small: according to the most conservative point estimate, for a one standard deviation deterioration in city shape, mobility declines by 1.2%. Results are more mixed when considering proximity. In the IV, bad shape is associated with lower proximity, but the coefficient is small and insignificant at conventional levels. The corresponding OLS is positive and significant, perhaps reflecting the fact that in the OLS bad shape correlates with city size, and large cities tend to have more amenities to begin with.

[14] I am thankful to the authors for granting me access to their data.

Interaction variables in Tables 9 and A15 Below I discuss the variables employed in the interaction specifications of Table 9 and A15. These tables present the same IV specifications, but Table A15 additionally controls for a proxy for city income (the number of banks in 1981).

In columns 1 through 3 I interact city shape with urban road length. In column 1 I consider the total length of urban roads in a city, obtained from Openstreetmap as discussed above. In column 2 I consider the total length of city urban roads, from the 1981 Census (town directory). In column 3 I consider the total length of urban roads in a state as of 1981, from the Ministry of Road Transport and Highways, accessed through the Centre for Monitoring India Economy. This figure is normalized by the total urban land area in a state, in square kilometers, as provided by the Centre for Industrial and Economic Research's Industrial Databooks (CIER, 1990).

In columns 4 and 5 I consider the proximity and grid conformity index from Akbar et al. (2018, 2019), discussed above.

In columns 6 through 8 I interact city shape with proxies for the availability of motor vehicles. In column 6 (7) I consider the number of city households with access to cars, reported in the 2011 (2001) Census (town-level tables). In column 8 I consider the total number of registered cars in a state in 1984, from the Ministry of Road Transport and Highways, accessed through the Centre for Monitoring India Economy. Year 1984 is the earliest year for which this figure is available for most states. This figure is normalized by the total urban land area in a state in 1981, in square kilometers, as provided by the Centre for Industrial and Economic Research's Industrial Databooks (CIER, 1990).

A6. Other variables

Geographic constraints For the purposes of constructing the city shape instrument, I code geographic constraints to urban expansion as follows. Following Saiz (2010), I consider land pixels as "undevelopable" when they are either occupied by a water body, or characterized by a slope above 15%. I draw upon two high-resolution sources: the Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) Global Digital Elevation Model (NASA and METI, 2011), with a resolution of 30 meters, and the Global MODIS Raster Water Mask (Carroll et al., 2009), with a resolution of 250 meters. I combine these two raster datasets to classify pixels as "developable" or "undevelopable". Figure 4 in the paper illustrates this classification for the Mumbai area.

Firm location and employment subcenters Data on the spatial distribution of employment in year 2005 is derived from the urban Directories of Establishments, pertaining to the 5th Economic Census. For this round, establishments with more than 10 employees were required to provide an additional "address slip", containing a complete address of the establishment, year of initial operation, and employment class. I geo-referenced all the addresses corresponding to cities in my sample through Google Maps API, retrieving consistent coordinates for approximately 240 thousand establishments in about 190 footprints.[15]

I utilize these data to compute the number of employment subcenters in each city, following the

[15] Results are similar to excluding firms whose address can only be approximately located by Google Maps.

two-stage, non-parametric approach described in McMillen (2001). Of the various methodologies proposed in the literature, this is particularly suitable for my context as it can be fully automated and replicated for a large number of cities. This procedure identifies employment subcenters as locations that have significantly larger employment density than nearby ones, and that have a significant impact on the overall employment density function in a city. This procedure is outlined in Section F of this Appendix. The number of employment subcenters calculated for year 2005 ranges from 1, for purely monocentric cities, to 9, for large cities such as Delhi and Mumbai. Consistent with results obtained in the US context by McMillen and Smith (2003), larger cities tend to have more employment subcenters.

Slum population Data on slums is drawn from the 1981 and 2011 Census, which provide slum population counts for selected cities. The Census defines slums as follows: all areas notified as "slum" by state or local Government; and any compact area with population above 300 characterized by "poorly built congested tenements, in unhygienic environment, usually with inadequate infrastructure and lacking in proper sanitary and drinking water facilities". Such areas are identified by Census Operations staff (Director of Census Operations, 2011).

Floor Area Ratios Data on the maximum permitted Floor Area Ratios for a small cross-section of Indian cities (55 cities in my sample)[16] is taken from Sridhar (2010), who collected them from individual urban local bodies as of the mid-2000s. FARs are expressed as ratios of the total floor area of a building over the area of the plot on which it sits. I consider the average of residential and non-residential FARs (but results are similar focusing on residential FARs only). For a detailed discussion of FARs in India, see Sridhar (2010) and Bertaud and Brueckner (2005).

While in this paper I take FARs as given, the question might arise on their determinants. Regressing FAR values on urban shape and area, I found weak evidence of FARs being more restrictive in larger cities, consistent with one of the stated objectives of regulators - curbing densities in growing cities.

B. A Simple Model of Spatial Equilibrium across Cities

I motivate the empirical analysis of the impacts of city shape drawing on a model of spatial equilibrium across cities (Rosen, 1979, Roback, 1982). I embed city shape in this framework by hypothesizing that households and firms may value city compactness when evaluating the trade-offs associated with different cities. In order to deliver the intuition and provide estimable equations, I focus on a simple version of the model, with Cobb-Douglas functional forms, following the exposition in Glaeser (2008). I then discuss caveats and extensions to be addressed in future research.

Model setup

The model features homogeneous households, firms, and developers.

Households have Cobb-Douglas utility $U(C,H,\theta) = \theta C^{1-\alpha} H^{\alpha}$ over a numéraire good C, housing H, and a city-specific "quality of life" parameter θ. The latter captures any utility cost or benefit

[16] Sridhar (2010) collects data for about 100 cities, but many of those cities are part of larger urban agglomerations, and do not appear as individual footprints in my panel, or are too small to be detected by night-time lights.

associated with living in a particular city that requires compensation through factor prices. It is useful to conceptualize θ as consisting of three components: "public services" θ_P, "transit accessibility" θ_T, and "consumption amenities" θ_A. All else being equal, better public services (such as electricity or water), greater accessibility, and better amenities (such as good climate) improve household utility. Denoting city shape with S, I assume that S can affect θ_P and θ_T, in line with the conjectures of urban planners. Households supply labor inelastically, receiving a city-specific wage W. Solving their utility-maximization problem, for a given city, yields the following indirect utility:

$$\log(W) - \alpha \log(p_h) + \log(\theta) = \log(\overline{\upsilon}). \tag{B.1}$$

where p_h is the rental price of housing.

Spatial equilibrium requires that indirect utility $\overline{\upsilon}$ be equalized across cities, otherwise workers would move.[17] This condition delivers the key intuition that households, in equilibrium, implicitly pay for better quality of life, as captured by θ, through lower wages (W) or through higher housing prices (p_h).

In the production sector, firms competitively produce a traded good Y, using labor N, traded capital K (which trades at price 1), a fixed supply of non-traded capital $\overline{Z}$[18], and a bundle of city-specific production amenities A. Their production function is $Y(N,K,\overline{Z},A) = AN^{\beta}K^{\gamma}\overline{Z}^{1-\beta-\gamma}$. Similar to households, firms may benefit from compact city shape through better access to services or because of greater accessibility, which I capture by allowing S to affect A via two components, A_P and A_T. Normalizing the price of traded capital to 1, the zero-profit condition for firms delivers the following labor demand curve:

$$(1-\gamma)\log(W) = (1-\beta-\gamma)(\log(\overline{Z}) - \log(N)) + \log(A) + \kappa_1. \tag{B.2}$$

Finally, developers competitively produce housing H, using land l and "building height" h. In each city there is a fixed supply of land $\overline{L}$, determined by planners.[19] Denoting the price of land with p_l, the developers' cost function reads $C(H) = c_0 h^{\delta} l - p_l l$, with $\delta > 1$.

By combining housing supply, obtained from the developers' maximization problem, with housing demand, resulting from the households' problem, one obtains the following housing market equilibrium condition:

$$(\delta-1)\log(H) = \log(p_h) - \log(c_0\delta) - (\delta-1)\log(N) + (\delta-1)\log(\overline{L}). \tag{B.3}$$

[17] The notion of spatial equilibrium across cities presumes that households are choosing across various locations. While mobility in India is lower than in other developing countries (Munshi and Rosenzweig, 2016), the observed pattern of migration to urban areas is compatible with this element of choice: as per the 2001 Census, about 38% of rural to urban internal migrants move to a location outside their district of origin, presumably choosing a city rather than simply moving to the closest available urban area.

[18] This is to ensure decreasing returns at the city level, which, in turn, is required to have a finite city size. Alternatively, one could assume congestion in amenities or decreasing returns in housing production.

[19] In this framework, the amount of land to be developed is assumed to be given in the short run. It can be argued that, in reality, this is an endogenous outcome of factors such as regulation, city growth, and geographic constraints. In my empirical analysis I incorporate city area as a control variable and I instrument it using historical population, thus abstracting from these issues.

Equilibrium

The system of equations given by the three optimality conditions (B.1), (B.2), and (B.3) can be solved for the three endogenous variables N, W, and p_h as functions of the city-specific productivity parameter A and consumption amenities θ. Denoting with F, G, D, and K constant functions of the model's deep parameters, this yields the following:

$$\log(N) = F_N \log(A) + G_N \log(\theta) + D_N \log(\overline{L}) + K_N \tag{B.4}$$

$$\log(W) = F_W \log(A) + G_W \log(\theta) + D_W \log(\overline{L}) + K_W \tag{B.5}$$

$$\log(p_h) = F_P \log(A) + G_P \log(\theta) + D_P \log(\overline{L}) + K_P \tag{B.6}$$

where $F_N, F_W, F_P > 0$; $G_N, G_P > 0$; and $G_W < 0$. Population, wages, and rents are all increasing functions of the city-specific productivity parameter A. Intuitively, higher A allows firms to pay higher wages, which attracts households and bids up rents. Similarly, population and rents are increasing in the "quality of life" parameter θ: better amenities attract households and bid up rents. Wages are decreasing in θ because firms prefer cities with higher production amenities, whereas households prefer cities with higher consumption amenities, and factor prices - W and p_h - strike the balance between these conflicting location preferences.

Reduced-form predictions

Consider now an exogenous shifter of urban geometry S, higher values denoting less compact shapes. I hypothesize that S may be part of the A or θ bundle as follows:

$$\log(A) = \kappa_A + \lambda_A S \tag{B.7}$$

$$\log(\theta) = \kappa_\theta + \lambda_\theta S. \tag{B.8}$$

Plugging (B.7) and (B.8) into (B.4), (B.5), (B.6) yields the following reduced-form equations:

$$\log(N) = B_N S + D_N \log(\overline{L}) + K_N \tag{B.9}$$

$$\log(W) = B_W S + D_W \log(\overline{L}) + K_W \tag{B.10}$$

$$\log(p_h) = B_P S + D_P \log(\overline{L}) + K_P. \tag{B.11}$$

Suppose that non-compact shape S decreases households' indirect utility in equilibrium, but does not directly affect firms' productivity (λ_A=0, λ_θ<0). This would be the case if, for example, households located in non-compact cities faced longer commutes, or were forced to live in a less preferable location so as to avoid long commutes, while firms' transportation costs were unaffected - possibly because of better access to transportation technology, or because of being centrally located within a city. In this case, the model predicts that $B_N < 0$, $B_W > 0$, $B_P < 0$. A city with poorer shape should have, all else equal, smaller population, higher wages, and lower housing rents. Intuitively, households prefer cities with good shapes, which drives rents up and bids wages down in these locations.

Suppose, instead, that poor city geometry enters both θ and A, i.e. in equilibrium it is associated with both lower household indirect utility and lower firm productivity (λ_A<0, λ_θ<0). This would be the case if the costs of longer commutes were borne by households but also by firms. This would imply $B_N < 0$, $B_W \gtrless 0$, $B_P < 0$. The model's predictions are similar, except that the effect on wages will be ambiguous, given that now both firms and households prefer to locate in compact cities. As both firms and households compete to locate in compact cities, the net effect on W depends on whether firms or households value low S relatively more (on the margin). If S affects households more than firms, then $B_W > 0$.

Denote with $\widehat{B_N}$, $\widehat{B_W}$, and $\widehat{B_P}$ the reduced-form estimates for the impact of S on, respectively, $\log(N), \log(W)$, and $\log(p_h)$. These estimates, in conjunction with plausible values for parameters α, β, and γ, can be used to back out λ_θ and λ_A, representing respectively the marginal willingness to pay for S and the marginal productivity impact of S. Totally differentiating the indirect utility of households (B.1) with respect to S yields:

$$\frac{\partial \log(\theta)}{\partial S} = \alpha \frac{\partial \log(p_h)}{\partial S} - \frac{\partial \log(W)}{\partial S} \tag{B.12}$$

suggesting that λ_θ can be estimated as

$$\widehat{\lambda_\theta} = \alpha \widehat{B_P} - \widehat{B_W}. \tag{B.13}$$

Totally differentiating the zero-profit condition (B.2) with respect to S yields:

$$\frac{\partial \log(A)}{\partial S} = (1-\beta-\gamma) \frac{\partial \log(N)}{\partial S} + (1-\gamma) \frac{\partial \log(W)}{\partial S} \tag{B.14}$$

suggesting that λ_A can be estimated as

$$\widehat{\lambda_A} = (1-\beta-\gamma)\widehat{B_N} + (1-\gamma)\widehat{B_W}. \tag{B.15}$$

Equations (B.9), (B.10), and (B.11) are taken to the data in Section 6 in the paper. Estimates of λ_A and λ_θ are provided in Section 8.

Discussion and extensions

The framework outlined above makes a number of simplifying assumptions and modelling choices. In what follows I discuss limitations of the current framework and extensions for future research.

The model features homogeneous and perfectly mobile households. As such, it has little to say about welfare consequences of bad shape, as all agents are marginal and indifferent in equilibrium. With indirect utility pinned down by utility in a reservation location, there are no welfare gains from improving shape, as higher rents accrue to landlords. In order to be able to make welfare and distributional statements, one would require a richer model incorporating landlords and tenants as well as heterogeneity in idiosyncratic location preferences or migration costs.

However, the willingness to pay parameter λ_θ can be viewed as an upper bound for welfare effects of deteriorating shape, to the extent that reality is somewhere in between the case with infinitely elastic or infinitely inelastic supply of urban dwellers. With a fixed total urban population at the country level, equilibrium indirect utility $\overline{\upsilon}$ will unambiguously increase everywhere if amenities improve in one city. In the Cobb-Douglas case, $\log \overline{\upsilon}$ increases proportionally to the average of $\log \theta$ across cities, so that λ_θ

coincides with the welfare impact of a one unit improvement in shape in all cities. The assumption of a fixed total population is extreme, as many migrants into cities are coming from the countryside rather than reallocating across cities. The alternative extreme assumption is that of a perfectly elastic supply of migrants to cities, with indirect utility being pinned down by a reservation utility in the countryside, which delivers the prediction that any improvement in amenities will result in larger urban populations but no welfare change. This provides a lower bound of zero for the welfare effect.

In a richer model with heterogeneous households, there will be welfare impacts on inframarginal households, as the Rosen-Roback conditions continue to hold for the marginal household. As discussed in Moretti (2011), the welfare impacts on inframarginal agents will depend on the relative elasticity of labor and of housing supply. A lower local elasticity of labor implies a larger incidence on households, with the elasticity of local labor supply being ultimately governed by households' idiosyncratic preferences for locations.

With its assumption of homogeneous agents, the model also rules out sorting. The estimated compensating differentials should be thus thought of as an underestimate of true equalizing differences for those with a strong preference for compact layouts, and an overestimate for those with weak preferences.

In the model, any cost and benefit of city shape that requires cross-city compensating differentials will be part of the θ and A bundles. However, the model is agnostic on the specific channels. City shape may reduce utility because of worse service delivery in disconnected cities or because of worse accessibility. The latter includes direct costs of commuting but also "indirect" costs stemming from traffic congestion, other externalities (such as the disutility from traffic noise or pollution), or other utility costs borne to cope with bad shape - for example, households may give up certain trips or may choose to live, work, or shop in less preferred locations in order to avoid lengthy commutes. Modelling these channels is challenging and distinguishing these different costs in the data would require much more granular data at the sub-city level than what is available for India.

The costs of bad shape that are directly related to accessibility could be accounted for more explicitly by nesting a within-city model in the cross-city framework. In such a framework, some of the costs associated with longer distances could be offset at the sub-city level by the local rent gradient. In Section C in this Appendix I provide a sketch of such a model. In a monocentric, open city with topographic constraints, households directly pay for commuting costs (that are linear in distance) out of their budgets. The predictions for rents and population are consistent with my empirical findings.

The housing sector in the model features constant housing supply elasticity $(1/(\delta-1))$ across cities. Allowing for heterogeneity in housing supply elasticity across cities would not change the predicted sign of the relationship between θ, A, and the endogenous variables. However, the magnitudes would be affected: to the extent that good shape affects households' indirect utility in equilibrium, in more inelastic cities the impacts on population and wages would be attenuated and the impact on rents would be amplified (in absolute terms). The current framework could be extended allowing housing supply elasticity to be jointly determined with city shape via geography, echoing the insights of Saiz (2010). This would provide a richer characterization of the relationship between shape and supply elasticity, but empirically disentangling the supply elasticity impact would require additional orthogonal sources of variation.

While the notion of compensating differentials based on rents and wages and the model's reduced-form predictions are very general, the calculation of λ_θ and λ_A rely on particular functional forms. This framework uses standard Cobb-Douglas functional form assumptions, which imply homothetic preferences and a constant housing expenditure share. The latter assumption is in line with a large literature (e.g. Ahlfeldt et al., 2015) and finds empirical support in the US (Davis and Ortalo-Magné, 2011). Whether this is plausible for developing countries is less clear and could be addressed in future work. *A priori*, it is also unclear whether the marginal willingness to pay for "good shape" should be the same across income levels. Poorer households without access to individual means of transportation may be the ones affected the most by bad shape, in line with my findings on slum dwellers making a larger share of the population in compact cities.

Finally, the model does not allow for externalities. In the presence of congestion in consumption, $\widehat{\lambda_\theta}$ would be capturing the equilibrium effect of shape, gross of congestion, providing a lower bound for λ_θ. Similarly, in the presence of agglomeration in production, production amenities will affect productivity both directly, through λ_A, and indirectly, through their effect on city size N. If compact cities have larger populations, this will make them more productive through agglomeration, which will amplify the direct productivity impact of compactness. In this case, estimates of the productivity impact of shape will be an upper bound for λ_A.

C. Spatial Equilibrium within the City and Topographic Constraints

In this Section I provide a framework that embeds irregular city shape in a model of spatial equilibrium within, rather than across cities. I present a simple version of a monocentric city model, augmented with topographic obstacles. This allows me to focus on the implications of city shape for the distribution of households within a city and for commuting. While data availability constraints prevent me from taking this model to the data, the cross-city implications of this model are consistent with my reduced-form results. The within-city model predicts that, for given transportation costs, constrained cities are characterized by a lower population, and by average rents that may be lower or higher depending on the location and the magnitude of the constraint.

Model setup

I draw on a simple version of the monocentric city model (Alonso, 1964; Mills, 1967; Muth, 1969; Brueckner, 1987) in which city inhabitants all commute to the CBD. I consider an "open city" version of this model, in which the population of each city is endogenously determined in a way that ensures spatial equilibrium across cities. Each individual earns a wage w and consumes L units of land, both of which are fixed across locations. City dwellers face linear commuting costs τd, where d is the distance from the CBD at which they choose to live. The rental cost per unit of land at a distance d from the CBD is $r(d)$, which is endogenously determined in the model. The utility function of city dwellers is $U(C,L)$ where consumption C is equal to wage income net of housing and commuting costs

Figure C.1: Population in a linear city with a constraint

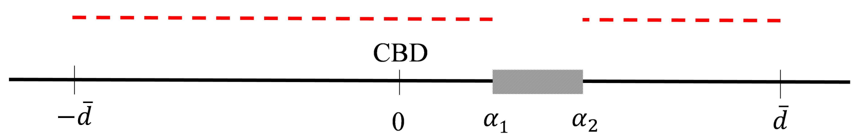

or $W - \tau d - r(d)L$. For a given city choice, inhabitants choose at which distance from the CBD to live by solving the following maximization problem:

$$\max_d U(w - \tau d - r(d)L, L) \tag{C.1}$$

which yields the Alonso-Muth condition as the first-order condition:

$$r'(d) = -\frac{\tau}{L}, \tag{C.2}$$

The rent function in the city is thus

$$r(d) = r(0) - \frac{\tau}{L}d \tag{C.3}$$

with rents declining with distance in a way that offsets the increase in transportation costs. Utility is equalized at any distance from the CBD.

Equilibrium in the constrained city

Now let us make some assumptions on the geometry of the city, and consider the case of a city with topographic or planning constraints. For the sake of simplicity and to provide closed-form solutions, I consider a linear city, in which people live on one dimension along a line. The intuitions and qualitative predictions of the model carry over to a two-dimensional city (as in the standard Alonso-Mills-Muth framework).

In the benchmark model without topographic constraints, individuals can locate on any point along the line; as a result, the distance-minimizing city structure is one in which inhabitants are symmetrically distributed along the line on either side of the CBD. In contrast, a constrained city is one in which certain locations are undevelopable. I model this by assuming that, on one side of the CBD, locations at distances between α_1 and α_2 from the CBD are unavailable, with $0 < \alpha_1 < \alpha_2$.[20] This layout is illustrated in Figure C.1. The plane in which the city is located is represented as the solid black line. Locations along the line are expressed as distances from the CBD, the position of which is normalized at 0. The constraint is represented by the hatched rectangle. For a given city population, the distance-minimizing city structure in the constrained city may become asymmetric, with a smaller fraction of the population locating on the constrained side of the line. The distribution of inhabitants under this city structure is depicted as the dashed red line in Figure C.1. The edge of the city on either side of the CBD is placed at some distance $\bar{d}$, that will be endogenously determined in the model.[21]

[20] The model's intuitions also apply to a city with multiple constraints.

[21] The benchmark, unconstrained city can be viewed as a special case of the constrained city for which $\alpha_1 = \alpha_2$ (i.e. the obstacle has size 0).

Below, I solve for the equilibrium population and rents in the constrained city. The first step is to solve the model for a city population of size N, which will be then endogenized. Assuming that N is sufficiently large relative to the size of the topographic obstacle,[22] the population in the constrained city will distribute itself around the CBD as in Figure C.1. On both sides of the CBD, the furthest occupied location will be at distance $\bar{d}$. The constrained side of the line, however, offers only $\bar{d}-(\alpha_2-\alpha_1)$ units of inhabitable land. N residents using L units of land each will require NL units of land in total, that are distributed across the two sides of the CBD.

We thus have $NL = \bar{d}+\bar{d}-(\alpha_2-\alpha_1)$, implying:

$$\overline{d} = \frac{NL+(\alpha_2-\alpha_1)}{2}. \tag{C.4}$$

In contrast, in the unconstrained city we would have $NL = 2\overline{d}$, as residents distribute themselves symmetrically on each side of the CBD, all the way to the edge of the city at distance $\overline{d}$.

Next, consider the rent function $r(d)$. Assume that rents at the city edge $\bar{d}$ are equal to $\underline{r}$, the opportunity cost of land. By setting $r(\bar{d}) = \underline{r}$ in (C.3) one can obtain $r(0) = \underline{r}+\frac{\tau}{L}\bar{d}$, which implies:

$$r(d) = \underline{r}+\frac{\tau}{L}\bar{d}-\frac{\tau}{L}d. \tag{C.5}$$

Plugging (C.4) in (C.5) yields:

$$r(d) = \underline{r}+\frac{\tau N}{2}+\frac{\tau(\alpha_2-\alpha_1)}{2L}-\frac{\tau}{L}d. \tag{C.6}$$

In the open-city framework, N is determined by utility-equalizing population flows across cities. Denoting the reservation utility as $\underline{U}$, spatial equilibrium across cities implies $U(w-\tau d-r(d)L, L) = \underline{U}$. Plugging (C.6) into the utility function, this condition becomes:

$$U\left(w-\underline{r}L-\frac{\tau NL}{2}-\frac{\tau(\alpha_2-\alpha_1)}{2}, L\right) = \underline{U}. \tag{C.7}$$

Let us now consider average rents in the constrained city. In order to derive simple closed-form solutions for N and $r(d)$, further assume that income net of commuting and housing costs in the reservation location is equal to $\underline{C}$:

$$w-\underline{r}L-\frac{\tau NL}{2}-\frac{\tau(\alpha_2-\alpha_1)}{2} = \underline{C}. \tag{C.8}$$

From (C.8) one can pin down the equilibrium N:

$$N = \frac{2(w-\underline{r}L-\underline{C})}{\tau L}-\frac{(\alpha_2-\alpha_1)}{L}. \tag{C.9}$$

Plugging (C.9) into (C.4) yields:

$$\overline{d} = \frac{(w-\underline{r}L-\underline{C})}{\tau} \tag{C.10}$$

which does not depend on the size or on the position of the topographic obstacles. Plugging the equilibrium $\overline{d}$ into (C.5) yields:

$$r(d) = \underline{r}+\frac{w-\underline{r}L-\underline{C}}{L}-\frac{\tau}{L}d. \tag{C.11}$$

[22] Specifically, NL has to be greater than $(\alpha_1+\alpha_2)$, otherwise the constraint will never be reached. This condition will be met provided that the city pays a high enough wage relative to transportation costs.

Figure C.2: Rents distribution in a linear city with a constraint

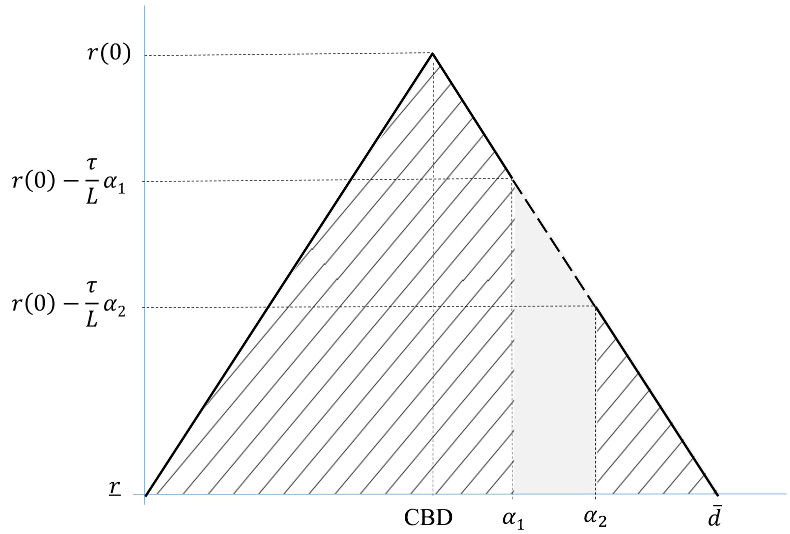

Comparative statics

Below I discuss the model's prediction for the equilibrium population and rents in a constrained versus unconstrained city. I show that population is unambiguously lower in constrained cities - a prediction which is borne in my data. On the other hand, whether rents are higher or lower in constrained cities depends on the position of the constraint, making it ultimately an empirical question.

From (C.7) and (C.9) it is apparent that the city's population N is smaller, the larger the size of the constraint $(\alpha_2 - \alpha_1)$. Intuitively, a city with topographic constraints is one in which, for a given maximum distance from the CBD, there are fewer locations available, and in equilibrium it will host a smaller population.

I next show that, all else being equal, average rents in the constrained city may be lower or higher than in the unconstrained city. Consider two cities that are identical in all parameters of the model, except for the fact that one is constrained and the other is unconstrained. The distribution of rents as a function of distance from the CBD in the constrained city is represented by the solid line in Figure C.2. The solid line plus the dashed line segment, taken together, represent rents in the unconstrained city.

Note that both cities have the same rent gradient $r(d)$ and the same equilibrium $\bar{d}$, but the constrained city is missing a portion of the distribution of rents, corresponding to the dashed segment. The hatched area in Figure C.2 corresponds to total rents in the constrained city; the hatched plus the solid area correspond to total rents in the unconstrained city. Rents per unit of land will be higher or lower in the constrained city depending on the size and position of the constraint. Intuitively, if the topographic obstacle precludes development close to the CBD, where rents would be high, average rents will be lower than in the unconstrained city. If the topographic obstacle precludes development far from the CBD, where rents would be low, average rents will be higher than in an unconstrained city. This intuition applies also to cases with multiple constraints that may introduce gaps in the rent distribution at different points.

This can be shown algebraically by computing average rents in the two cities, which can be easily done by calculating the areas of the relevant triangles and rectangles in Figure C.2. Total rents in the unconstrained city, denoted as R_U, can be calculated as:

$$R_U = \frac{2\bar{d}\,[r(0) - \underline{r}]}{2} = \frac{(w - \underline{r}L - \underline{C})^2}{\tau L}. \tag{C.12}$$

Denoting the equilibrium population in the unconstrained city with N_U, the average rent per unit of occupied land in the unconstrained city is:

$$\frac{R_U}{LN_U} = \frac{R_U\,\tau}{2\,(w - \underline{r}L - \underline{C})} = \frac{w - \underline{r}L - \underline{C}}{2L}. \tag{C.13}$$

Total rents in the constrained city, denoted as R_C, are equal to total rents in the unconstrained city minus the area of the solid trapezoid in Figure C.2, which we can denote as A:

$$R_C = R_U - A = R_U - (\alpha_2 - \alpha_1)\left[\left(r(0) - \frac{\tau}{L}\alpha_2 - \underline{r}\right) + \frac{\frac{\tau}{L}\,(\alpha_2 - \alpha_1)}{2}\right]. \tag{C.14}$$

Denoting the equilibrium population in the constrained city with N_C, the average rent per unit of land in the constrained city is thus:

$$\frac{R_C}{LN_C} = \frac{R_U - A}{L\left(N_U - \frac{(\alpha_2 - \alpha_1)}{L}\right)}. \tag{C.15}$$

Average rents in the constrained city are lower than in the unconstrained city when (C.15) is smaller than (C.13), or equivalently when:

$$\alpha_2 - \alpha_1 < \frac{LN_U}{R_U}A. \tag{C.16}$$

Plugging in the expressions for N_U, R_U and A, this inequality simplifies to:

$$\alpha_1 + \alpha_2 < \frac{w - \underline{r}L - \underline{C}}{\tau}. \tag{C.17}$$

Whether a constrained city has lower average rents than an unconstrained city depends on the size of the constraint $(\alpha_2 - \alpha_1)$ and the position of the constraint α_1. All else being equal, when the obstacle is close to the CBD (α_1 is small), the condition above is more likely to be satisfied; intuitively, topography is preventing development in a location that would be a high-rent one due to its proximity to the center. Furthermore, for a given topography, the condition above is more likely to be satisfied when wages are higher or transportation costs are lower. High wages and low transportation costs attract a larger population and make the city more spread out (leading to a larger $\bar{d}$); as a result, locations at distance α_1 from the CBD become relatively more central and demand a higher rent.

Note that the analysis above holds transportation costs constant across constrained and unconstrained cities. In a richer model, one could assume that transportation costs per unit distance are higher in cities that have irregular layouts due to topographic constraints. A standard comparative statics result in the open-city version of the monocentric city model is that cities with higher transportation costs have lower rents and smaller populations (Brueckner, 1987), which would further align the theoretical predictions with my empirical findings.

D. Signing the OLS Bias

In this Section I follow up on the discussion in Section 5 in the paper and illustrate analytically that the OLS bias from estimating the impact of city shape on population and other outcomes has an ambiguous sign. I also discuss under what conditions the bias for the impact of shape on population is positive.

Denoting city population with N and city shape with S, consider the following version of the estimating equation, where for simplicity I have dropped subscripts and additional regressors:

$$logN = \alpha_1 S + u_1 \tag{D.1}$$

In what follows, I show analytically that $cov(S, u_1) \neq 0$. City shape is the result of an interaction of exogenous determinants (such as topographic obstacles) and endogenous determinants. While the endogenous determinants are manifold, to fix ideas, consider two factors: local institutional capacity and highways connecting cities. With good local institutional capacity, urban planners can encourage compact development through well-enforced master plans and land use regulations. Highways connecting into a city can affect urban form by encouraging urban development along transit corridors, which has been associated with sprawl (Baum-Snow, 2007) and potentially deteriorates urban shape. In Section 5 in the paper I discuss more factors that affect city shape and that would generate selection effects similar to those highlighted here. Denoting (exogenous) geographic predictors of city shape as $\tilde{S}$, local institutional capacity as $Inst$, and road infrastructure as $Infra$, S can be written as a function of its determinants:

$$S = \beta \tilde{S} + \delta_1 Inst + \delta_2 Infra + \eta \tag{D.2}$$

Assume that $\beta > 0$ (potential shape predicts actual shape), $\delta_1 < 0$ (better institutional capacity makes cities more compact), and $\delta_2 > 0$ (highways cause sprawl). Further assume that $cov(Inst, \eta) = cov(Infra, \eta) = cov(S, \eta) = 0$.

The endogeneity problem associated with $Inst$ and $Infra$ stems from the fact that city size N affects both local institutional capacity and highways. Larger cities have both greater institutional capacity and more infrastructural investment. We thus have:

$$Inst = \gamma_1 logN + \xi_1 \tag{D.3}$$

$$Infra = \gamma_2 logN + \xi_2 \tag{D.4}$$

where $\gamma_1 > 0$ and $\gamma_2 > 0$ and $cov(logN, \xi_1) = cov(logN, \xi_2) = 0$.

Plugging (D.3) and (D.4) into (D.2) one obtains:

$$S = \beta \tilde{S} + \alpha_2 logN + \tilde{\eta} \tag{D.5}$$

where $\alpha_2 = \delta_1 \gamma_1 + \delta_2 \gamma_2$ and $\tilde{\eta} = \delta_1 \xi_1 + \delta_2 \xi_2 + \eta$.

In equilibrium, city shape is a function of exogenous geographic predictors plus a term that depends endogenously on population. The sign of α_2 is ambiguous because $\delta_1 < 0$ and $\delta_2 > 0$. If the effect of highways is stronger than that of institutional capacity, α_2 will be positive. The descriptive patterns in

Table A2 as well as the OLS impacts in Table3 indicate a positive correlation between city size N and shape S, suggesting that empirically α_2 is positive.

Solving the system of equations given by (D.5) and (D.1) one obtains:

$$S = \frac{\beta}{1-\alpha_2\alpha_1}\tilde{S} + \frac{\alpha_2 u_1}{1-\alpha_2\alpha_1} + \frac{\tilde{\eta}}{1-\alpha_2\alpha_1} \tag{D.6}$$

Going back to estimating equation (D.1), one can now compute the covariance between the error term u_1 and S as follows:

$$cov(S,u_1) = cov\Big(\frac{\alpha_2}{1-\alpha_2\alpha_1}u_1, u_1\Big) = \sigma^2_{u_1}\big(\frac{\alpha_2}{1-\alpha_2\alpha_1}\big) \tag{D.7}$$

In general, (D.7) has an ambiguous sign. As a result, the OLS bias in estimating the impact of S in (D.1) could be positive or negative. However, if $\alpha_2 > 0$ (the equilibrium correlation between shape and population is positive) and $\alpha_1 < 0$ (the structural effect of shape on population is negative), then $(\frac{\alpha_2}{1-\alpha_2\alpha_1})$ will be unambiguously positive and the OLS estimate will be biased towards positive values. This is indeed what the estimates show, with negative IV impacts and positive OLS impacts of city shape on population growth.

E. Single-Instrument Approach

In this Section I outline an alternative implementation of my instrumental variables strategy, which I employ in the robustness check presented in Table 8 and discussed in Section 7.

Recall from Section 5 in the paper that in the benchmark estimation I consider city shape and area as two endogenous regressors, and I employ as instrumental variables "potential" city shape and projected historical population. Potential shape is calculated based on the relative position of topographic obstacles encountered as a city grows along a predicted expansion path, which in turn depends on the city's own projected population growth. Below I present an alternative strategy that does not rely on projected historical population. Relative to the baseline approach, this entails two differences: first, the instrument is constructed using a completely mechanical model for city expansion that postulates the same rate of expansion for all cities. Second, city area is not directly controlled for in the estimating equations, but instead both right- and left-hand side variables are normalized by area.

The first step is to determine $\widehat{r}_{c,t}$, i.e. the predicted city radius within which the potential footprint is constructed. Under this alternative approach I do so by postulating that cities expand at the same rate, equivalent to the average expansion rate across all cities in the sample. Specifically, the steps involved are the following:

(i) Denoting the area of city c's actual footprint in year t as $area_{c,t}$, I pool together the 1951-2010 panel of cities and estimate the following regression:

$$\log(area_{c,t}) = \theta_c + \gamma_t + \varepsilon_{c,t} \tag{E.1}$$

where θ_c and γ_t denote city and year fixed effects.

(ii) From the regression above, I obtain $\widehat{area_{c,t}}$, and corresponding $\widehat{r_{c,t}} = \sqrt{\frac{\widehat{area_{c,t}}}{\pi}}$.

The second modification relative to the baseline approach is in the estimating equations. Instead of controlling for city area explicitly, this approach relies on normalizing both right- and left-hand side variables by city area, regressing population density on the normalized version of the shape index.

Define population density[23] as

$$d_{c,t} = \frac{pop_{c,t}}{area_{c,t}}$$

and denote the normalized version of shape as nS.

The alternative estimating equation becomes:

$$d_{c,t} = a \cdot nS_{c,t} + \mu_c + \rho_t + \eta_{c,t} \tag{E.2}$$

which includes the endogenous regressor $nS_{c,t}$. This is a counterpart of equation (8) in the paper.

The corresponding first-stage equation is

$$nS_{c,t} = \beta \cdot \widetilde{nS_{c,t}} + \lambda_c + \gamma_t + \varepsilon_{c,t}. \tag{E.3}$$

where $\widetilde{nS_{c,t}}$ is namely the normalized shape index computed for the potential footprint.

F. Nonparametric Employment Subcenter Identification (McMillen, 2001)

In order to compute the number of employment subcenters in each city, used as the dependent variable in Table 10, I employ the two-stage, non-parametric approach described in McMillen (2001). This procedure identifies employment subcenters as locations that have significantly larger employment density than nearby ones, and that have a significant impact on the overall employment density function in a city. The data on firms' location used as input in this procedure is discussed in Section A6 above.

The procedure outlined below is performed separately for each city in the 2005 sample. As units of observation within each city, I consider grid cells of 0.01 degree latitude by 0.01 degree longitude, with an area of approximately one square kilometer. I calculate a proxy for employment density in each cell, by considering establishments from the 2005 Economic Census located in that cell and summing their reported number of employees.[24] In order to define the CBD using a uniform criterion for all cities, I consider the centroid of the 1951 footprint. Results are similar using the 2005 centroid as an alternative definition.

In the first stage of this procedure, "candidate" subcenters are identified as those grid cells with significant positive residuals in a smoothed employment density function. Let y_i be the log employment density in grid cell i; denote with x_i^N its distance north from the CBD, and with x_i^E its distance east. Denoting the error term with ε_i, I estimate:

$$y_i = f(x_i^N, x_i^E) + \varepsilon_i \tag{F.1}$$

[23] Note that this does not coincide with population density as defined by the Census, which reflects administrative boundaries.

[24] The Directory of Establishments provides establishment-level employment only by broad categories, indicating whether the number of employees falls in the 10-50, 51-100, or 101-500 range, or is larger than 500. In order to assign an employment figure to each establishment, I consider the lower bound of the category.

using locally weighted regression, employing a tricube kernel and a 50% window size. This flexible specification allows for local variations in the density gradient, which are likely to occur in cities with topographic obstacles. Denoting with $\widehat{y_i}$ the estimate of y for cell i, and with $\widehat{\sigma_i}$ the corresponding standard error, candidate subcenters are grid cells such that $(y_i - \widehat{y_i}) / \widehat{\sigma_i} > 1.96$.

The second stage of the procedure selects those locations, among candidate subcenters, that have significant explanatory power in a semiparametric employment density function estimation. Let D_{ij} be the distance between cell i and candidate subcenter j, and denote with $DCBD_i$ the distance between cell i and the CBD. With S candidate subcenters, denoting the error term with u_i, the semi-parametric regression takes the following form:

$$y_i = g(DCBD_i) + \sum_{j=1}^{S} \delta_j^1 (D_{ji})^{-1} + \delta_j^2 (-D_{ji}) + u_i. \tag{F.2}$$

In the specification above, employment density depends non-parametrically on the distance to the CBD, and parametrically on subcenter proximity, measured both in levels and in inverse form. This parametric specification allows us to conduct convenient hypothesis tests on the coefficients of interest δ_j^1 and δ_j^2. (F.2) is estimated omitting cells i corresponding to one of the candidate subcenters or to the CBD. I approximate $g(.)$ using cubic splines.

If j is indeed an employment subcenter, the variables $(D_j)^{-1}$and/or $(-D_j)$ should have a positive and statistically significant impact on employment density y. One concern with estimating (F.2) is that, with a large number of candidate subcenters, the distance variables D_{ij} can be highly multicollinear. To cope with this problem, a stepwise procedure is used to select which subcenter distance variables to include in the regression. In the first step, all distance variables are included. At each step, the variable corresponding to the lowest t-statistic is dropped from the regression, and the process is repeated until all subcenter distance variables in the regression have a positive coefficient, significant at the 20% level. The final list of subcenters includes the sites with positive coefficients on either $(D_j)^{-1}$ or $(-D_j)$.

References

[1] Ahlfeldt, G. M., S. J., Redding, D. M. Sturm and N. Wolf (2015), "The Economics of Density: Evidence from the Berlin Wall", Econometrica, 83 (6), 2127-2189.

[2] Akbar, P., Couture, V., Duranton, G., and A. Storeygard (2019), "Mobility and Congestion in Urban India", working paper.

[3] Akbar, P., Couture, V., Duranton, G., and A. Storeygard (2018), "Accessibility in Urban India", working paper.

[4] Alonso, W. (1964), *Location and Land Use*, Cambridge: Harvard University Press.

[5] Angel, S., J. Parent, and D. L. Civco (2009), "Ten Compactness Properties of Circles: A Unified Theoretical Foundation for the Practical Measurement of Compactness", The Canadian Geographer, 54 (4), 441-461.

[6] Balk, D. L., U. Deichmann, G. Yetman, F. Pozzi, S. I. Hay and A. Nelson (2006), "Determining Global Population Distribution: Methods, Applications and Data", Advances in Parasitology, 62, 119-156.

[7] Barr, J., T. Tassier and R. Trendafilov (2011), “Bedrock Depth and the Formation of the Manhattan Skyline, 1890-1915”, Journal of Economic History, 71 (4), 1060-1077.

[8] Baruah, N., V. J. Henderson, and C. Peng (2017), “Colonial Legacies: Shaping African Cities”, SERC Urban and Spatial Programme Discussion Paper.

[9] Baum-Snow, N. (2007), “Did Highways Cause Suburbanization?”, Quarterly Journal of Economics 122 (2), 775-805.

[10] Bertaud, A. and J. K. Brueckner (2005), “Analyzing Building-Height Restrictions: Predicted Impacts and Welfare Costs”, Regional Science and Urban Economics, 35 (2), 109-125.

[11] Brueckner, J. K. (1987), “The Structure of Urban Equilibria: A Unified Treatment of the Muth-Mills Model”, in E. S. Mills (ed.), *The Handbook of Regional and Urban Economics*, Amsterdam: North Holland Press.

[12] Carroll, M., J. Townshend, C. DiMiceli, P. Noojipady, and R. Sohlberg (2009), “A New Global Raster Water Mask at 250 Meter Resolution”, International Journal of Digital Earth, 2 (4), 291-308.

[13] Centre for Industrial and Economic Research (CIER) (1990), *Industrial Databook 1990*, New Delhi: CIER.

[14] Centre for Monitoring India Economy, *https://www.districtsofindia.com/*.

[15] Chauvin, J. P., E. Glaeser, Y. Ma and K. Tobio (2017), “What is Different about Urbanization in Rich and Poor countries? Cities in Brazil, China, India and the United States”, Journal of Urban Economics, 98, 17-49.

[16] CIESIN - Columbia University, IFPRI, The World Bank, and CIAT (2011), *Global Rural-Urban Mapping Project, Version 1 (GRUMPv1): Settlement Points,* Palisades, NY: NASA Socioeconomic Data and Applications Center (SEDAC).

[17] Davis, M.A. and F. N. Ortalo-Magné (2011), “Household Expenditures, Wages, Rents”, Review of Economic Dynamics, 14 (2), 248-261.

[18] Dev, S. (2006), “Rent Control Laws in India: A Critical Analysis”, CCS Working Paper No. 158, Centre for Civil Society, New Delhi.

[19] Director of Census Operations, *Census of India 2011*, New Delhi: Office of the Registrar General & Census Commissioner.

[20] Fernandes, A. and G. Sharma (2012), “Determinants of Clusters in Indian Manufacturing: The Role of Infrastructure, Governance, Education, and Industrial Policy”, IGC working paper.

[21] Glaeser, E. (2008),*Jall Cities, Agglomeration and Spatial Equilibrium*, Oxford: Oxford University Press.

[22] Greenstone, M. and R. Hanna (2014), “Environmental Regulations, Air and Water Pollution, and Infant Mortality in India”, American Economic Review, 104 (10), 3038-72.

[23] Henderson, M., E. Yeh, P. Gong and C. Elvidge (2003), “Validation of Urban Boundaries Derived from Global Night-time Satellite Imagery”, International Journal of Remote Sensing, 24 (3), 595-609.

[24] Hengl, T., J. Mendes de Jesus, G. B. M. Heuvelink, M. Ruiperez Gonzalez, M. Kilibarda et al. (2017), “SoilGrids250m: Global Gridded Soil Information Based on Machine Learning”, PLoS ONE 12(2): e0169748.

[25] Iyer, L. (2010), “Direct versus Indirect Colonial Rule in India: Long-Term Consequences”, Review of Economics and Statistics, 92 (4), 693-713.

[26] Joshi, P. K., B. M. Bairwa, R. Sharma, V. S. P. Sinha (2011), "Assessing Urbanization Patterns over India Using Temporal DMSP-OLS Night-time Satellite Data", Current Science, 100 (10), 1479-1482.

[27] McMillen, D. P. (2001), "Nonparametric Employment Subcenter Identification", Journal of Urban Economics, 50 (3), 448-473.

[28] McMillen, D. P. and S. C. Smith (2003), "The Number of Subcenters in Large Urban Areas", Journal of Urban Economics, 53 (3), 321-338.

[29] Mills, E. S. (1967) "An Aggregative Model of Resource Allocation in a Metropolitan Area", American Economic Review, 57 (2), 197-210.

[30] Mitra, A. (1980), *Population and Area of Cities, Towns, and Urban Agglomerations, 1872-1971*, Bombay: Allied.

[31] Moretti, E. (2011), "Local Labor Markets", in O. Ashenfelter and D. Card, eds., Handbook of Labor Economics, 14, 1237-1313.

[32] Munshi, K. and M. Rosenzweig (2016), "Networks and Misallocation: Insurance, Migration, and the Rural-Urban Wage Gap", American Economic Review, 106 (1), 46-98.

[33] Muth, R. F. (1969), *Cities and Housing*, Chicago: University of Chicago Press.

[34] NASA and Ministry of Economy, Trade and Industry of Japan (METI), Land Processes Distributed Active Archive Center (LP DAAC) (2011), *ASTER Global Digital Elevation Model, Version 2*, USGS/Earth Resources Observation and Science (EROS) Center, Sioux Falls, South Dakota.

[35] Nordhaus, W., Q. Azam, D. Corderi, K. Hood, N. Victor, M. Mohammed, A. Miltner and J. Weiss (2006), "The G-Econ Database on Gridded Output: Methods and Data", http://gecon.yale.edu/

[36] Office of the Registrar General of India and Census Commissioner (2011), Provisional Results of 2011 Census, Office of the Registrar General, Government of India, New Delhi.

[37] Parent, J., Civco, D. L., and S. Angel (2009),"Shape Metrics", ESRI working paper.

[38] Roback, J. (1982), "Wages, Rents and the Quality of Life", Journal of Political Economy, 90 (6), 1257-1278.

[39] Rosen, S. (1979), "Wage-Based Indexes of Urban Quality of Life", in P. Mieszkowski and M. Strazheim (eds.), *Current Issues in Urban Economics*, Baltimore: Johns Hopkins University Press.

[40] Roychowdhury, K., S. D. Jones and C. Arrowsmith (2009), "Assessing the Utility of DMSP/OLS Night-time Images for Characterizing Indian Urbanization", 2009 IEEE Urban Remote Sensing Joint Event, Shanghai, China.

[41] Saiz, A. (2010), "The Geographic Determinants of Housing Supply", Quarterly Journal of Economics, 125 (3), 1253-1296.

[42] Small, C., F. Pozzi and C. D. Elvidge (2005), "Spatial Analysis of Global Urban Extent from DMSP-OLS Night Lights", Remote Sensing of Environment, 96 (3), 277-291.

[43] Small, C. and C. D. Elvidge (2013), "Night on Earth: Mapping Decadal Changes of Anthropogenic Night Light in Asia", International Journal of Applied Earth Observation and Geoinformation, 22, 40-52.

[44] Sridhar, K. S. (2010), "Impact of Land Use Regulations: Evidence From India's Cities", Urban Studies, 47 (7), 1541-1569.

[45] U.S. Army Map Service (ca. 1950), *India and Pakistan Topographic Maps, Series U502, 1:250,000,* U.S. Army Map Service Topographic Map Series.

Table A1: List of most and least compact cities

Rank	City	Shape (normalized)
	Top 10 most compact cities	
1	Rajkot, Gujarat	0.924
2	Kannur, Kerala	0.934
3	Bhopal, Madhya Pradesh	0.943
4	Lucknow, Uttar Pradesh	0.943
5	Meerut, Uttar Pradesh	0.943
6	Thrissur, Kerala	0.945
7	Nashik, Maharashtra	0.948
8	Jaipur, Rajasthan	0.948
9	Jabalpur, Madhya Pradesh	0.952
10	Gwalior, Madhya Pradesh	0.956
	Top 10 least compact cities	
1	Asansol, West Bengal	1.625
2	Jharia-Dhanbad, Jharkhand	1.180
3	Kolkata, West Bengal	1.128
4	Ludhiana, Punjab	1.124
5	Surat, Gujarat	1.111
6	Aurangabad, Maharashtra	1.108
7	Visakhapatnam, Ahndra Pradesh	1.108
8	Patna, Bihar	1.100
9	Amritsar, Punjab	1.100
10	Chennai, Tamil Nadu	1.081

Note: Sample of cities with million-plus population in 2011, ranked by normalized shape in 2010. The normalized shape index has a mean of 0.96 and a standard deviation of 0.07 in 2010.

Table A2: Descriptive correlations

	(1)	(2)	
Dependent variable:	*Shape, 2010*	*Δ Shape, 2010-1950*	*Obs.*
Panel A: 1951 city size quartiles			
Quartile I	-1.188***	-1.152***	351
	(0.195)	(0.255)	
II	-0.789***	-0.885***	351
	(0.192)	(0.250)	
III	-0.109	-0.491*	351
	(0.194)	(0.297)	
IV	2.514***	3.075***	351
	(0.403)	(0.538)	
Panel B: Public services and accessibility			
Share households with electricity, 2011	7.150***	13.50***	351
	(2.080)	(2.876)	
Share households with tap water, 2011	1.892***	4.092***	351
	(0.576)	(0.800)	
Share households with cars, 2011	10.72***	17.92***	351
	(2.295)	(3.248)	
Urban road length, km, 2019	-0.00136***	0.00100***	351
	(0.000477)	(0.000182)	
District avg. distance to work, km, 2011	0.176*	0.567***	208
	(0.101)	(0.201)	
Panel C: Pre-determined characteristics			
Elevation, 100 m	-0.0200	0.0151	351
	(0.0326)	(0.0569)	
Distance from the coast, km	0.000136	0.000508	351
	(0.000324)	(0.000501)	
Distance from nearest river or lake, km	5.73e-05	0.00649	351
	(0.00510)	(0.00816)	
Distance from nearest mineral deposit, km	-0.000623	-0.00336	351
	(0.00133)	(0.00244)	
Ruggedness, m	0.00143**	0.00207*	351
	(0.000709)	(0.00116)	
Bedrock depth, m	-0.0230	0.00629	351
	(0.0220)	(0.0392)	
Crop suitability, tons per hectare	0.120	-0.406	351
	(0.274)	(0.423)	
Panel D: Non-pre-determined characteristics			
Shape in 1950, km	0.818**	-0.228	351
	(0.324)	(0.808)	
Distance from state headquarters, km	-0.000623	-0.00147*	351
	(0.000562)	(0.000813)	
Distance from district headquarters, km	-0.00988*	-0.00676	351
	(0.00533)	(0.00776)	
Distance from nearest city, km	-0.00342	-0.00513	351
	(0.00210)	(0.00375)	
British direct rule	0.226	0.0105	351
	(0.156)	(0.287)	
State Capital	0.957	4.103***	351
	(1.068)	(1.123)	
Control	Area 2010	Area 1950	

Notes: this table reports pairwise correlations between levels and changes in shape and city attributes. Each row reports a coefficient from an OLS regression of shape in 2010 (col. 1) and the 2010-1950 difference in shape (col. 2) on the attribute indicated in each row, controlling respectively for city area in 2010 (col.1) and in 1950 (col. 2). A description of the variables is provided in Section A.4 in the Appendix. Summary statistics are in Table A3. Robust standard errors in parentheses.*** p<0.01,** p<0.05,* p<0.1.

Table A3: Additional summary statistics

	Obs.	*Mean*	*St. Dev.*	*Min*	*Max*
Elevation, m	351	262.20	230.88	0	1590
Distance from the coast, km	351	400.41	334.06	0.14	1315.88
Distance from nearest river or lake, km	351	15.09	16.47	0.01	102.69
Distance from nearest mineral deposit, km	351	77.21	62.18	0.05	315.18
Ruggedness, m	351	102.39	132.38	0	1000.00
Bedrock depth, m	351	5.74	3.98	1.36	19.77
Crop suitability, tons per hectare	351	1.43	0.32	0.23	1.93
Initial shape, km	351	1.01	0.71	0.35	5.82
British direct rule dummy	351	0.66			
State capital dummy	351	0.05			
Distance from state headquarters, km	351	296.56	190.41	0	998.00
Distance from district headquarters, km	351	15.53	28.25	0	138.00
Distance from nearest city, km	351	40.68	41.75	0	342.00
Share households with electricity, 2011	351	0.94	0.06	0.65	1
Share households with tap water, 2011	351	0.56	0.21	0.06	0.94
Share households with cars, 2011	351	0.08	0.05	0.01	0.29
Urban road length, km, 2019	351	773	2613	3.2	35148
Average distance to work, km, 2011	208	5.73	1.26	2.99	11.34

Notes: this table provides summary statistics for the city-level variables employed in Table A2 and in the robustness checks. These variables are described in Section A.4 in the Appendix.

Table A4: First stage and impact of city shape on population, panel results

	(1)	(2)	(3)	(4)	(5)	(6)
	First Stage		IV	OLS	IV	OLS
Dependent variable:	*Shape, km*	*Log area*	*Log population*			
Potential shape, km	1.397***	0.151***				
	(0.228)	(0.0451)				
Log projected population	-1.188***	0.297**				
	(0.269)	(0.117)				
Shape, km			-0.0975**	0.0250***	-0.107**	0.0247***
			(0.0387)	(0.00788)	(0.0442)	(0.00792)
Log area			0.783***	0.165***	0.827***	0.168***
			(0.182)	(0.0309)	(0.211)	(0.0331)
Observations	6,173	6,173	1,325	1,325	1,135	1,135
AP F stat shape	86.96	86.96	69.36		57.34	
AP F stat area	17.76	17.76	14.12		11.36	
KP F stat	21.13	21.13	16.37		13.29	
City FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
Sample	Full	Full	Full	Full	Long diff	Long diff

Notes: this table presents the main results in panel format. Each observation is a city-year. Cols. 1 and 2 are similar to cols. 3 and 4 in Table 2, but employ the full panel of cities. Cols. 3 and 4 present the panel version of cols. 1 and 2 of Table 3, estimated using the full panel of cities. Cols. 5 and 6 repeat the same specifications, but for the panel of 351 cities employed in the main long-differences specification. Cols. 1 and 2 show the first stage, estimated over years 1950 and 1992 through 2010. Cols. 3 through 6 show the IV (odd cols.) and OLS (even cols.) estimates of the impact of city shape (in km) and log city area on log population, using data from Census years 1951, 1991, 2001, and 2011. Angrist-Pischke and Kleibergen-Paap F statistics are reported. All specifications include city and year fixed effects. Standard errors clustered at the city level in parentheses. *** p<0.01,** p<0.05,* p<0.1.

Table A5: Robustness to alternative luminosity thresholds

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	First Stage		IV	OLS	First Stage		IV	OLS
Dependent variable:	*Δ Shape, km*	*Δ Log area*	*Δ Log population*	*Δ Log population*	*Δ Shape, km*	*Δ Log area*	*Δ Log population*	*Δ Log population*
Δ Potential shape, km	1.947***	0.214***			1.806***	0.240***		
	(0.191)	(0.0471)			(0.221)	(0.0523)		
Δ Log projected population	-1.940***	0.118			-1.985***	0.0664		
	(0.469)	(0.123)			(0.455)	(0.144)		
Δ Shape, km			-0.0901**	0.0271***			-0.125**	0.0199**
			(0.0377)	(0.00862)			(0.0543)	(0.00857)
Δ Log area			0.839***	0.197***			0.909***	0.226***
			(0.225)	(0.0349)			(0.255)	(0.0431)
Observations	374	374	374	374	320	320	320	320
AP F stat shape	33.42	33.42	33.42		26.34	26.34	26.34	
AP F stat area	10.83	10.83	10.83		8.42	8.42	8.42	
KP F stat	16.49	16.49	16.49		11.77	11.77	11.77	
Luminosity threshold	30	30	30	30	40	40	40	40
Mean dep var in levels, 2010	5.01	123.00	568,361	568,361	4.40	106.23	708,939	708,939
Mean dep var in levels, 1950	0.98	3.49	98,122	98,122	1.02	3.85	113,205	113,205

Notes: this table presents estimates of the first stage and the impact of shape on population, obtained using different definitions of urban footprints. Cols.1, 2, 5, and 6 report the first stage (analogous to cols. 1 and 2 in Table 2). Cols. 3 and 7 (4 and 8) report the IV (OLS) impact of city shape on population (similar to Table 3). The dependent variables and regressors are all defined as long differences 2010-1950 (2011-1951 for population). The luminosity threshold used to define urban areas is 30 in cols. 1 through 4 and 40 in cols. 5 through 8 (the baseline in the paper is 35). Angrist-Pischke and Kleibergen-Paap F statistics are reported. Robust standard errors in parentheses. *** p<0.01,** p<0.05,* p<0.1.

Table A6: First stage, alternative shape indicators

Shape metric	*A. Remoteness*		*B. Spin*		*C. Range*	
	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	*Δ Shape, km*	*Δ Log area*	*Δ Shape, km^2*	*Δ Log area*	*Δ Shape, km*	*Δ Log area*
Δ Potential shape	1.110***	0.178***	1.128***	0.00266	2.698***	0.0998***
	(0.165)	(0.0432)	(0.321)	(0.00189)	(0.368)	(0.0226)
Δ Log projected population	-0.533*	0.232**	-1.032	0.464***	-7.567***	0.0443
	(0.276)	(0.116)	(6.823)	(0.106)	(1.558)	(0.130)
Observations	351	351	351	351	351	351
AP F stat shape	26.74	26.74	27.44	27.44	30.26	30.26
AP F stat area	9.19	9.19	19.57	19.57	10.13	10.13
KP F stat	11.7	11.7	16.62	16.62	13.64	13.64
Avg. shape 1950	0.75		1.06		2.95	
Avg. shape 2010	3.45		25.72		13.57	

Notes: this table presents estimates of the first stage for alternative shape metrics. Odd (even) cols. are analogous to col. 1 (2) in Table 2. The dependent variables and regressors are all defined as long differences 2010-1950. The unit of the shape metrics is km, except for the spin index that is in square km. The shape indexes are discussed in Section A.2 in the Appendix. Remoteness (cols. 1 and 2) is the average distance to the centroid. Spin (cols. 3 and 4) is the average squared length of distances to the centroid. Range (cols. 5 and 6) is the maximum distance between two points on the outline of the city. Angrist-Pischke and Kleibergen-Paap F statistics are reported. Robust standard errors in parentheses. *** p<0.01,** p<0.05,* p<0.1.

Table A7: Impact of city shape on population, robustness to alternative shape indicators

Dependent variable:	*Δ Log population, 2011-1951*					
Shape metric	***A. Remoteness***		***B. Spin***		***C. Range***	
	(1)	(2)	(3)	(4)	(5)	(6)
	IV	OLS	IV	OLS	IV	OLS
Δ Shape	-0.118**	0.0306***	-0.00129	0.000891***	-0.0277**	0.00675***
	(0.0557)	(0.00997)	(0.000810)	(0.000264)	(0.0125)	(0.00220)
Δ Log area	0.817***	0.212***	0.587***	0.243***	0.812***	0.216***
	(0.226)	(0.0339)	(0.128)	(0.0294)	(0.222)	(0.0329)
Observations	351	351	351	351	351	351
AP F stat shape	26.74		27.44		30.26	
AP F stat area	9.19		19.57		10.13	
KP F stat	11.7		16.62		13.64	

Notes: this table presents estimates of the relationship between city shape and population for alternative shape metrics. Odd (even) cols. are analogous to col. 1 (2) in Table 3. The corresponding first stage is reported in Table A6. The regressors are defined as long differences 2010-1950. The unit of the shape metrics is km, except for the spin index that is in square km. Angrist-Pischke and Kleibergen-Paap F statistics are reported. Robust standard errors in parentheses. *** p<0.01,** p<0.05,* p<0.1.

Table A8: Impact of city shape on rents, robustness

Dependent variable:	*Δ Log rent 2008-2006, excluding bottom 25%*					
	(1)	(2)	(3)	(4)	(5)	(6)
	IV	OLS	IV	OLS	IV	OLS
	All districts		*Only districts with one city*		*Only top city per district*	
Δ Shape, km	-0.663	0.00421	-0.532	-0.00729	-0.769	0.0188
	(0.557)	(0.0487)	(0.333)	(0.0692)	(0.704)	(0.0493)
Δ Log area	-2.535	-0.0125	-1.354	-0.103	-2.138	-0.0574
	(2.300)	(0.0927)	(1.131)	(0.112)	(2.282)	(0.100)
Observations	262	262	134	134	215	215
AP F stat shape	9.60		14.77		5.11	
AP F stat area	3.00		6.12		2.80	
KP F stat	1.67		2.93		1.20	

Notes: this table is analogous to Table 5, but the district averages of rents exclude the bottom 25% of the rents distribution. Robust standard errors in parentheses. *** p<0.01,** p<0.05,* p<0.1.

Table A9: Impact of city shape on rents, IHDS data

Dependent variable:	*Δ Log rent 2010-2005*		*Δ Log rent residual 2010-2005*	
	(1)	(2)	(3)	(4)
	IV	OLS	IV	OLS
Δ Shape, km	-0.0624	0.0121	-0.0382	0.0250
	(0.0761)	(0.0164)	(0.0789)	(0.0190)
Δ Log area	-0.292	0.203*	-0.279	0.193+
	(0.736)	(0.116)	(0.766)	(0.128)
Observations	111	111	111	111
AP F stat shape	6.85		6.85	
AP F stat area	4.11		4.11	
KP F stat	2.25		2.25	
Source	IHDS	IHDS	IHDS	IHDS

Notes: this table is analogous to Table 5, cols.1 and 2, but uses rents from a different source. Each observation is a district. In cols. 1 and 2 the dependent variable is the 2010-2005 long difference in log monthly total rents, averaged at the district level, from the IHDS dataset. In cols. 3 and 4 the dependent variable is the long difference in the log average rent residual, from a hedonic regression of rent on housing characteristics discussed in Section A in the Appendix. Robust standard errors in parentheses. *** p<0.01,** p<0.05,* p<0.1.

Table A10: Pairwise correlations between instruments and city characteristics

	(1)	(2)
Dependent variable:	*Δ Potential shape, km, 2010-1950*	*Δ Log projected population, 2010-1950*
Elevation, 100 m	-0.0245	0.0127
	(0.0396)	(0.0121)
Distance from the coast, km	-0.000303	-0.000242***
	(0.000262)	(9.16e-05)
Distance from nearest river or lake, km	-0.00307	0.00340
	(0.00473)	(0.00212)
Distance from nearest mineral deposit, km	-0.00301*	-0.000135
	(0.00169)	(0.000675)
Ruggedness, m	0.000699	0.000368*
	(0.000507)	(0.000221)
Bedrock depth, m	0.00455	-0.0113*
	(0.0232)	(0.00633)
Crop suitability	0.299	0.247*
	(0.306)	(0.135)
Observations	351	351

Notes: this table reports estimates of the relationship between the instruments and time-invariant city characteristics. Each row reports a coefficient from an OLS regression of the 2010-1950 long differences in potential shape (col. 1) and log projected population (col. 2) on the controls indicated in each row. The controls are described in Section A.4 in the Appendix. Robust standard errors in parentheses.*** p<0.01,** p<0.05,* p<0.1.

Table A11: IV impact of city shape on population, robustness to sample cuts

	Dependent variable: Δ Log population, 2011-1951						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Δ Shape, km	-0.104**	-0.0999**	-0.155*	-0.175**	-0.0857**	-0.0932**	-0.0854**
	(0.0475)	(0.0486)	(0.0800)	(0.0748)	(0.0412)	(0.0422)	(0.0408)
Δ Log area	0.856***	0.863***	1.103***	0.961***	0.804***	0.818***	0.812***
	(0.246)	(0.265)	(0.415)	(0.321)	(0.234)	(0.220)	(0.233)
Observations	337	334	242	204	318	316	316
AP F stat shape	29.02	24.16	8.91	9.22	27.33	28.51	27.27
AP F stat area	8.72	7.76	4.13	4.03	9.14	10.13	8.62
KP F stat	12.36	10.92	6.97	9.96	12.18	14.09	12.10
Excluded cities	Mountainous	Coastal	River/lake	Mineral	Top 90% ruggedness	Top 90% bedrock depth	Top 90% crop suitability

Notes: this table reports the same IV specification as in Table 3, col.1, for various sample cuts discussed in Section 7. A description of the controls is provided in Section A.4 in the Appendix. Robust standard errors in parentheses.*** p<0.01,** p<0.05,* p<0.1.

Table A12: Robustness to confounding trends, non-predetermined characteristics

Characteristic:	*A. Initial shape*			*B. British direct rule*			*C. State capital*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	First Stage		IV	First Stage		IV	First Stage		IV
Dependent variable:	*Δ Shape, km*	*Δ Log area*	*Δ Log population*	*Δ Shape, km*	*Δ Log area*	*Δ Log population*	*Δ Shape, km*	*Δ Log area*	*Δ Log population*
Δ Potential shape, km	1.566***	0.319***		1.943***	0.239***		1.700***	0.222***	
	(0.203)	(0.0480)		(0.250)	(0.0490)		(0.186)	(0.0489)	
Δ Log projected population	-1.542***	-0.112		-2.228***	0.0389		-2.025***	0.0555	
	(0.439)	(0.127)		(0.486)	(0.132)		(0.396)	(0.132)	
Δ Shape, km			-0.271**			-0.0963**			-0.111***
			(0.116)			(0.0446)			(0.0418)
Δ Log area			1.337***			0.851***			0.814***
			(0.456)			(0.238)			(0.224)
Control	1.764***	-0.409***	1.060**	-0.0654	-0.221**	-0.0109	6.228***	0.271	0.940***
	(0.435)	(0.111)	(0.468)	(0.259)	(0.0904)	(0.0915)	(1.173)	(0.209)	(0.285)
Observations	351	351	351	351	351	351	351	351	351
AP F stat shape	6.09	6.09	6.09	26.33	26.33	26.33	33.92	33.92	33.92
AP F stat area	3.96	3.96	3.96	8.85	8.85	8.85	9.55	9.55	9.55
KP F stat	10.02	10.02	10.02	12.78	12.78	12.78	14.38	14.38	14.38

Notes: this table extends the robustness checks of Table 6, showing estimates of the first stage and of the impact of shape on population, controlling for initial shape (cols. 1 through 3), a British direct rule dummy (cols. 4 through 6), and a state capital dummy (cols. 7 through 9). Cols.1, 2, 4, 5, 7, and 8 report the first stage (analogous to cols. 1 and 2 in Table 2). Cols. 3, 6, and 9 report the IV impact of city shape on population (similar to Table 3, col.1). The dependent variables and regressors are all defined as long differences 2010-1950 (2011-1951 for population). Angrist-Pischke and Kleibergen-Paap F statistics are reported. Robust standard errors in parentheses. *** p<0.01,** p<0.05,* p<0.1.

Table A13: Robustness to sample cuts, non-predetermined characteristics

Excluded cities:	***A. Shrinking***			***B. Fast growing***			***C. Slow growing***			***D. Constrained***		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	First Stage		IV	First Stage		IV	First Stage		IV	First Stage		IV
Dependent variable:	*Δ Shape*	*Δ Log area*	*Δ Log population*	*Δ Shape*	*Δ Log area*	*Δ Log population*	*Δ Shape*	*Δ Log area*	*Δ Log population*	*Δ Shape*	*Δ Log area*	*Δ Log population*
Δ Potential shape, km	1.945***	0.238***		1.669***	0.232***		1.931***	0.225***		1.861***	0.239***	
	(0.250)	(0.0495)		(0.171)	(0.0542)		(0.250)	(0.0484)		(0.257)	(0.0553)	
Δ Log projected population	-2.255***	0.00689		-2.214***	-0.0382		-2.242***	0.0118		-2.214***	0.00894	
	(0.507)	(0.135)		(0.412)	(0.147)		(0.511)	(0.134)		(0.492)	(0.149)	
Δ Shape, km			-0.126**			-0.0728*			-0.124**			-0.125**
			(0.0510)			(0.0426)			(0.0503)			(0.0567)
Δ Log area			1.034***			0.578**			1.028***			0.969***
			(0.265)			(0.227)			(0.275)			(0.268)
Observations	336	336	336	316	316	316	316	316	316	316	316	316
AP F stat shape	21.25	21.25	21.25	25.12	25.12	25.12	21.12	21.12	21.12	22.52	22.52	22.52
AP F stat area	7.17	7.17	7.17	6.13	6.13	6.13	6.64	6.64	6.64	6.83	6.83	6.83
KP F stat	10.49	10.49	10.49	9.38	9.38	9.38	9.74	9.74	9.74	9.84	9.84	9.84

Notes: this table extends the robustness checks of Table A11, showing estimates of the first stage and of the impact of shape on population, excluding particular sets of cities, discussed in Section 7. Cols.1, 2, 4, 5, 7, 8, 10, and 11 report the first stage (analogous to cols. 1 and 2 in Table 2). Cols. 3, 6, 9, and 12 report the IV impact of city shape on population (similar to col. 1 in Table 3). The dependent variables and regressors are all defined as long differences 2010-1950 (2011-1951 for population). Angrist-Pischke and Kleibergen-Paap F statistics are reported. Robust standard errors in parentheses. *** p<0.01,** p<0.05,* p<0.1.

Table A14: Falsification test with lagged outcomes, wages and rents

Dependent variable	*Δ Log rents, 2008-2006*	*Δ Log rents, 2008-2006*	*Δ Log wages, 1995-1992*	*Δ Log wages, 1998-1994*	*Δ Log wages, 1995-1992*	*Δ Log wages, 1998-1994*
	(1)	(2)	(3)	(4)	(5)	(6)
Potential shape, km, 2000					-0.000151	
					(0.210)	
Potential shape, km, 2005					-0.0549	
					(0.217)	
Potential shape, km, 2007						-0.221*
						(0.125)
Potential shape, km, 2008	-0.0117					
	(0.166)					
Potential shape, km, 2010	-0.00202					0.120
	(0.132)					(0.0988)
Δ Potential shape, km, 2005-2000			0.00863			
			(0.196)			
Δ Potential shape, km, 2010-2005				-0.0465		
				(0.0339)		
Δ Potential shape, km, 2010-2009		0.0139				
		(0.0697)				
Projected population, 2000					-0.304	
					(0.691)	
Projected population, 2005					0.379	
					(0.669)	
Projected population, 2007						0.0337
						(0.910)
Projected population, 2008	3.771**					
	(1.709)					
Projected population, 2010	-3.731**					0.0464
	(1.701)					(0.914)
Δ Projected population, 2005-2000			0.290			
			(0.490)			
Δ Projected population, 2010-2005				0.500		
				(0.553)		
Δ Projected population, 2010-2009		-6.622*				
		(3.409)				
Observations	303	303	168	191	168	191

Notes: this table presents a falsification test similar to that of Table 7, to show that the instrument is not correlated with past changes in rents and wages. Robust standard errors in parentheses. *** p<0.01,** p<0.05,* p<0.1.

Table A15: Heterogeneous effects of transit, robustness

	Dependent variable: Δ Log population, 2011-1951							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Shape, km	-0.248**	-0.432**	-0.193**	-0.224**	-0.195**	-0.315**	-0.296**	-0.252**
	(0.115)	(0.219)	(0.0925)	(0.0886)	(0.0822)	(0.125)	(0.117)	(0.111)
Δ Shape · Roads	7.48e-06**	0.000170**	0.0108**	0.321***	0.407*	0.000356**	0.000775***	0.000190*
	(3.44e-06)	(8.07e-05)	(0.00474)	(0.121)	(0.227)	(0.000141)	(0.000296)	(0.000103)
Δ Log area	1.195***	1.517**	1.044***	1.050***	0.973***	1.480***	1.437***	1.338***
	(0.459)	(0.699)	(0.369)	(0.329)	(0.327)	(0.521)	(0.493)	(0.445)
Observations	336	336	336	123	123	246	246	246
AP F stat interaction	1079.77	35.7	358.48	37.59	31.18	644.28	693.9	328.63
AP F stat shape	11.04	5.04	11.6	5.61	6.60	6.98	7.62	7.38
AP F stat area	5.13	3.18	6.93	9.55	12.39	3.63	3.84	4.35
KP F stat	8.19	5.57	9.74	8.66	8.06	6.36	6.61	6.95
Interaction variable	Roads 2019	Roads1981	State roads 1981	Proximity	Grid roads	Cars 2011	Cars 2001	State cars 1984

Notes: this table reports the same specifications of Table 9, but additionally controls for the number of banks in 1981. Robust standard errors in parentheses.*** p<0.01,** p<0.05,* p<0.1.

Table A16: Impact of city shape on infrastructure and transit

Panel A: Roads

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS	IV	OLS
Dependent variable:	*Log roads*		*Log motorways*		*Log roads per capita*		*Log motorways per capita*		*Δ Log roads*		*Δ Log roads per capita*	
Shape, km	-0.161**	-0.0138	-0.707*	-0.490***	-0.0683	-0.0923***	-0.615*	-0.568***				
	(0.0728)	(0.0133)	(0.364)	(0.171)	(0.0507)	(0.0202)	(0.343)	(0.178)				
Log area	1.823***	1.215***	5.047***	4.416***	0.419**	0.639***	3.642**	3.840***				
	(0.237)	(0.0497)	(1.553)	(0.838)	(0.188)	(0.0745)	(1.503)	(0.856)				
Δ Shape, km									-0.0236	0.0815***	-0.0180	0.0659**
									(0.0529)	(0.0296)	(0.0536)	(0.0298)
Δ Log area									0.626	0.397***	0.377	0.277***
									(0.398)	(0.0917)	(0.456)	(0.102)
Observations	351	351	351	351	351	351	351	351	335	335	335	335
AP F stat shape	8.44		8.44		8.44		8.44		32.68		32.68	
AP F stat area	25.08		25.08		25.08		25.08		11.32		11.32	
KP F stat	6.73		6.73		6.73		6.73		15.69		15.69	

Panel B: Akbar et al. (2019)

	(1)	(2)	(3)	(4)	(5)	(6)
	IV	OLS	IV	OLS	IV	OLS
Dependent variable:	*Grid conformity*		*Mobility*		*Proximity*	
Shape, km	0.0149**	0.00519**	-0.0352*	-0.0165***	-0.0172	0.0158***
	(0.00639)	(0.00251)	(0.0202)	(0.00461)	(0.0210)	(0.00468)
Log area	-0.0940**	-0.0325**	0.130	0.0529**	0.174	-0.0310
	(0.0369)	(0.0131)	(0.0969)	(0.0238)	(0.117)	(0.0296)
Observations	128	128	128	128	128	128
AP F stat shape	1.78		1.78		1.78	
AP F stat area	4.38		4.38		4.38	
KP F stat	2.76		2.76		2.76	

Notes: this table reports estimates of the impact of city shape on infrastructure-related variables described in Section 9 and Section A.5 in the Appendix. In all cols. other than cols. 9 through 12, the regressors are city shape, in km, and log city area, measured in 2010. Panel A, cols. 1 through 8 considers the length of roads (motorways) in a city's 2010 lit-up shape, as reported in 2019 in Openstreetmap. In cols. 5 through 8 the dependent variable is the log of road length normalized by city population ins 2011. In cols. 9 through 12 the dependent variable is the log difference of 2019 Openstreetmap roads and 1981 city roads (from the Census) and the regressor s are 2010-1950 changes in city shape and log city area. Panel B considers indexes from Akbar et al. (2019), measured in 2016. Grid conformity (cols. 1 and 2) is a measure the regularity of a city's primary road grid. Mobility (col. 3 and 4) is a speed-based index of vehicular mobility. Proximity (cols. 5 and 6) is an index of distance accessibility from Akbar et al (2018). Means of the dependent variables are reported in Table 9. Estimation is by IV in odd columns, and OLS in even columns. Angrist-Pischke and Kleibergen-Paap F statistics are reported. Robust standard errors in parentheses.*** p<0.01,** p<0.05,* p<0.1.