Cities in Bad Shape: Urban Geometry in India*

Mariaflavia Harari†

October 2017

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 the geometric properties of cities over time using satellite imagery of night-time lights and historical maps. I then propose an instrument for urban shape that combines geography with a mechanical model for city expansion: cities are predicted to expand in circles of increasing sizes, and city shape is predicted by obstacles within each circle. With this instrument, I investigate how city shape affects consumer welfare and firm productivity, in a spatial equilibrium framework. Cities with more compact shapes have larger population, lower wages, and higher rents, consistent with compact shape being a consumption amenity. A one standard deviation deterioration in shape entails a 4% welfare loss for households but no productivity loss for firms. I also consider policy responses to deteriorating shape. The adverse effects of unfavorable topography are exacerbated by building height restrictions and mitigated by road infrastructure.

JEL: R10, R30

*I am grateful to 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, the Central European University and Stanford for helpful comments and discussions.

†The Wharton School, University of Pennsylvania. 1467 Steinberg Hall-Dietrich Hall, 3620 Locust Walk, Philadelphia, PA, 19104-6302. Phone: (215) 5733503. Email: harari@wharton.upenn.edu.

1 Introduction

In urban economics we often offer stylized representations of cities that are circular. Real-world cities, however, often depart significantly from this assumption. Geographic or regulatory constraints can result in asymmetric or fragmented urban development patterns.

While the economics literature has devoted very little attention to this particular feature of urban form, the spatial layout of cities is very much on the minds of town planners, who take decisions affecting the spatial structure of cities in the medium and long-term through master plans and land use regulations. Given the persistence of urban form, such policy decisions can have long-lasting effects on the location choices of consumers and firms and, more generally, on the spatial configuration of economic activity in a city.

In particular, the geometry of urban footprints is an important determinant of intra-urban commuting efficiency. All else equal, a city with a more compact geometry will be characterized by shorter potential within-city trips and more cost-effective transport networks, which, in turn, can affect the productivity of firms and the welfare of consumers (Bertaud, 2004; Cervero, 2001).

These considerations are especially relevant for cities in developing countries. Their rapid and often uncontrolled expansion can lead to sprawled and irregular spatial configurations. This can adversely impact urban connectivity, in a context in which most inhabitants cannot afford individual means of transportation. Cities in developing countries currently host 1.9 billion residents (around 74 percent of the world's urban population), and this figure is projected to rise to 4 billion by 2030 (UN, 2015).

This paper investigates empirically the causal economic impacts of urban geometry on consumers and firms in the context of India, exploiting plausibly exogenous variation in city shape driven by geographic barriers.

The first contribution of this study is to provide an economic framework to measure the welfare and productivity impacts of urban shape at the city level. The city lends itself as a natural unit of analysis, since urban planners typically have to consider the city as a whole when designing spatial planning policies. Drawing on a model of spatial equilibrium a la Roback-Rosen, I examine how consumers and firms are affected by urban geometry in their location choices across cities, and in particular, how much they value urban shapes conducive to shorter within-city trips. By examining the impact of city shape on population, wages, and housing rents, I quantify the loss from deteriorating urban geometry in a revealed preference setting.

A second contribution is related to the data and, in particular, the measurement of the spatial properties of urban footprints over time. I quantify the geometric properties of urban footprints using quantitative indicators of shape used in urban planning. Essentially, these indicators measure the extent to which the shape of a polygon departs from that of a circle, higher values indicating a less compact urban footprint and longer within-city distances. To illustrate that more compact cities are associated with shorter distances, consider as an example the cities of Kolkata and Bangalore. Kolkata has a distinctive elongated layout, stretching along the North-South axis, whereas Bangalore, roughly

shaped like a pentagon, has a more compact layout. Controlling for city area, I find that the average linear distance between any two points in the city is over 40 percent longer in Kolkata than it is in Bangalore.[1]

Given the persistence of urban form and the long-term nature of spatial planning decisions, investigating questions related to city shape requires observing the dynamics of urban footprints for a large enough sample of cities and over a sufficiently long time horizon, drawing on retrospective data. With its numerous cities in rapid expansion and the world's second largest urban population (UN, 2015), India represents a relevant setting for researching these issues. However, systematic data on Indian cities and their spatial structures is not readily available. I assemble a novel panel dataset that covers over 450 Indian cities and includes detailed information on each city's spatial properties and micro-geography, as well as economic outcomes from the Census and other sources. In particular, I trace the evolution of the footprints of Indian cities over the past decades by combining newly geo-referenced historical maps (1951) with satellite imagery of night-time lights (1992-2010). One of the patterns emerging from these data is that rapidly growing cities have a tendency to become less compact in shape over time.

A third contribution of the paper concerns the identification strategy. Estimating the causal impact of city shape on economic outcomes is challenging, given that 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 geography, city growth, and policy choices, such as land use regulations and infrastructural investment. In order to overcome this endogeneity problem, I propose a novel instrument for urban geometry that combines geography with a mechanical model for city expansion. The underlying idea is that, as cities expand in space over time, they face different geographic constraints - steep terrain or water bodies (Saiz, 2010) - 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; in its simplest version, such model postulates a common growth rate for all cities. 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 compute the geometric properties of the "potential footprint" and I then instrument the geometric properties of the *actual* city footprint in that given year with the shape properties of the *potential* footprint. The resulting instrument varies at the city-year level, allowing me to control for time-invariant city characteristics through city fixed effects. The identification of the impact of shape thus relies on changes in shape that a given city undergoes over time, as a result of hitting geographic obstacles. This instrument's explanatory power is not limited to extremely constrained topographies (e.g., coastal or mountainous cities) in my sample.

With this instrument in hand, I examine the aggregate responses of population, wages, and housing rents, measured at the city-year level, to changes in shape. Through the lens of the model, these

[1] See Section 3.2 and Figure 3 for a more detailed discussion of this example.

responses allow me to establish whether households and firms value compact city shape when making location choices across cities, and to quantify this value. I document that city shape, a feature of urban form previously overlooked in the economics literature, can have substantial economic implications, particularly for consumers.

First, my findings indicate that consumers value city compactness as a "consumption amenity" that enhances their indirect utility. All else equal, more compact cities experience faster population growth. There is also evidence that consumers are paying a premium for living in more compact cities, in terms of lower wages and higher housing rents. Households locating in non-compact cities require a substantial compensation: a one-standard deviation deterioration in city shape, corresponding to a 720 meter increase in the average within-city round-trip distance, entails a loss equivalent to a 4percent decrease in income.

Second, I find that city shape does not affect the productivity of firms in equilibrium. Thus, compact city shape can be likened to a pure "consumption amenity", but not to a "production amenity". This does not automatically indicate that city compactness is *ex ante* irrelevant for firms. Rather, my results indicate that, in equilibrium, firms are able to optimize against "bad" shape, in a way that consumers cannot.

The margin through which firms are able to neutralize the effects of bad shape could be their location choices within cities. Drawing evidence on the street addresses of establishments, I show that firms located in non-compact cities tend to cluster in few employment sub-centers. It is then consumers who have to bear the costs of longer commutes to work, and who require a compensation for these longer trips through wages and rents.

Third, I consider the role of policy. I investigate some of the possible policy responses to a "bad" city shape, and ways in which planners can influence city shape. On the one hand, I consider infrastructure. I find that the negative effects of deteriorating shape on population are mitigated by road infrastructure, suggesting that infrastructural investment could be a policy tool to counteract the effects of poor geometry. This also supports the interpretation that intra-urban commuting is the primary channel through which non-compact shape affects consumers.

On the other hand, I consider land use regulations as co-determinants of city shape. I find that more permissive vertical limits for buildings, in the form of higher Floor Area Ratios (FARs), result in cities that are less spread out in space and more compact than their topographies would predict. Holding geography constant, increasing FARs by one[2] in an average-sized city improves shape by an amount that is associated with a 4 percent welfare gain.

Taken together, these results provide policy lessons for developing countries that prepare to accommodate a large urban expansion, and in which policy makers are concerned about haphazard urban growth. In India, in particular, there is growing concern about existing land use regulations and urban planning practices that are viewed as potentially distortive of urban form and conducive to sprawl.[3]

[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.

[3]Beyond FARs, examples include the Urban Land Ceiling Act, claimed to hinder intra-urban land consolidation, and regulations preventing the conversion of land from one use to another.

As I document, rapid urban growth can be accompanied by a deterioration in urban shape, which has a negative impact on the welfare of consumers. The latter should be incorporated in cost-benefit analyses when designing spatial planning policies.

The rest of the paper is organized as follows. Section 2 provides some background on urbanization in India and reviews related literature. Section 3 describes the main data sources and the shape indicators I employ. Section 4 outlines the conceptual framework. Section 5 presents my empirical strategy and instrument construction. In Section 6 I discuss my main empirical results, concerning the implications of city shape for the spatial equilibrium across cities. Section 7 presents additional empirical results on responses to city shape, including interactions with policy. Section 8 concludes.

2 Background and Previous Literature

India represents a promising setting to study urban spatial structures for several reasons. First, as most developing countries, India is experiencing fast urban growth. According to the 2011 Census, the urban population amounts to 377 million, representing 31 percent of the total population, and it is predicted that another 250 million will join the urban ranks by 2030 (McKinsey, 2010). This growth in population has been accompanied by a significant physical expansion of urban footprints, typically beyond urban administrative boundaries (Indian Institute of Human Settlements, 2013; World Bank, 2013). This setting thus provides a unique opportunity to observe the shapes of cities as they evolve over time.

Secondly, unlike most other developing countries, India has a large number of growing cities. This provides enough power for an econometric approach based on a city-year panel.

The challenges posed by rapid urban expansion on urban form have gained increasing importance in India's policy debate, making it particularly relevant to investigate these matters from an economics perspective. Sprawl, lengthy commutes and limited urban mobility are often cited among the perceived harms of rapid urbanization (e.g., World Bank, 2013). There is also a growing concern that existing land use regulations might contribute to distorting urban form (World Bank, 2013, Sridhar, 2010, Glaeser, 2011). In particular, sprawl has been linked to vertical limits in the form of restrictive Floor Area Ratios (Bertaud, 2002; Bertaud and Brueckner, 2005; Brueckner and Sridhar, 2012; Glaeser, 2011; Sridhar, 2010; World Bank, 2013). Another example is given by 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).

Literature directly related to the geometric layout of cities is scant, but a number of literature 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, typically 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. A large body of empirical literature investigates urban sprawl (see Glaeser and Kahn, 2004), typically in the US context, suggesting longer commutes as one of its potential costs (Bajari and Kahn, 2004). Although some studies identify sprawl with non-contiguous development (Burchfield et al., 2006), 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 density as an outcome variable.[4]

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., 2009a, 2009b). 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). 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.

In terms of methodology, my work is related to that of Burchfield et al. (2006), who also 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. Saiz (2010) also examines geographic constraints to city expansion, by computing the amount of developable land within 50 km radii from US city centers and relating it to the elasticity of housing supply. I use the same notion of geographic constraints, but I employ them in a novel way to construct a time-varying instrument for city shape.

This paper also complements a growing literature on infrastructure, transit and urban expansion in developing countries. Most of this literature considers the impact of roads connecting different cities in a country (Baum-Snow et al., 2017; Morten and Oliveira, 2016; Storeygard, 2016). Differently from these studies, I consider the internal structure of cities rather than infrastructural connections across cities. At a more micro level, Akbar et al. (2017) examine transit times within Indian cities using current traffic data. Differently from their paper, my focus is on the spatial configuration of cities and on their long-term dynamics. My unit of observation of choice is the city as a whole - the level at which the decisions of urban planners are typically taken.

3 Data Sources

I assemble a panel dataset covering over 450 Indian cities, for a period ranging from 1951 to 2011. 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 more detailed description of data sources and methods can be found in the Appendix (Section A). Summary statistics are provided in Table 1.

3.1 Urban Footprints

The first step in constructing my dataset is to trace the footprints of Indian cities at different points in time and measure their geometric properties. I retrieve the boundaries of urban footprints

[4] In this respect, my work is also related to that of Bento et al. (2005), who incorporate a measure of city shape in their investigation of the link between urban form and travel demand in a cross-section of US cities. Differently from their approach, I rely on a panel of cities and I address the endogeneity of city shape in an instrumental variables framework.

from two sources. First, a set of historical maps of India (U.S. Army Map Service, ca. 1950), that I georeferenced and used to trace the boundaries of urban areas as of 1951. A fragment of one such map, with the detail of the city of Mumbai, is displayed in Figure 1.

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 km. While night-time lights have been used by economists mostly for purposes other than urban mapping (Henderson et al., 2012), they can be used to delineate urban areas by considering spatially contiguous lighted areas surrounding a city's coordinates, with luminosity above a pre-defined threshold. This approach is illustrated in Figure 2.[5] Although this process is not immune to measurement error, it should be noted that all the geometric properties of urban footprints, including both area and shape, will be instrumented throughout my analysis. Amongst other things, this instrumental variables approach addresses issues of non-classical measurement error in the extents of urban footprints - for instance, due to the correlation between income and luminosity.

3.2 Shape Metrics

For each footprint-year, I compute indicators of city shape (Angel et al., 2009a, 2009b) used in urban planning to proxy for the length of within-city trips. Such indicators are all based on the distribution of points within a polygon. Conditional on footprint area, higher values of these indexes (measured in kilometers) indicate longer within-city distances and less compact footprints.

(i) The *remoteness* index is the average distance between all interior points in a polygon and the polygon's centroid. It can be considered a proxy for the average length of commutes to the urban center.

(ii) The *spin* index is the average of the squared distances between interior points and the centroid. Relative to the remoteness index, it gives more weight to the polygon's extremities.

(iii) The *disconnection* index is the average distance between all pairs of interior points. It can be viewed as a proxy for commutes within the city, without focusing on those to or from the center. As I discuss in Section 6 below, this will be my benchmark indicator.

(iv) The *range* index is the maximum distance between two points on the polygon's perimeter, representing the longest possible commute trip within a city.

It should be noted that all of the above indexes are 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. When I follow this approach, city area is separately instrumented for (see Section 5.2). Alternatively, it is possible to normalize each index, computing a version that is invariant to the area of the polygon, and obtaining what I define *normalized remoteness*, *normalized spin,* etc.[6] One way to interpret these normalized metrics is as deviations of a polygon's shape from that of a circle, the shape that minimizes all the indexes above.

[5] The choice of luminosity threshold is discussed in detail in the Appendix, Section A. My results are robust to alternative luminosity thresholds (available upon request).

[6] See the Appendix, Section A, for details on the normalization.

In order to provide a visual example of how these metrics map into urban shape, Figure 3 displays the footprints of Bengaluru and Kolkata as of year 2005, where Bengaluru's footprint has been rescaled so that they have the same area, along with the corresponding shape metrics. Among India's largest cities, they are the ones with, respectively, the "best" and the "worst" geometry based on the above indicators. The difference in the disconnection index between Kolkata and (rescaled) Bengaluru is 6.2 km. All else being equal, if Kolkata had the same compact shape as Bengaluru, the average potential trip within the city would be shorter by 6.2 km.[7]

Importantly, these metrics should be viewed as proxies for the length of *potential* intra-urban trips as driven by the city's layout, abstracting from the actual distribution of households or jobs within the city. Commuting trips that are realized in equilibrium can be thought of as subsets of those potential trips, that depend on household location choices within each city. Detailed data at the sub-city level is, in general, difficult to obtain for India. To have a rough sense of the mapping between city shape and realized commuting, I draw on 2011 Census data on distances from residence to place to work, available at the district level.[8] For year 2011, the correlation between my benchmark measure of shape - the disconnection index - and the population-weighted average distance to work in the corresponding district is 0.208 (p-value 0.001). As expected, this correlation is positive.

3.3 Outcome Data: Population, Wages, Rents

Population data at the city level for the period 1871-2011 is obtained from the Census, available at 10-year intervals.[9]

Data on wages are taken from the Annual Survey of Industries (ASI), waves 1990, 1994, 1995, 1997, 1998, 2009, and 2010. These are a series of repeated cross-sections covering manufacturing plants in the formal sector.[10] This selective coverage may affect the interpretation of my results, to the extent that this sector is systematically over- or underrepresented in cities with worse shapes. I provide some suggestive evidence on the relationship between city shape and the local industry mix in Appendix Table A4. 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. I construct a rough proxy for the rental price of housing drawing on the National Sample Survey (rounds 2005-

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

[8] 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 km, 2-5 km, 6-10 km, 11-20 km, 21-30 km, 31-50 km, or above 51 km). As discussed in Section 3.3, the matching between cities and districts in India is not one to one, but this correlation is robust to different approaches for matching cities to districts.

[9] Note that "footprints", as retrieved from the night-time lights dataset, do not always have an immediate Census counterpart. The matching between footprints and population totals requires intermediate steps, detailed in the Appendix.

[10] A potential alternative source of wages data covering also the informal sector is the National Sample Survey. However, these data are more problematic to match to cities because few waves are representative at a sufficiently fine geographic level (see Appendix Section A).

2006, 2006-2007, and 2007-2008), in which households are asked about the amount spent on rent and about the floor area of their dwelling.[11] The average yearly rent per square meter is 603 Rs., at current prices. 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). 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, which is *a priori* less likely to include observations from rent-controlled housing.

It should be noted that, for both wages and rents, I rely on sources that are not at the city, but at the coarser district level. Following Greenstone and Hanna (2014), I use district urban averages as proxies for city-level averages. As discussed in the Appendix, Section A, 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).

3.4 Other Data

For the purposes of constructing the instrument, I code geographic constraints 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 percent. Data on water bodies and slope is drawn from high-resolution raster datasets (details are provided in the Appendix, Section A). Figure 4 illustrates this classification for the Mumbai area.

Data on the location of firms within cities in 2005 is drawn from the 5th Economic Census, which collects street addresses and employment class of all urban productive establishments. I geo-referenced addresses corresponding to cities in my sample through Google Maps, retrieving consistent coordinates for approximately 240 thousand establishments in about 190 cities.

Data on Floor Area Ratios are drawn from Sridhar (2010), who collects a cross-section of the maximum allowed FAR levels as of the mid-2000s, for 55 cities in my sample,[12] disaggregating by residential and non-residential FAR.

As this data description shows, retrieving and assembling together city-year level data for Indian cities is not a straightforward exercise. Moreover, it should be noted that data at a more disaggregated level than the city is typically not available for India. In particular, I cannot observe population densities and commuting patterns within cities.[13]

4 Conceptual Framework

I frame the empirical question of the economic value of city shape drawing on a simple model of spatial equilibrium across cities, with production and consumption amenities (Rosen 1979; Roback, 1982). In this framework, consumers and firms optimally choose in which city to locate, and, in equilibrium, they are indifferent across cities with different amenities. I hypothesize that households and firms

[11] For those who own, an imputed figure is provided. Results are very similar when excluding owners from the sample (available upon request).

[12] 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.

[13] I cannot infer within-city density patterns through the night-time lights dataset, which does not display enough variation in luminosity within urban areas.

may value the "compactness" of a city as one of these amenities, as they incorporate considerations on the relative ease of within-city trips when evaluating the trade-offs associated with different cities. I then use this framework to guide my empirical analysis, the goal of which is to establish whether compact city shape is an amenity for consumers and/or for firms. In the model, wages and housing rents equalize the equilibrium utility across cities, and are informative of the value that agents attribute to different amenities. This modeling approach is attractive because it allows me to shed light on the economic value of city shape for consumers and for firms, by observing the aggregate responses of population, wages, and housing rents, measured at the city-year level, to changes in urban shape.[14]

I follow the exposition of the model by Glaeser (2008). Households consume a composite good C and housing H. They supply labor inelastically, receiving a city-specific wage W. Their utility depends on labor income net of housing costs, and on a city-specific bundle of consumption amenities θ.[15] City compactness may be a part of this bundle. To fix ideas, one can imagine cities with more compact shapes to have a better functioning transportation network, which in turn reduces transportation costs for consumers in the city. However, the model *per se* is agnostic about the specific channels by which city shape enters the utility of consumers.[16] For a given city of residence, their optimization problem reads:

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

where p_h is the rental price of housing, and

$$U(C,H,\theta) = \theta C^{1-\alpha} H^{\alpha}. \tag{2}$$

In equilibrium, indirect utility V must be equalized across cities, otherwise workers would move:[17]

$$V(W - p_h H, H, \theta) = \overline{v} \tag{3}$$

which, given the functional form assumptions, yields the condition:

$$\log(W) - \alpha \log(p_h) + \log(\theta) = \log(\overline{v}). \tag{4}$$

The intuition for (4) is that consumers, in equilibrium, pay for amenities θ through lower wages (W)

[14]Given that the unit of analysis of choice is the city, I employ a model of spatial equilibrium across, as opposed to within cities. A question may however arise on how irregular city shapes could be incorporated in a framework of spatial equilibrium within cities (à la Alonso-Mills-Muth), and whether my empirical results are compatible with the predictions of such an alternative model. I refer the interested reader to the Appendix (Section B), in which I outline a simple model of a monocentric city with a topographic constraint and I show that the model's predictions are consistent with my empirical findings.

[15]In general, amenities may be exogenous or endogenous. For simplicity, in the model I treat amenities as exogenous parameters. In the empirical analysis, however, I discuss specific reasons for the possible endogeneity of city shape and consider a plausibly exogenous shifter of city shape as my instrument.

[16]Empirically, pinning down these channels directly would require more disaggregated data than what is available for India. However, some evidence on mechanisms can be inferred from heterogeneous effects, explored in Section 6.5 below. In particular, I provide empirical evidence suggesting that the urban transit channel is involved. Alternative channels through which city shape might affect consumers are also briefly discussed in Section 8.

[17]The notion of spatial equilibrium across cities presumes that consumers 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 percent 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.

or through higher housing prices (p_h).[18] The extent to which housing rents net of wages rise with an amenity is a measure of the extent to which that amenity increases utility, relative to the marginal utility of income. Holding indirect utility $\overline{v}$ constant, differentiating this expression with respect to some exogenous variable S - which could be (instrumented) city geometry - yields:

$$\frac{\partial \log(\theta)}{\partial S} = \alpha \frac{\partial \log(p_h)}{\partial S} - \frac{\partial \log(W)}{\partial S}. \tag{5}$$

This equation provides a way to evaluate the amenity value of S: the overall impact of S on utility can be found as the difference between the impact of S on housing prices, multiplied by the share of housing in consumption α, and the impact of S on wages.

Firms in the production sector also choose optimally in which city to locate. Each city is a competitive economy that produces a good Y, using labor N and a bundle of local production amenities A, that may include city compactness. The technology of firms also requires traded capital K and a fixed supply of non-traded capital $\overline{Z}$.[19] Firms solve the following profit maximization problem:

$$\max_{N,K} \left\{ Y(N,K,\overline{Z},A) - WN - K \right\} \tag{6}$$

where

$$Y(N,K,\overline{Z},A) = AN^{\beta}K^{\gamma}\overline{Z}^{1-\beta-\gamma}. \tag{7}$$

In equilibrium, firms earn zero expected profits. Under these functional form assumptions, the maximization problem for firms yields the following labor demand condition:

$$(1-\gamma)\log(W) = (1-\beta-\gamma)(\log(\overline{Z}) - \log(N)) + \log(A) + \kappa_1. \tag{8}$$

Finally, developers produce housing H, using land l and "building height" h. In each city there is a fixed supply of land $\overline{L}$, as a result of land use regulations.[20] Denoting the price of land with p_l, their maximization problem reads:

$$\max_{H} \{ p_h H - C(H) \} \tag{9}$$

where

$$H = l \cdot h \tag{10}$$

$$C(H) = c_0 h^{\delta} l - p_l l \ , \ \delta > 1. \tag{11}$$

Construction profits are equalized across cities. By combining the housing supply equation, resulting

[18] This simple model assumes perfect mobility across cities. With migration costs, agents other than the marginal migrant will not be indifferent across locations and will not be fully compensated for disamenities. This would lead to larger gaps in wages net of housing costs than if labor were perfectly mobile. These gaps would reflect differences in amenities as well as migration costs or idiosyncratic preferences for particular locations (Kline and Moretti, 2014).

[19] Introducing a fixed supply of non-traded capital is one way to ensure decreasing returns at the city level, which, in turn, is required in order to have a finite city size. This fixed capital assumption could be dropped by assuming, for instance, decreasing returns in the production of housing (Glaeser, 2008).

[20] 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, when city area is explicitly controlled for, it is instrumented using historical population, thus abstracting from these issues (see Section 5.2).

from the developers' maximization problem, with the housing demand equation, resulting from the consumers' problem, one obtains the following housing market equilibrium condition:

$$(\delta-1)\log(H)=\log(p_h)-\log(c_0\delta)-(\delta-1)\log(N)+(\delta-1)\log(\overline{L}). \tag{12}$$

Using the three optimality conditions for consumers (4), firms (8), and developers (12), the model can be solved for the three endogenous variables N, W, and p_h, representing, respectively, population, wages, and housing prices, as functions of the city-specific productivity parameter and consumption amenities. Denoting all constants with K, this yields the following:

$$\log(N)=\frac{(\delta(1-\alpha)+\alpha)\log(A)+(1-\gamma)\left(\delta\log(\theta)+\alpha(\delta-1)\log(\overline{L})\right)}{\delta(1-\beta-\gamma)+\alpha\beta(\delta-1)}+K_N \tag{13}$$

$$\log(W)=\frac{(\delta-1)\alpha\log(A)-(1-\beta-\gamma)\left(\delta\log(\theta)+\alpha(\delta-1)\log(\overline{L})\right)}{\delta(1-\beta-\gamma)+\alpha\beta(\delta-1)}+K_W \tag{14}$$

$$\log(p_h)=\frac{(\delta-1)\left(\log(A)+\beta\log(\theta)-(1-\beta-\gamma)\log(\overline{L})\right)}{\delta(1-\beta-\gamma)+\alpha\beta(\delta-1)}+K_P. \tag{15}$$

These conditions translate into the following predictions:

$$\frac{d\log(N)}{d\log(A)}>0,\ \frac{d\log(N)}{d\log(\theta)}>0 \tag{16}$$

$$\frac{d\log(W)}{d\log(A)}>0,\ \frac{d\log(W)}{d\log(\theta)}<0 \tag{17}$$

$$\frac{d\log(p_h)}{d\log(A)}>0,\ \frac{d\log(p_h)}{d\log(\theta)}>0. \tag{18}$$

Population, wages, and rents are all increasing functions of the city-specific productivity parameter A. Population and rents are increasing in the amenity parameter θ as well, whereas wages are decreasing in it. The intuition is that firms prefer cities with higher production amenities, whereas consumers prefer cities with higher consumption amenities. Factor prices - W and p_h - strike the balance between these conflicting location preferences.

Consider now an exogenous shifter of urban geometry S, higher values of S denoting less compact shapes. Suppose that non-compact shape is purely a consumption disamenity, which decreases consumers' utility, all else being equal, but does not directly affect firms' productivity:

$$\frac{\partial\theta}{\partial S}<0,\ \frac{\partial A}{\partial S}=0. \tag{19}$$

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 one should observe the following:

$$\frac{dN}{dS}<0,\frac{dW}{dS}>0,\frac{dp_h}{dS}<0. \tag{20}$$

A city with poorer shape should have, *ceteris paribus,* a smaller population, higher wages, and lower housing rents. Intuitively, consumers prefer cities with good shapes, which drives rents up and bids wages down in these locations.

Suppose, instead, that poor city geometry is both a consumption and a production disamenity, i.e., it depresses both the utility of consumers and the productivity of firms:

$$\frac{\partial \theta}{\partial S} < 0, \ \frac{\partial A}{\partial S} < 0. \tag{21}$$

This would be the case if the costs of longer commutes were borne by households but also by firms. This would imply the following:

$$\frac{dN}{dS} < 0, \frac{dW}{dS} \gtrless 0, \frac{dp_h}{dS} < 0. \tag{22}$$

The model's predictions are similar, except that the effect on wages will be ambiguous, given that now both firms and consumers prefer to locate in compact cities. With respect to the previous case, there is an additional force that tends to bid wages up in compact cities: competition among firms for locating in those cities. The net effect on W depends on whether firms or consumers value low S relatively more (on the margin). If S is more a consumption than it is a production disamenity, then $\frac{dW}{dS} > 0$.

To strengthen the exposition of this point, assume now that:

$$\log(A) = \kappa_A + \lambda_A S \tag{23}$$

$$\log(\theta) = \kappa_\theta + \lambda_\theta S. \tag{24}$$

Parameter λ_θ captures, in log points, the utility loss from a marginal increase in S. Parameter λ_A captures the impact of a marginal increase in S on city-specific productivity. Plugging (23) and (24) into (13), (14), and (15) yields:

$$\log(N) = B_N S + D_N \log(\overline{L}) + K_N \tag{25}$$

$$\log(W) = B_W S + D_W \log(\overline{L}) + K_W \tag{26}$$

$$\log(p_h) = B_P S + D_P \log(\overline{L}) + K_P \tag{27}$$

where

$$B_N := \frac{(\delta(1-\alpha)+\alpha)\lambda_A + (1-\gamma)\delta\lambda_\theta}{\delta(1-\beta-\gamma)+\alpha\beta(\delta-1)}, \ D_N := \frac{(1-\gamma)\alpha(\delta-1)}{\delta(1-\beta-\gamma)+\alpha\beta(\delta-1)} \tag{28}$$

$$B_W := \frac{(\delta-1)\alpha\lambda_A - (1-\beta-\gamma)\delta\lambda_\theta}{\delta(1-\beta-\gamma)+\alpha\beta(\delta-1)}, \ D_W := \frac{-(1-\beta-\gamma)\alpha(\delta-1)}{\delta(1-\beta-\gamma)+\alpha\beta(\delta-1)} \tag{29}$$

$$B_P := \frac{(\delta-1)\lambda_A + (\delta-1)\beta\lambda_\theta}{\delta(1-\beta-\gamma)+\alpha\beta(\delta-1)}, \ D_P := \frac{-(1-\beta-\gamma)(\delta-1)}{\delta(1-\beta-\gamma)+\alpha\beta(\delta-1)}. \tag{30}$$

Note that $(28), (29), (30)$ imply:

$$\lambda_A = (1-\beta-\gamma)B_N + (1-\gamma)B_W \tag{31}$$

$$\lambda_\theta = \alpha B_P - B_W. \tag{32}$$

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 $\beta, \gamma,$ and α, can be used to back out λ_A and λ_θ:

$$\widehat{\lambda_A} = (1-\beta-\gamma)\widehat{B_N} + (1-\gamma)\widehat{B_W} \quad (33)$$

$$\widehat{\lambda_\theta} = \alpha\widehat{B_P} - \widehat{B_W}. \quad (34)$$

Through this approach one can capture the overall, net effect of S, in equilibrium, on the marginal city dweller.

This model makes a number of simplifying assumptions. First, it does not address heterogeneity across consumers in tastes and skills, by which agents will sort themselves into locations based on their preferences. The estimated differences in wages and rents across cities should be thus thought of as an underestimate of true equalizing differences for those with a strong taste for the amenity of interest, in this case compact layouts, and an overestimate for those with weak preferences.[21]

Second, the model could be extended to allow for congestion or agglomeration. In particular, in the presence of congestion externalities in consumption, the indirect utility of consumers will negatively depend on city size N as well. If shape is a consumption amenity and more compact cities have larger populations, they will also be more congested; this congestion effect will tend to reduce the positive impact of compact shape on utility. The implication is that, when estimating the consumption amenity value of compact shape through equation (34), I would be capturing the equilibrium effect of shape, gross of congestion; hence, $\widehat{\lambda_\theta}$ will be a *lower* bound for λ_θ.

Similarly, in the presence of agglomeration externalities 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 production amenity value of compact shape, obtained from equation (33), will be an *upper* bound for λ_A.[22]

Reduced-form estimates for B_N, B_W, B_P are presented in Sections 6.2 and 6.3. Section 6.4 provides estimates for parameters $\lambda_A, \lambda_\theta$.

5 Empirical Strategy

The model suggests that the city-level responses of population, wages, and rents to city shape are informative of whether consumers and firms value city compactness as a production or consumption amenity. In the next section, I examine these responses empirically, by estimating empirical counterparts of equations (25)-(27) for a panel of city-years. Denote the city with c and the year with t; let $area_{c,t}$ be the area of the urban footprint, and recall that S is an indicator for city shape. The specification of interest is:

$$\log(Y_{c,t}) = a \cdot S_{c,t} + b \cdot \log(area_{c,t}) + \eta_{c,t} \quad (35)$$

[21] While a richer model would allow to capture these nuances, the scope of my empirical analysis is limited by the lack of disaggregated data. Some indirect evidence of sorting is discussed in Section 6.5, in which I examine slum populations across cities with different geometries.

[22] My identification strategy does not allow me to pin down the pure amenity value of compact shape, net of congestion / agglomeration, as I would require an additional source of exogenous variation in city size.

where the outcome variable $Y \in (N, W, p_H)$.

The main concern in estimating the relationship between city shape $S_{c,t}$ and city-level outcomes $Y_{c,t}$ 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 local geography, city growth, and policy choices. City shape is affected by urban planning both directly, through master plans, and indirectly - for instance, land use regulations can encourage land consolidation, resulting in a more compact development pattern. Moreover, investments in road infrastructure can encourage urban growth along transport corridors, generating distinctive spatial patterns of development. Such policy choices are likely to be jointly determined with the outcome variables at hand. To see how this could bias estimates of parameter a, consider the response of population to city shape. Faster growing cities could be subject to more stringent urban planning practices, due to a perceived need to prevent haphazard growth, which, in turn, may result in more compact urban shapes. This would create a spurious *positive* correlation between compactness and population, and would bias my estimates towards finding a positive response. On the other hand, faster growing cities may be expanding in a more chaotic and unplanned fashion, generating a "leapfrog" pattern of development, which translates into less compact shapes. This would create a spurious *negative* correlation between compactness and population.

Another concern is that compact shape could be systematically correlated with other amenities or disamenities. For example, there may be some unobserved factor - e.g., better institutions and law enforcement - that causes cities to have both better quality of life and better urban planning practices, which result in more compact shapes. In this case, I may observe a response of population, wages and rents compatible with compact shape being a consumption amenity even if shape were not an amenity *per se*. For the reasons discussed above, a naïve estimation of (35) would suffer simultaneity bias in a direction that is *a priori* ambiguous.

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 the city's growth. My benchmark specifications include city and year fixed effects, that account for time-invariant city characteristics and for India-wide trends in population and other outcomes.

Details of the instrument construction and estimating equations are provided in Sections 5.1 and 5.2, respectively. Section 5.3 discusses in more depth the identification strategy and possible threats to identification.

5.1 Instrumental Variable Construction

The instrument is constructed by combining variation in topographic constraints with a mechanical model for city expansion in time. This is based on the idea that, as cities expand in space and over time,

they hit different geographic obstacles that constrain their shapes by preventing expansion in some of the possible directions. 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. More 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 properties of the *potential* footprint and use these as instruments for the corresponding shape properties 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. In its simplest form, this mechanical model postulates a common growth rate for all cities.

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 1951 (from the U.S. Army maps) and then in every year t between 1992 and 2010 (from the night-time lights dataset). I take as a starting point the minimum bounding circle of the 1951 city footprint (Figure 5a). To construct the instrument for city shape in 1951, 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 1951. 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 simple, mechanical model for city expansion in space. I consider two versions of this model: a "city-specific" one and a "common rate" one.

City-specific: I postulate that the rate of expansion of $\widehat{r}_{c,t}$ varies across cities, depending on their 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, 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 1951-2010 panel of cities and run the following regression:

$$\log(area_{c,t}) = \alpha \cdot \log(\widehat{pop_{c,t}}) + \beta \cdot \log\left(\frac{pop_{c,1951}}{area_{c,1951}}\right) + \gamma_t + \varepsilon_{c,t}. \tag{36}$$

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

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

$$\widehat{r_{c,t}} = \sqrt{\frac{\widehat{area_{c,t}}}{\pi}}. \tag{37}$$

The interpretation of the circle with radius $\widehat{r_{c,t}}$ from Figures 5c and 5d is thus the following: this is the area the city would occupy if it continued to grow as in 1871-1951, if its density remained the

same as in 1951, and if the city expanded freely and symmetrically in all directions, in a fashion that optimizes the length of within-city trips.

Common rate: In this alternative, simpler version of the model, the rate of expansion of $\widehat{r}_{c,t}$ is the same for all cities, and equivalent to the average expansion rate across all cities in the sample. 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:

$$\log(area_{c,t}) = \theta_c + \gamma_t + \varepsilon_{c,t} \tag{38}$$

where θ_c and γ_t denote city and year fixed effects. From the regression above, I obtain an alternative version of $\widehat{area_{c,t}}$, and corresponding $\widehat{r_{c,t}} = \sqrt{\frac{\widehat{area_{c,t}}}{\pi}}$.

The richer "city-specific" model yields an instrument that has better predictive power in the first stage. The "common rate" model yields a weaker first stage, but provides arguably a cleaner identification as it does not rely on historical projected population, a variable that may be correlated with present-day outcomes.

5.2 Estimating Equations

Consider a generic shape metric S. Denote with $S_{c,t}$ the shape metric computed for the *actual* footprint observed for city c in year t, and with $\widetilde{S_{c,t}}$ the shape metric computed for the *potential* footprint of city c in year t, namely the largest contiguous patch of developable land within the predicted radius $\widehat{r_{c,t}}$.

Double-Instrument Specification

The first set of specifications that I consider are empirical counterparts of equations (25)-(27). As I clarify below, in this specification there are two variables that I treat as endogenous and that I instrument for: city shape and area. For this reason, throughout the paper I refer to this approach as to the "double-instrument" specification. Consider outcome variable $Y \in (N, W, p_H)$ and let $area_{c,t}$ be the area of the urban footprint. The counterparts of equations (25)-(27), augmented with city and year fixed effects, take the following form:

$$\log(Y_{c,t}) = a \cdot S_{c,t} + b \cdot \log(area_{c,t}) + \mu_c + \rho_t + \eta_{c,t}. \tag{39}$$

The two endogenous regressors in equation (39) are $S_{c,t}$ and $\log(area_{c,t})$. These are instrumented using respectively $\widetilde{S_{c,t}}$ and $\log(\widehat{pop_{c,t}})$ - the same projected historical population used in the city-specific model for urban expansion, step (i), described above. This results in the following two first-stage equations:

$$S_{c,t} = \sigma \cdot \widetilde{S_{c,t}} + \delta \cdot \log(\widehat{pop_{c,t}}) + \omega_c + \varphi_t + \theta_{c,t} \tag{40}$$

and

$$\log(area_{c,t}) = \alpha \cdot \widetilde{S_{c,t}} + \beta \cdot \log(\widehat{pop_{c,t}}) + \lambda_c + \gamma_t + \varepsilon_{c,t}. \tag{41}$$

The counterpart of $\log(area_{c,t})$ in the conceptual framework is $log(\overline{L})$, where $\overline{L}$ is the amount of land

which regulators allow to be developed in each period. It is plausible that regulators set this amount based on projections of past city growth, which rationalizes the use of projected historical population as an instrument.

One advantage of this approach is that it allows me to analyze the effects of shape and area considered separately - recall that the non-normalized shape metrics are mechanically correlated with footprint size. However, a drawback is that it requires not only an instrument for shape, but also one for area. Moreover, there is a concern that historical population might be correlated with current outcomes, leading to possible violations of the exclusion restriction. This motivates me to complement this specification with an alternative one, that does not explicitly include city area in the regression, and therefore does not require including projected historical population among the instruments. I denote this as the "single-instrument" specification discussed below.

Single-Instrument Specification

When focusing on population as an outcome variable, a natural way to control for area is to normalize both dependent and independent variables, considering, respectively, the normalized shape metric (see Section 3.2) and population density. Define population density[23] as

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

and denote the normalized version of shape metric S with nS. I then estimate the following, more parsimonious specification:

$$d_{c,t} = a \cdot nS_{c,t} + \mu_c + \rho_t + \eta_{c,t} \tag{42}$$

which includes the endogenous regressor $nS_{c,t}$. I instrument $nS_{c,t}$ with $\widetilde{nS_{c,t}}$, namely the normalized shape metric computed for the potential footprint. The corresponding first-stage equation is

$$nS_{c,t} = \beta \cdot \widetilde{nS_{c,t}} + \lambda_c + \gamma_t + \varepsilon_{c,t}. \tag{43}$$

The same approach can be followed for other outcome variables representing quantities distributed in space - such as the number of employment centers in a city. This approach is less demanding, as it does not require using projected historical population. For this reason, when estimating the single-instrument specification, I choose to construct the shape instrument using the "common rate" model for city expansion (see Section 5.1).

While population density is a meaningful and easily interpretable outcome *per se*, it does not seem as natural to normalize factor prices - wages and rents - by city area. For these outcome variables, the more parsimonious alternative to the double-instrument specification takes the following form:

$$\log(Y_{c,t}) = a \cdot S_{c,t} + \mu_c + \rho_t + \eta_{c,t} \tag{44}$$

where $Y \in (W, p_H)$. This equation does not explicitly control for city area (other than indirectly through

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

city and year fixed effects). Again, the endogenous regressor $S_{c,t}$ is instrumented using $\widetilde{S_{c,t}}$, resulting in the following first-stage equation:

$$S_{c,t} = \sigma \cdot \widetilde{S_{c,t}} + \omega_c + \varphi_t + \theta_{c,t}. \tag{45}$$

All of the specifications discussed above include year and city fixed effects. Although the bulk of my analysis, presented in Section 6, relies on both cross-sectional and temporal variation, a limited number of outcomes, analyzed in Section 7, are available only for a cross-section of cities. In these cases, I resort to cross-sectional versions of the specifications above. In all specifications I cluster standard errors at the city level, to account for arbitrary serial correlation over time in cities.

5.3 Discussion

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.

The exclusion restriction requires that, conditional on city and year fixed effects, this particular time-varying function of geography is only affecting the outcomes of interest though the constraints that it posits to urban shape. The main threats to identification are related to the possibility that the "moving geography" characteristics used in the construction of the instrument directly affect location choices and the outcomes considered, in a time-varying way.

A possible channel of violation of the exclusion restriction is the inherent amenity or disamenity value of geography, to the extent that it may be time-varying. 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 - for example, a fragmented coastline could be a landscape amenity for households; terrain ruggedness could be a production disamenity for firms. If this inherent value changes over time, this poses a threat to identification and the bias could go in different directions. For example, suppose that some of the topographic obstacles that make cities less compact are also inherent "landscape amenities" for consumers - as it may be the case for coasts or lakes. Assume further that this amenity value increases over time. If cities with landscape amenities are also becoming less compact over time, then the consumption disamenity effect of bad shape is partially offset by an increasing consumption amenity effect. This would lead me to underestimate the consumption amenity value of compact shape.[24] On the other hand, suppose that the geographic features that are making city footprints less compact also provide some natural advantage for trade. For example, certain coastal configurations may be particularly favorable to the creation of successful commercial ports (Bleakley and Lin, 2012). Suppose that this is a production amenity that becomes obsolete over time, as naval technology changes. If cities with such configurations are becoming less compact over time, then the production disamenity effect of bad shape detected by my instrument would be confounded by

[24] Note that landscape amenities are likely to be luxury goods, so their relevance for a developing country may be limited.

the declining production amenity value of the city's coastal configuration. In this scenario, I would overestimate the productive disamenity value of bad shape.

I address these concerns through two robustness checks: first, I show that my results are unchanged when I exclude mountainous and coastal cities (Appendix Table A1). The latter are the two most obvious examples of cities where geography could have an intrinsic (dis)amenity value. Second, I show that my results are robust to controlling for year fixed effects interacted with the city's initial shape at the beginning of the panel (Appendix Tables A2 and A3). This more conservative specification allows cities that have different initial geometries to be on different trends, allowing for the initial topographical configuration to have time varying effects.

In general, it should be noted that the instrument is constructed in a way that makes it unlikely to pick up intrinsic effects of geography. The instrument captures a very specific feature of geography: whether the spatial layout of topographic obstacles allows for compact development or not. What makes cities less compact, as captured by the instrument, is neither the generic presence nor the magnitude of topographic constraints. Rather, it is the geometry of developable terrain, once topographic obstacles are excluded. Changes in this geometry over time are dictated primarily by the position of newly encountered topographic obstacles relative to previously encountered ones.[25]

While the presence of large topographic obstacles – such as coasts or mountains – could have a direct amenity value, it is unlikely that the relative position of minor obstacles has. The bulk of the variation in the instrument comes, indeed, from these types of topographic configurations. This is confirmed by the robustness of my first-stage results to excluding mountainous and coastal cities (Appendix Table A1).

A second threat to identification is related to the supply of real estate. Another way in which geographic obstacles may directly affect outcomes such as population and housing rents is by limiting the availability of developable land, constraining housing supply. Albeit in a different context, Saiz (2010) shows that US cities constrained by water bodies and steep terrain have higher housing prices and a more inelastic housing supply. If cities with "bad shapes", as predicted by topography, are also cities with land scarcity and a more inelastic housing supply, this would tend to bias my results towards a positive relationship between non-compactness and housing rents, and lead me to underestimate the value of non-compactness as a consumption disamenity.

A priori, this concern is mitigated by the specific way in which the instrument is constructed. As argued above, the instrument is not based on the share of land that is undevelopable (the main explanatory variable in Saiz, 2010), nor on the magnitude of topographic obstacles, and the results are not driven by cities that are especially land constrained due to being coastal or mountainous. More importantly, the data indicates that Indian cities are not particularly land-constrained by those topographic layouts that induce a non-compact geometry. When I examine empirically the relationship between a city's area and the shape of the "potential footprint" (Table 2, cols. 3 and 6), I find that, all

[25]For example, suppose all topographic obstacles are concentrated East of the center of the city. Newly encountered obstacles along the city's predicted expansion path will be adjacent to previously encountered ones. This does 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.

else being equal, cities that are constrained into "bad" shapes by topographic obstacles tend to occupy, if anything, larger areas than their compact counterparts. In other words, such topographic constraints do not seem to prevent development, but rather induce cities to grow into less compact and potentially more land-consuming shapes. Finally, it should be noted that housing supply elasticity effects of this kind would lead to a positive relationship between rents and non-compactness. As discussed in Section 6.3, I find that non-compact cities are less expensive, suggesting that, even if a housing supply elasticity effect is in place, the "disamenity" effect prevails.

6 Empirical Results: Amenity Value of City Shape

In this section, I address empirically the question of how city shape affects the spatial equilibrium across cities. Recall that, according to the model's predictions discussed in Section 4, if city shape is valued as a consumption amenity by consumers, cities with less compact shapes should be characterized by lower population, higher wages, and lower rents.

6.1 First Stage

Table 2 presents results from estimating the first-stage relationship between city shape and the geography-based instrument described in Section 5.1, for the full sample of city-years for which geometry is observed.[26] This is an interesting descriptive exercise in itself, as it sheds light on the land consumption patterns of Indian cities as a function of their geography. Each observation is a city-year. Panels, A, B, C, and D each correspond to one of the four shape metrics discussed in Section 3.2: respectively, remoteness, spin, disconnection, and range.[27] Higher values of these indexes represent less compact shapes. Cols. 1 and 4 report the first-stage for normalized shape (eq. 43), which is the explanatory variable used in the single-instrument specification. Recall that normalized shape is an area-invariant measure of shape obtained when normalizing a given shape metric by footprint radius. In this specification, the construction of the potential footprint is based on the common rate model for city expansion, outlined in Section 5.1. Cols. 2, 3 and 5, 6 report the first-stage estimates for footprint shape (eq. 40) and area (eq. 41), which are relevant for the double-instrument specification. The dependent variables are city shape, measured in km, and log city area, in square km. The corresponding instruments are the shape of the potential footprint and log projected historical population, as described in Section 5.2. The construction of the potential footprint is based on the city-specific model for city expansion discussed in Section 5.1.

Table 2, panel A focuses on the remoteness index. The remoteness of the *potential* footprint is a highly significant predictor of the remoteness index computed for the *actual* footprint, both in the normalized (col. 1) and non-normalized version (col. 2). Similarly, in col. 3, projected historical population predicts footprint area. As expected, the city-specific, double-instrument specification is a

[26] Most outcome variables considered in the subsequent analysis are observed in a subsample of cities and years, resulting in a smaller sample.

[27] Recall that remoteness (panel A) is the average length of trips to the centroid; spin (panel B) is the average squared length of trips to the centroid; disconnection (panel C) is the average length of within-city trips; range (panel D) is the maximum length of within-city trips. Summary statistics are reported in Table 1.

better predictor for city shape, as highlighted by the higher Angrist-Pischke F statistic. Nevertheless, throughout the paper I will also report results from the common-rate, single instrument specification for robustness. Col. 3 reveals another interesting pattern: the area of the *actual* footprint is positively affected by the remoteness of the *potential* footprint. While this partly reflects the mechanical correlation between shape metric and footprint area, it also suggests that cities which are surrounded by topographic obstacles tend to expand more in space. An interpretation of this result is that the presence of topographic constraints induces a "leapfrog" development pattern, which is typically more land-consuming. It could also reflect an inherent difficulty in planning land-efficient development in constrained contexts, which could result in less parsimonious land use patterns. The results for the remaining shape indicators, reported in panels B, C, and D, are qualitatively similar.

6.2 Population

My main results on population and city shape are reported in Table 3. As in Table 2, each observation is a city-year and each panel corresponds to a different shape metric.

Cols. 1 and 4 report the IV results from estimating the single-instrument specification (eq. 42), which links population density, measured in thousand inhabitants per square km, to (instrumented) normalized shape. The corresponding first-stage estimates are reported respectively in cols. 1 and 4 of Table 2. Recall that normalized shape metrics capture the departure of a city's shape from an ideal circular shape and are invariant to city area, higher values implying longer trips. The IV estimates of the single-instrument specification indicate that less compact cities are associated with a decline in population density. The magnitudes of this effect are best understood in terms of standardized coefficients. Consider the remoteness index (panel A), representing the length, in km, of the average trip to the footprint's centroid. A one-standard deviation increase in normalized remoteness is associated with a decline in population density of almost one standard deviation.

Cols. 2 and 5 of Table 3 report the IV results from estimating the double-instrument specification (eq. 39), which links population to city area and shape, separately instrumented for. The corresponding first-stage estimates are reported, respectively, in cols. 2, 3 and 5, 6 of Table 2. Cols. 3 and 6 of Table 3 report the corresponding OLS estimates. Interestingly, the OLS relationship between population and shape, conditional on area, appears to be positive due to an equilibrium correlation between city size and bad geometry: larger cities are typically also less compact. This arises from the fact that an expanding city has a tendency to deteriorate in shape. Intuitively, a new city typically arises in a relatively favorable geographic location; as it expands in space, however, it inevitably reaches areas with less favorable geography. Once shape is instrumented by geography (col. 2), less compact cities are associated with a decrease in population, conditional on (instrumented) area, city, and year fixed effects. To understand the magnitudes of this effect, consider the remoteness index (panel A). A one-standard deviation increase in normalized remoteness for the average-sized city (which has radius 4.5 km), corresponds to roughly 0.26 km. Holding constant city area, this 0.26 km increase in the average trip to the centroid is associated with an approximate 3 percent decline in population. The results obtained with the double-instrument specification, together with the first-stage estimates in Table 2,

indicate that the observed decline in population density (Table 3, col.1) is driven both by a decrease in population (Table 3, col.2) and by an increase in footprint area (Table 2, col.3).

The results for the remaining shape indicators, reported in panels B, C, and D, are qualitatively similar. The fact that these indexes are mechanically correlated with one another prevents me from including them all in the same specification. However, a comparison of the magnitudes of the IV coefficients of different shape metrics on population suggests that the most salient spatial properties are remoteness (Table 3A) and disconnection (Table 3C), which capture, respectively, the average trip length to the centroid and the average trip length within the footprint. This is plausible, since these two indexes are closer proxies for urban commute patterns. Non-compactness in the periphery, captured by the spin index (Table 3B), appears to have a precisely estimated zero effect on population in the double-instrument specification, whereas the effect of the range index (Table 3D), capturing the longest possible trip within the footprint, is significant but small in magnitude. For brevity, in the rest of my analysis I will mostly focus on the disconnection index, which measures the average within-city trip, without restricting one's attention to trips leading to the centroid. This index is the most general indicator for within-city commutes, and seems suitable to capture trip patterns in polycentric as well as monocentric cities. Unless otherwise specified, in the rest of the tables "shape" will indicate the disconnection index.

As a robustness check, in Appendix Table A1, I re-estimate the double-instrument specification, excluding from the sample cities with severely constrained topographies, namely those located on the coast (panel A) or in high-altitude areas (panel B). Such cities make up about 9 percent of cities in my sample. Out of 457 cities in the initial year of the panel (1951), those located on the coast and in mountainous areas are respectively 24 and 17. Both the first-stage (cols. 1, 3, and 4) and the IV estimates of the effect of shape on population density (col.2) and population (col. 5) are minimally affected by excluding these cities. This shows that the instrument has explanatory power also in cities without extreme topographic constraints,[28] and that my IV results are not driven by a very specific subset of compliers.

Another robustness check is provided in Appendix Table A2. I re-estimate the IV impact of shape on density and population (cols. 4 and 5 from Table 3C), including year fixed effects interacted with each city's shape at the beginning of the panel. This more conservative specification allows cities with different initial geometries to follow different time trends. Results are qualitatively similar to those obtained in Table 3. This mitigates the concern that diverging trends across cities with different geometries might be confounding the results.

6.3 Wages and Rents

The results presented thus far suggest that consumers are affected by city shape in their location choices and that they dislike non-compact shapes. The next step is to quantify the value of "good shape" drawing on the Rosen-Roback framework, as discussed in Section 4.

[28] Recall that my instrument - the shape of the "potential" footprint - is not based on the severity of topographic constraints *per se,* nor on the total share of land lost to such constraints, but is mostly driven by the way in which constraints are positioned in space relative to one another.

Results on wages and rents are reported in Tables 4 and 5, respectively. As discussed in Section 3.3, wages and rents are observed at the district level and the matching between districts and cities is not one to one. For robustness, I show results from two samples: one that includes all matched districts, and a smaller sample restricted to cities that have a one-to-one correspondence with districts (i.e., excluding districts with more than one city).

In Table 4, I report the OLS and IV relationship between average wages and city shape. The dependent variable is the log urban average of individual yearly wages in the city's district, in thousand 2014 Rupees. Cols. 1 and 4 report the IV results from estimating the single-instrument specification (eq. 44), that does not explicitly control for city area. Cols. 2 and 5 report the IV results from estimating the double-instrument specification (eq. 39), which is conditional on instrumented city area. The construction of the potential footprint is based on the common rate model for city expansion in cols. 1, 4, and on the city-specific one in cols. 2, 5 (see Section 5.1). These estimates indicate that less compact shapes, as captured by higher values of the disconnection index, are associated with higher wages both in the OLS and in the IV. This pattern is consistent across different specifications and city-district matching approaches. Appendix Table A3, panel A, shows that these results are also robust to including year fixed effects interacted by initial shape.

This positive estimated impact is compatible with the interpretation that consumers are paying a premium, in terms of foregone wages, in order to live in cities with better shapes. Moreover, if interpreted through the lens of the simple model outlined in Section 4, it suggests that city shape is more a consumption than it is a production amenity. When city area is explicitly included in the regression, the equilibrium relationship between area and wages is negative, which is also consistent with the model's prediction (eq.17). This "amenity" interpretation is subject to a caveat, related to sorting. Cities with different shapes could attract different types of firms and workers, and differences in wages across cities could reflect differences in the skill composition of the workforce (Combes et al., 2008). For instance, low-income and low-skill workers could disproportionately locate in cities with more compact shapes, that are also friendlier to commuters with limited individual transport options.[29] In this case, the finding that compact cities have lower wages may partly reflect systematic differences in workers' productivity. When I examine industry shares, I do not find that compact cities attract particular sectors (Appendix Table A4). However, this does not address heterogeneity at the worker level, which I cannot control for with the information available in the ASI data.

Table 5 reports the same set of specifications for housing rents. In panel A, the dependent variable considered is the log of yearly housing rent per square meter, in 2014 Rupees, averaged throughout all urban households in the district. In panel B, the dependent variable is analogous, but constructed by averaging only the upper half of the distribution of urban housing rents in each district. This addresses the concern that the average reported rent is a downward-biased estimate of the market rental rate due to rent control policies. Many of these estimates appear to be noisy or only borderline significant, with

[29]Indeed, Section 6.5 provides some evidence that compact cities have a larger share of slum dwellers. However, as argued in Section 6.5, it is unlikely that the wages of slum dwellers are included in the wages data I employ, since ASI data cover the formal sector only.

p-values between 0.10 and 0.15, while those in col. 5 are significant at the 10 percent level. This is partly due to small sample size, as only three consecutive years are included. However, a consistent pattern emerges: the impact of disconnected shape on rents is negative in the IV and close to zero in the OLS. Appendix Table A3, panel B, shows that these results are qualitatively similar including year fixed effects interacted by initial shape. This is consistent with the interpretation that consumers are paying a premium in terms of higher housing rents in order to live in cities with better shapes. The estimated relationship between city area and rents is also negative, consistent with the predictions of the conceptual framework (eq. 15).

6.4 Interpreting Estimates through the Lens of the Model

The signs of the estimated coefficients from Tables 3, 4, and 5 are consistent with the interpretation that consumers view compact shape as an amenity. In this sub-section, I use these reduced-form estimates in conjunction with the model to back out the implied welfare loss associated with poor city geometry, as measured through the disconnection index.

Recall that a one-unit increase in shape metric S has a welfare effect equivalent to a decrease in income of λ_θ log points, which, as derived in Section 4 (eq. 34), can be estimated as

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

where α is the share of consumption spent on housing.

My most conservative point estimates for $\widehat{B_W}$ and $\widehat{B_P}$, from the double-instrument specification as estimated in Tables 4 and 5, respectively, amount to 0.038 and -0.516. 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. The implied $\widehat{\lambda_\theta}$ is -0.12. A one standard deviation increase in normalized disconnection for the average-sized city corresponds to an increase in the average within-city trip of about 360 meters. Interpreting this as *potential* commuting length, this suggests that an increase in one-way commutes of 360 meters entails, on average, a welfare loss equivalent to a 0.04 log points or a 4 percent decrease in income.

As a reference, one could compare this figure with the *actual* cost of an extra 360 meters in one's daily one-way commute.[30] Postulating one round-trip per day (720 meters), 5 days per week, this amounts to 225 extra km per year. Assuming that trips take place on foot, at a speed of 4.5 km per hour, a one-standard deviation deterioration in shape amounts to 50 extra commute hours per year, which is equivalent to 2.3 percent of the yearly wage. This figure is roughly 53 percent of the welfare cost I estimate. Assuming instead that trips occur by car, postulating a speed of 25 km per hour, the direct cost of increased commute length amounts to 1.3 percent of the yearly wage, or roughly 30 percent of the welfare cost estimated above.[31]

[30] In order to evaluate this magnitude, it would be interesting to compare this figure to estimates of the value of other amenities. However, the literature on urban amenities in developing countries is limited and I am not aware of such estimates available for India.

[31] To compute the time-cost component of commuting, I estimate hourly wages by dividing the average yearly wage in my sample (93,950 Rs.) by 312 yearly working days and 7 daily working hours, obtaining a figure of 43 Rs. per hour. For the calculation based on car trips, I postulate a fuel efficiency of 5 liters per 100 km, and fuel prices of 77 Rs. per liter. These

The estimated welfare loss appears to be large, if compared to the immediate time and monetary costs of commuting. This is consistent with the interpretation that commuting is perceived as a particularly burdensome activity.[32] It should be emphasized, however, that deteriorating shape may or may not be associated with longer *realized* commutes in equilibrium. When the layout of a city deteriorates, making within-city trips potentially lengthier, consumers may adjust through a range of margins, one of which is their location choices within cities, both in terms of residence and employment. Households may be hurt by poor shape either through longer commutes, or because they are forced to live or work in less preferable locations in order to avoid long commutes.[33]

Next, consider the effect of city shape S on firms. The signs of the reduced-form estimates for B_W, B_N, and B_P are, in principle, compatible with city shape behaving like a production amenity or disamenity. The effect of S on productivity is pinned down by eq. 33 from Section 4:

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

where parameters β and γ represent the shares of labor and tradeable capital in the production function postulated in eq. 7. My most conservative point estimates for $\widehat{B_N}$ and $\widehat{B_W}$ are -0.099 and 0.038, from Tables 3C and 4, respectively. Setting β to 0.4 and γ to 0.3, the implied $\widehat{\lambda_A}$ is -0.003, which indicates a productivity loss of about 0.001 log points for a one standard deviation deterioration in city disconnection. These estimates appear very small. Overall, they suggest that city shape in equilibrium is not affecting firms' productivity, and that the cost of disconnected shape is borne mostly by consumers. This does not indicate that city compactness is *ex ante* irrelevant for firms. Rather, these results indicate that firms do not require a compensation through factor prices for poor city geometry, whereas households do. In equilibrium, firms may be able to optimize against "bad" shape, in a way that consumers cannot. This may be related to the relative location of households and firms within cities. This hypothesis is explored in Section 7.1, that investigates how firms respond to city shape in their location choices within cities, by looking at the spatial distribution of employment.

6.5 Channels and Heterogeneous Effects

The results presented so far provide evidence that, in equilibrium, city shape matters for consumers. In this section, I seek to shed light on the mechanisms through which city shape affects consumers, and on the categories of consumers who are affected the most by poor urban geometry.

Infrastructure If transit is the main channel through which urban shape matters, road infrastructure and the availability of means of transportation may mitigate the adverse effects of poor geometry. In Table 6, I investigate this by interacting city shape with a number of indicators for infrastructure. For

figures are drawn from U.S. Energy Information Administration (2014) and from the Shell India website, accessed in 2014. The 25 km per hour postulated speed is based on Ministry of Urban Development (2008). All monetary values are expressed in 2014 Rupees.

[32] The behavioral literature has provided evidence in support of this interpretation, albeit in the context of developed countries. For example, Stutzer and Frey (2008) estimate that individuals commuting 23 minutes one way would have to earn 19 percent more per month, on average, in order to be fully compensated for the disutility of commuting.

[33] In the empirics, lack of systematic data on residential and workplace locations within cities prevents a detailed investigation of these patterns.

ease of interpretation, I focus on the single-instrument specification (eq. 42), which links normalized shape to population density, and report IV estimates, using both models for city expansion. The shape indicators considered are the disconnection (Panel A) and range index (Panel B). While disconnection is a general indicator for city shape, the range index, capturing the maximum length of within-city trips, appears to be more suitable to capture longer, cross-city trips, which might be more likely to require motorized means of transportation.

This exercise is subject to a number of caveats. An obvious identification challenge lies in the fact that infrastructure is not exogenous, but rather jointly determined with urban shape. I partly address the endogeneity of city infrastructure by employing state-level proxies or measures of infrastructure observed at the beginning of the panel. Another concern is that the effect of infrastructure might be confounded by differential trends across cities of different income levels: cities that have better infrastructure could be cities that also start out with higher income levels, and follow different time trends. To mitigate this problem, I also consider a specification which includes a time-varying proxy for city income: year fixed effects interacted with the number of banks in 1981, as reported in the 1981 Census Town Directory.

In cols. 1, 2, and 3, instrumented normalized shape is interacted with urban road density in 1981, as reported by the 1981 Census Town Directory; this is the first year in which the Census provides this figure. To cope with the potential endogeneity of this variable, in cols. 4, 5, and 6, I consider instead state urban road density in 1991, provided by the Ministry of Transport and Highways. Although the level of statistical significance varies, the coefficients of all three interaction terms are positive. In particular, the interaction between city shape and urban road density is highly significant across specifications and shape indicators (cols. 1, 2, and 3). Overall, this can be interpreted as suggestive evidence that infrastructure mitigates the negative effects of poor geometry. Reassuringly, estimates are qualitatively similar when I include year fixed effects interacted with number of banks, as a time-varying proxy for city income.

Slum population A complementary question relates to who bears the costs of "bad" shape. Compact cities may be more favorable to the poor because they reduce distance, particularly in countries where they cannot afford individual means of transportation. However, higher rental prices in compact cities may reduce the housing floor space that the poor can afford and price them out of the formal market (Bertaud, 2004). In Table 7, I investigate the relationship between city shape and slum prevalence using data from the more recent Census waves.[34] The dependent variables are log slum population (cols. 1, 2 and 3) and log share of slum population (cols. 4, 5 and 6) in a given city-year. I find that cities with less compact shapes have overall fewer slum dwellers, both in absolute terms and relative to total population. Two interpretations are possible. The first is that higher equilibrium rents in compact cities are forcing more households into sub-standard housing. The second interpretation relates to sorting of poorer migrants into cities with more compact shapes, possibly because of their lack of

[34]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).

individual means of transport and consequent higher sensitivity to commute lengths.[35]

7 Empirical Results: Endogenous Responses to City Shape

In this section I consider examples of private and policy responses to deteriorating city geometry.

7.1 Polycentricity

A private kind of response to deteriorating city shape is firms' location choices within cities. As cities grow into larger and more disconnected footprints, businesses might choose to locate further apart from each other, and/or possibly form new business districts.[36] I shed light on this hypothesis by analyzing the spatial distribution of firms listed in the Urban Directories of the 2005 Economic Census. For each urban productive establishment, the latter provides a street address and number of employees, by coarse bins. I detect employment sub-centers within cities, by employing the two-stage, non-parametric approach developed by McMillen (2001), detailed in the Appendix (Section C). This procedure appears suitable for my context, given that it can be fully automated and replicated for a large number of cities. 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. 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.

In Table 8, I estimate the relationship between number of employment centers, city area, and shape, in a cross-section of footprints observed in 2005.[37] I report the IV results from estimating cross-sectional versions of the single-instrument specification (col. 1) and the double-instrument specification (col. 2). The dependent variables are the number of subcenters per square km and the log number of employment subcenters, respectively. Col. 3 reports the OLS version of the specification in col. 2. Due to the small sample size and the limitations of cross-sectional inference, results are noisy and suffer from weak instruments. Nevertheless, several interesting qualitative patterns are detected. Consistent with theories of endogenous subcenter formation, and with the results obtained in the US context by McMillen and Smith (2003), larger cities tend to have more employment subcenters (cols. 2 and 3). Interestingly, conditional on area, less compact cities do not appear to be more polycentric: if anything, they have fewer subcenters (col. 2) This is consistent with the following interpretation: as cities grow into more disconnected shapes, firms continue to cluster in a few locations within a city, and pay higher wages to compensate their employees for the longer commutes they face. This is in line with the finding that more disconnected cities are characterized by higher wages (Section 6.3). More generally, this firm location pattern is consistent with the finding that poor shape entails large losses

[35] 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, in a way consistent with my findings on slum dwellers. Recall, however, that my wages sample covers the formal sector only and is therefore unlikely to include slum dwellers.

[36] Models of endogenous subcenter formations 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), McMillen and Smith (2003), and Agarwal et al. (2012) for a review of the literature on polycentricity.

[37] The sample size is quite small, due to inconsistencies in reported addresses and difficulty in geocoding them.

for consumers, but has negligible impacts on firms (Section 6.4). If places of work were as dispersed as places of residence, commutes should be relatively short, regardless of shape, and poor geometry would have negligible effects. However, if firms are less dispersed than households are, workers will tend to face longer potential commutes to work as cities grow into less compact layouts.

7.2 Floor Area Ratios

The evidence presented so far indicates that poor city shape reduces consumer welfare. Given that most cities cannot expand radially due to their topographies, a question arises on what kind of land use regulations best accommodate city growth. This section provides evidence on the interactions between topography, city shape, and land use regulations. I focus on the most controversial among land use regulations currently in place in urban India: Floor Area Ratios (FARs).

FARs are restrictions on building height expressed in terms of the maximum allowed ratio of a 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 restrictive FARs in Indian cities to suburbanization and sprawl (Sridhar, 2010), as measured by administrative data sources. Bertaud and Brueckner (2005) analyze the welfare impact of FARs in the context of a monocentric city model, estimating that low FARs in Bengaluru carry a welfare cost ranging between 1.5 and 4.5 percent.

I employ data on the maximum allowed residential and non-residential FARs in a cross-section of Indian cities in the mid-2000s, from Sridhar (2010).[38] In Table 9, I explore the interaction between topography and FARs in determining city shape and area. The three first-stage equations presented in Table 2 are augmented with an interaction between each instrument and FAR levels. In cols. 1, 2, and 3, I consider the average of residential and non-residential FARs, whereas in cols. 4, 5, and 6, I focus on residential FARs. The main coefficients of interest are the interaction terms. The interaction between potential shape and FARs in cols. 1 and 4 is negative, and significant in col.1, indicating that cities with higher FARs have a more compact shape than their topography would predict. In particular, estimates in col. 1 indicate that, holding geography constant, increasing FARs by 1 improves normalized shape by 0.09, slightly more than one standard deviation. As shown in Section 6.4, for the average-sized city, a one standard deviation improvement in shape entails a 4 percent welfare increase. The interaction between projected population and FARs appears to have a negative impact on city area (cols. 3 and 6), indicating that laxer FARs cause cities to expand less in space than their projected growth would imply. This is in line with the results obtained by Sridhar (2010), who finds a cross-sectional correlation between restrictive FARs and sprawl using administrative, as opposed to remotely-sensed, data. This interaction term has a negative impact on city shape as well (cols. 2 and 5), suggesting that higher FARs can also slow down the deterioration in city shape that city growth entails. Overall, this suggests that if growing and/or potentially constrained cities are allowed to build vertically, they will do so, rather than expand horizontally and face topographic obstacles.

[38] Given that FARs are updated vert infrequently, the mid-2000s data I employ are a reasonable proxy for FAR values in place throughout the sample period.

In Appendix Table A5 I investigate the impact of FARs, interacted with city shape, on population and density. The specifications proposed here are equivalent to those in Table 3, augmented with interactions between the explanatory variables and FARs. Results are mixed and generally noisy, but there is some suggestive evidence (cols. 2 and 5) that laxer FARs may mitigate the negative impact of non-compactness on population, as highlighted by a positive coefficient on the interaction of instrumented shape and residential FARs. Long potential commutes may indeed matter less in cites that permit taller buildings, since this allows more consumers to live in central locations.

While these regressions take FARs as given, a question might arise on their determinants. Appendix Table A6 reports cross-sectional estimates for FAR values regressed on urban shape and area measured in 2005. There is some weak evidence of FARs being more restrictive in larger cities, consistent with one of the stated objectives of regulators - curbing densities in growing cities. At the same time, FARs appear to be less restrictive in non-compact cities, which could indicate a willingness of policy makers to allow for taller buildings in areas with constrained topographies.[39]

8 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. I find that less compact urban layouts, conducive to longer within-city distances, have sizable welfare costs, and that city compactness affects the spatial equilibrium across cities. City shape thus impacts not only the quality of life in cities, but also influences rural to urban migration patterns by affecting city choice. I also find that city shape does not affect the productivity of firms. This suggests that firms, in equilibrium, are able to optimize against "bad" shape, in a way that consumers cannot, possibly by clustering in the center of cities.

As India prepares to accommodate an unprecedented urban growth in the next decades, the challenges posed by urban expansion are gaining increasing importance in India's policy discourse. On the one hand, policy makers are concerned by 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 a potential source of distortions in urban form (Glaeser, 2011; Sridhar, 2010). My findings contribute to informing this 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 improve urban mobility directly but also to prevent the deterioration in connectivity that fast city growth entails. I find evidence that road infrastructure may mitigate the impact of disconnected city shape. My findings also suggest that urban connectivity can be indirectly improved through another channel: promoting more compact development through 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.[40]

[39] Interestingly, I find no evidence of FARs being driven by geology, including their earthquake proneness. Results are available upon request.

[40] 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

I find that restrictive FARs, among the most controversial of such regulations, result in less compact footprints. While a comprehensive cost-benefit analysis of these types of policies goes beyond the scope of this paper, my results suggest that distortive effects on urban morphology should be accounted for when evaluating the costs of regulations.

In future research, it would be interesting to explore the range of potential mechanisms through which city shape affects consumers. While I provide suggestive evidence that transit is the main channel captured by the shape indexes employed in this paper, urban geometry affects many kinds of urban utilities delivered through spatial networks, including electricity, water, and sewerage (Bertaud, 2004). Moreover, particular urban layouts may promote the separation of a city in different, disconnected neighborhoods and/or administrative units, which could have implications both in terms of political economy and residential segregation. More disaggregated data at the sub-city level will be required to investigate these ramifications.

References

[1] Agarwal, A., G. Giuliano, and C.L. Redfearn (2012), "Strangers in Our Midst: the Usefulness of Exploring Polycentricity", Annals of Regional Science, 48 (2), 433-450.

[2] Akbar, P., V. Couture, G. Duranton, E. Ghani and A. Storeygard (2017), "Accessibility and mobility in urban India", working paper.

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

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

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

[6] Bajari, P. and M. E. Kahn (2004), "The Private and Social Costs of Urban Sprawl: the Lot Size Versus Commuting Tradeoff", working paper.

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

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

[9] 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.

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

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

[12] 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.

relative freedom outside the administrative boundaries of cities.

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

[14] 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.

[15] 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.

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

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

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

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

[22] Glaeser, E. and M. Kahn (2004), "Sprawl and Urban Growth" in *The Handbook of Regional and Urban Economics*, 4, eds. V. Henderson and J. Thisse, Amsterdam: North Holland Press.

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

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

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

[26] Kline, P. and E. Moretti (2014) "People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Policies", Annual Review of Economics, 6, 629-662.

[27] McKinsey Global Institute (2010), "Globalisation and Urban Growth in the Developing World with Special Reference to Asian Countries".

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

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

[30] Ministry of Urban Development (2008), *Study on Traffic and Transportation Policies and Strategies in Urban Areas in India*, Government of India: New Delhi.

[31] Morten, M. and J. Oliveira (2016), "Paving the Way to Development: Costly Migration and Labor Market Integration", working paper.

[32] Munshi, K. and M. Rosenzweig (2016), "Networks and Misallocation: Insurance, Migration, and the Rural-Urban Wage Gap", 108 (1), 46-98.

[33] Stutzer, A. and B. S. Frey (2008), "Stress That Doesn't Pay: The Commuting Paradox", The Scandinavian Journal of Economics, 110 (2), 339-366.

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

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

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

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

[38] 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.

[39] United Nations Department of Economic and Social Affairs, Population Division (2015), "World Urbanization Prospects: The 2014 Revision".

[40] 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.

[41] U.S. Energy Information Administration (2014), Report on India, *http://www.eia.gov/countries/analysisbriefs/India/*.

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

Tables and Figures

Table 1: Descriptive Statistics, full sample

	Obs.	Mean	St.Dev.	Min	Max
Area, km^2	6276	62.63	173.45	0.26	3986.02
Remoteness, km	6276	2.42	2.22	0.20	27.43
Spin, km^2	6276	12.83	39.79	0.05	930.23
Disconnection, km	6276	3.30	3.05	0.27	38.21
Range, km	6276	9.38	9.11	0.86	121.12
Norm. remoteness	6276	0.71	0.06	0.67	2.10
Norm. spin	6276	0.59	0.18	0.50	6.81
Norm. disconnection	6276	0.97	0.08	0.91	2.42
Norm. range	6276	2.74	0.35	2.16	7.17
City population	1440	422869	1434022	5822	22085130
City population density (per km^2)	1440	15011	19124	432	239179
Avg. yearly wage, thousand 2014 Rs.	2009	93.95	66.44	13.04	838.55

Table 2: First Stage

	(1) OLS	(2) OLS	(3) OLS	(4) OLS	(5) OLS	(6) OLS
	Norm. shape of actual footprint	Shape of actual footprint, km	Log area of actual footprint, km^2	Norm. shape of actual footprint	Shape of actual footprint, km	Log area of actual footprint, km^2
	A. Shape Metric: Remoteness			***C. Shape Metric: Disconnection***		
Norm. shape of potential footprint	0.0414***			0.0663***		
	(0.0141)			(0.0239)		
Shape of potential footprint, km		0.349***	0.0564***		1.380***	0.153***
		(0.0903)	(0.0190)		(0.222)	(0.0457)
Log projected historic population		0.392**	0.488***		-1.162***	0.305***
		(0.159)	(0.101)		(0.264)	(0.117)
Observations	6,174	6,174	6,174	6,174	6,174	6,174
F stat shape	9.348	40.042	40.042	7.684	70.539	70.539
F stat area		12.815	12.815		14.955	14.955
	B. Shape Metric: Spin			***D. Shape Metric: Range***		
Norm. shape of potential footprint	0.0249**			0.0754**		
	(0.0102)			(0.0342)		
Shape of potential footprint, km		0.573***	-7.09e-05		1.919***	0.0645***
		(0.201)	(0.000756)		(0.347)	(0.0212)
Log projected historic population		2.451	0.571***		-4.146***	0.309***
		(2.406)	(0.100)		(0.953)	(0.116)
Observations	6,174	6,174	6,174	6,174	6,174	6,174
F stat shape	6.463	35.093	35.093	5.249	74.898	74.898
F stat area		28.671	28.671		16.748	16.748
Model for $\hat{r}$	common rate	city-specific	city-specific	common rate	city-specific	city-specific
City FE	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES

Notes: this table reports estimates of the first-stage relationship between city shape and area, and the instruments discussed in Section 5.1. Each observation is a city-year. Cols. (1) and (4) report the results from estimating equation (43). Cols. (2), (3) and (5), (6) report results from estimating equations (40), (41). The dependent variables are normalized shape (dimensionless), shape, in km, and log area, in km^2, of the actual city footprint. The corresponding instruments are: normalized shape of the potential footprint, shape of the potential footprint, in km, and log projected historic population. The construction of the potential footprint is based on a common rate model for city expansion in cols. (1) and (4), and on a city-specific one in col. (2), (3), (5), (6) - see Section 5.1. Angrist-Pischke F statistics are reported in cols. (2), (3), and (5), (6). Shape is measured by different indexes in different panels. Remoteness (panel A) is the average length of trips to the centroid. Spin (panel B) is the average squared length of trips to the centroid. Disconnection (panel C) is the average length of within-city trips. Range (panel D) is the maximum length of within-city trips. City shape and area are calculated from maps constructed from the DMSP/OLS Night-time Lights dataset (1992-2010) and U.S. Army maps (1951). Estimation is by OLS. All specifications include city and year fixed effects. Standard errors are clustered at the city level.*** p<0.01,** p<0.05,* p<0.1.

Table 3: Impact of City Shape on Population

	(1) IV Population density	(2) IV Log population	(3) OLS Log population	(4) IV Population density	(5) IV Log population	(6) OLS Log population
	A. Shape Metric: Remoteness			***C. Shape Metric: Disconnection***		
Norm. shape of actual footprint	-315.2***			-254.6***		
	(94.19)			(80.01)		
Shape of potential footprint, km		-0.137**	0.0338***		-0.0991**	0.0249***
		(0.0550)	(0.0108)		(0.0386)	(0.00785)
Log area of actual footprint, km^2		0.785***	0.167***		0.782***	0.167***
		(0.182)	(0.0306)		(0.176)	(0.0305)
Observations	1,329	1,329	1,329	1,329	1,329	1,329
F stat shape	13.107	40.042		12.014	70.539	
	B. Shape Metric: Spin			***D. Shape Metric: Range***		
Norm. shape of actual footprint	-110.9***			-84.14***		
	(37.62)			(27.70)		
Shape of potential footprint, km		-0.00101	0.000887***		-0.0284***	0.00763***
		(0.000657)	(0.000277)		(0.0110)	(0.00227)
Log area of actual footprint, km^2		0.547***	0.197***		0.746***	0.171***
		(0.101)	(0.0278)		(0.164)	(0.0293)
Observations	1,329	1,329	1,329	1,329	1,329	1,329
F stat shape	7.226	35.093		10.898	74.898	
Model for f̂	common rate	city-specific		common rate	city-specific	
City FE	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES

Notes: this table reports estimates of the relationship between city shape and population. Each observation is a city-year. Col. (1) and (4) report the results from estimating equation (42) (single-instrument specification). The dependent variable is population density, in thousands of inhabitants per km^2. The explanatory variable is normalized shape (dimensionless). Cols. (2), (3), (5) and (6) report the results from estimating equation (39) (double-instrument specification). The dependent variable is log city population. The explanatory variables are log city area, in km^2, and city shape, in km. In cols. (1), (4) and (2), (5) the estimation is by IV. The instruments are discussed in Section 5.2. Angrist-Pischke F statistics for the shape variable are reported. The construction of the potential footprint is based on a common rate model for city expansion in cols. (1) and (4) and on a city-specific one in col. (2) (5) - see Section 5.1. Col. (3), (6) report the same specification as col. (2), (5), estimated by OLS. Shape is measured by different indexes in different panels. Remoteness (panel A) is the average length of trips to the centroid. Spin (panel B) is the average squared length of trips to the centroid. Disconnection (panel C) is the average length of within-city trips. Range (panel D) is the maximum length of within-city trips. City shape and area are calculated from maps constructed from the DMSP/OLS Night-time Lights dataset (1992-2010) and U.S. Army maps (1951). Population is drawn from the Census of India (1951, 1991, 2001, 2011). All specifications include city and year fixed effects. Standard errors are clustered at the city level.*** p<0.01,** p<0.05,* p<0.1.

Table 4: Impact of City Shape on Wages

	Dependent variable: log wage					
	(1)	(2)	(3)	(4)	(5)	(6)
	IV	IV	OLS	IV	IV	OLS
	All districts			*Only districts with one city*		
Shape of actual footprint, km	0.109***	0.0381	0.0538***	0.0996***	0.0626	0.0586***
	(0.0275)	(0.0386)	(0.0169)	(0.0336)	(0.0536)	(0.0146)
Log area of actual footprint, km^2		-0.409	-0.0478		-0.167	-0.00936
		(0.390)	(0.0351)		(0.465)	(0.0503)
Observations	1,943	1,943	1,943	1,021	1,021	1,021
F stat shape	10.317	14.922		7.016	8.530	
Model for $\hat{r}$	common rate	city-specific		common rate	city-specific	
City FE	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES

Notes: this table reports estimates of relationship between city shape and average wages. Each observation is a city-year. The dependent variable is the log urban average of individual yearly wages in the city's district, in thousand 2014 Rupees. The explanatory variables are city shape, in km, and log city area, in km^2. Cols. (1) and (4) report the results from estimating equation (44) (single-instrument specification) by IV. Cols. (2) and (5) report the results from estimating equation (39) (double-instrument specification) by IV. Instruments are described in Section 5.2. Angrist-Pischke F statistics for the shape variable are reported. Cols. (3) and (6) report the same specification, estimated by OLS. The construction of the potential footprint is based on a common rate model for city expansion in cols. (1) and (4), and on a city-specific one in cols. (2) and (5) - see Section 5.1. In cols. (4), (5), (6), the sample is restricted to districts containing only one city. Shape is captured by the disconnection index, which measures the average length of trips within the city footprint, in km. City shape and area are calculated from maps constructed from the DMSP/OLS Night-time Lights dataset (1992-2010). Wages are from the Annual Survey of Industries, waves 1990, 1994, 1995, 1997, 1998, 2009, 2010. All specifications include city and year fixed effects. Standard errors are clustered at the city level. *** p<0.01,** p<0.05,* p<0.1.

Table 5: Impact of City Shape on Housing Rents

	(1) IV	(2) IV	(3) OLS	(4) IV	(5) IV	(6) OLS
	All districts			*Only districts with one city*		
	A. Dependent variable: log yearly rent per square meter					
Shape of actual footprint, km	-0.0107	-0.595	-0.000118	-0.0101	-0.516*	-0.00996
	(0.152)	(0.391)	(0.0493)	(0.208)	(0.282)	(0.0719)
Log area of actual footprint, km^2		-1.473	-0.0101		-0.865	-0.0527
		(1.400)	(0.0821)		(0.864)	(0.105)
Observations	839	840	841	445	445	445
F stat shape	8.466	11.130		5.135	16.817	
	B. Dependent variable: log yearly rent per square meter, upper 50%					
Shape of actual footprint, km	-0.0174	-0.674	-0.00786	-0.0196	-0.591*	-0.00904
	(0.171)	(0.426)	(0.0563)	(0.239)	(0.321)	(0.0807)
Log area of actual footprint, km^2		-1.580	0.00568		-0.974	-0.0357
		(1.529)	(0.0946)		(0.982)	(0.120)
Observations	839	840	841	445	445	445
F stat shape	8.466	11.130		5.135	16.817	
Model for $\hat{f}$	common rate	city-specific		common rate	city-specific	
City FE	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES

Notes: this table reports estimates of relationship between city shape and average housing rents. Each observation is a city-year. In panel A, the dependent variable is the log urban average of housing rent per square meter in the city's district, in 2014 Rupees. In panel B the dependent variable is analogous, but the average is calculated considering only the top 50% of the district's distribution of rents per square meter. The explanatory variables are city shape, in km, and log city area, in km^2. Cols. (1) and (4) report the results from estimating equation (44) (single-instrument specification) by IV. Cols. (2) and (5) report the results from estimating equation (39) (double-instrument specification) by IV. Instruments are described in Section 5.2. Angrist-Pischke F statistics for the shape variable are reported. Cols. (3) and (6) report the same specification, estimated by OLS. The construction of the potential footprint is based on a common rate model for city expansion in cols. (1) and (4), and on a city-specific one in cols. (2) and (5) - see Section 5.1. In cols. (4), (5), (6), the sample is restricted to districts containing only one city. Shape is captured by the disconnection index, which measures the average length of trips within the city footprint, in km. City shape and area are calculated from maps constructed from the DMSP/OLS Night-time Lights dataset (1992-2010). Housing rents are from the NSS Household Consumer Expenditure Survey, rounds 62 (2005-2006), 63 (2006-2007) and 64 (2007-2008). All specifications include city and year fixed effects. Standard errors are clustered at the city level. *** p<0.01,** p<0.05,* p<0.1.

Table 6: Interactions of City Shape with Infrastructure

Dependent variable: population density						
	(1)	(2)	(3)	(4)	(5)	(6)
	IV	IV	IV	IV	IV	IV
				A. Shape Metric: Disconnection		
Norm. shape	-327.6***	-306.2***	-204.8***	-275.6***	-254.3***	-136.3***
	(103.8)	(98.46)	(54.16)	(93.29)	(89.68)	(45.02)
Norm. shape x urban road density, 1981	1.955**	1.848**	1.485***			
	(0.827)	(0.779)	(0.451)			
Norm. shape x state urban road density, 1991				3.15e-06	0.000138	0.000638**
				(0.000339)	(0.000319)	(0.000315)
Observations	1,126	1,126	1,126	1,118	1,118	1,118
F stat shape	8.732	8.461	17.333	22.856	21.120	30.246
F stat interaction term	44.642	45.421	48.661	263.610	254.551	418.753
				B. Shape Metric: Range		
Norm. shape	-99.33***	-94.23***	-52.63***	-83.63***	-78.21***	-35.37***
	(34.36)	(33.23)	(12.09)	(28.61)	(27.96)	(9.469)
Norm. shape x urban road density, 1981	1.671**	1.607**	1.231***			
	(0.831)	(0.797)	(0.430)			
Norm. shape x state urban road density, 1991				0.000109	0.000214	0.000757**
				(0.000388)	(0.000368)	(0.000315)
Observations	1,126	1,126	1,126	1,118	1,118	1,118
F stat shape	7.705	7.372	23.996	15.483	14.375	43.605
F stat interaction term	49.083	49.834	50.459	278.326	266.603	485.781
Model for $\hat{r}$	common rate	common rate	city-specific	common rate	common rate	city-specific
City FE	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES
Year FE x Banks in 1981	NO	YES	YES	NO	YES	YES

Notes: this table investigates the impact of shape, interacted with infrastructure, on population density. Each observation is a city-year. All cols. report IV estimates of a specification similar to equation (42) (single-instrument specification), augmented with an interaction between normalized shape and different infrastructure proxies. The dependent variable is population density, in thousands of inhabitants per km^2. The construction of the potential footprint is based on a common rate model for city expansion in cols. (1), (2), (4), (5) and on a city-specific one in cols. (3) and (6) - see Section 5.1. Angrist-Pischke F statistics for the shape variable and for the interaction term are reported in each column. The shape metrics considered are normalized disconnection in Panel A, and normalized range in Panel B. Disconnection is the average length of within-city trips. Range is the maximum length of within-city trips. In cols. (1) , (2) and (3), normalized shape is interacted with urban road density in the city, as reported in the 1981 Census. In cols. (4), (5), and (6), normalized shape is interacted with state urban road density in year 1991 (source: Ministry of Road Transport and Highways). City shape and area are calculated from maps constructed from the DMSP/OLS Night-time Lights dataset (1992-2010) and U.S. Army maps (1951). Population is drawn from the Census of India (1951, 1991, 2001, 2011). All specifications include city and year fixed effects. Specifications in cols. (2), (3), (5) and (6) also include year fixed effects interacted with the number of banks in 1981. Standard errors are clustered at the city level. *** p<0.01,** p<0.05,* p<0.1.

Table 7: Impact of City Shape on Slum Population

	(1) IV	(2) IV	(3) OLS	(4) IV	(5) IV	(6) OLS
	Log slum population			***Log slum population share***		
Shape of actual footprint, km	-0.0253	-0.134**	-0.0353**	-0.0502***	-0.113**	-0.0498***
	(0.0171)	(0.0540)	(0.0158)	(0.0185)	(0.0561)	(0.0159)
Log area of actual footprint, km^2		0.592	0.134*		0.264	0.0501
		(0.492)	(0.0781)		(0.633)	(0.0884)
Observations	946	946	946	946	946	946
F stat shape	1016.198	50.126		1016.198	50.126	
Model for $\hat{r}$	common rate	city-specific		common rate	city-specific	
City FE	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES

Notes: this table reports estimates of relationship between city shape and slum population. Each observation is a city-year. Cols. (1) and (4) report the results from estimating equation (44) (single-instrument specification) by IV. Cols. (2) and (5) report the results from estimating equation (39) (double-instrument specification) by IV. Cols. (3) and (6) report the same specification, estimated by OLS. The explanatory variables are shape, in km, and log city area, in km^2. Instruments are described in Section 5.2. Angrist-Pischke F statistics for the shape variable are reported. The construction of the potential footprint is based on a common rate model for city expansion in cols. (1) and (4) and on a city-specific one in cols. (2) and (5) - see Section 5.1. The dependent variables are log slum population in the city (cols. (1), (2), (3)), log of the share of slum to total population in the city (cols. (4), (5), (6)). Data on slums is drawn from the 1991, 2001, and 2011 Census. Shape is captured by the disconnection index, which measures the average length of trips within the city footprint, in km. City shape and area are calculated from maps constructed from the DMSP/OLS Night-time Lights dataset (1992-2010). All specifications include city and year fixed effects. Standard errors are clustered at the city level. *** p<0.01,** p<0.05,* p<0.1.

Table 8: Impact of City Shape on the Number of Employment Subcenters, 2005

	(1) IV	(2) IV	(3) OLS
	Subcenters /km^2	Log subcenters	Log subcenters
Norm. shape of actual footprint	-0.317		
	(0.455)		
Shape of actual footprint, km		-0.0623*	-0.0579***
		(0.0377)	(0.0154)
Log area of actual footprint, km^2		0.606***	0.571***
		(0.124)	(0.0567)
Observations	187	187	187
F stat shape	4.34	6.50	
Model for $\hat{r}$	common rate	city-specific	

Notes: This table investigates the cross-sectional relationship between city shape, city area, and the number of employment subcenters in year 2005. Each observation is a city in year 2005. The dependent variables are the number of subcenters per km^2 (col. (1)) and the log number of employment subcenters (cols. (2) and (3)). Col. (1) reports estimates from a cross-sectional version of equation (42) (single-instrument specification), estimated by IV. Col. (2) reports estimates from a cross-sectional version of equation (39) (double-instrument specification), estimated by IV. Col. (3) presents the same specification, estimated by OLS. The construction of the potential footprint is based on a common rate model for city expansion in col. (1), and on a city-specific one in in col. (2) - see Section 5.1 Angrist-Pischke F statistics for the shape variable are reported. Shape is measured by the disconnection index, as the average length of trips within the city footprint, in km. The procedure used to determine the number of subcenters in each city is drawn from McMillen (2001) and detailed in the Appendix. Data on the spatial distribution of employment is derived from the urban Directories of Establishments, from the 2005 Economic Census. City shape and area are calculated from maps constructed from the DMSP/OLS Night-time Lights dataset, in year 2005.

Table 9: Impact of FARs on City Shape

	(1) OLS	(2) OLS	(3) OLS	(4) OLS	(5) OLS	(6) OLS
	Norm. shape of actual footprint	Shape of actual footprint km	Log area of actual footprint km²	Norm. shape of actual footprint	Shape of actual footprint km	Log area of actual footprint km²
Norm. shape of potential footprint	0.435***			0.337***		
	(0.123)			(0.117)		
Norm. shape of potential footprint x FAR	-0.0908**			-0.0571		
	(0.0452)			(0.0371)		
Log projected historic population		2.998	1.984**		1.322	1.086*
		(2.758)	(0.795)		(1.841)	(0.647)
Log projected historic population x FAR		-1.958*	-0.686**		-1.315*	-0.342
		(1.023)	(0.319)		(0.779)	(0.270)
Shape of potential footprint, km		0.156	-0.186		0.235	0.0928
		(1.182)	(0.233)		-0.671	(0.187)
Shape of potential footprint, km x FAR		0.665	0.135		0.617**	0.0244
		(0.487)	(0.106)		(0.279)	(0.0810)
Observations	1,183	1,183	1,183	1,183	1,183	1,183
Model for f̂	common rate	city-specific	city-specific	common rate	city-specific	city-specific
FAR	average	average	average	residential	residential	residential
City FE	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES

Notes: this table reports estimates of the first-stage relationship between Floor Area Ratios, city shape, and area. Each observation is a city-year. Cols. (1), (2), (3) and (4), (5), (6) report the same specifications reported in Table 2, panel C, cols. (4), (5), (6), with the addition of interactions between each instrument and FARs. FARs are drawn from Sridhar (2010) and correspond to the maximum allowed Floor Area Ratios in each city 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. Cols. (1), (2), (3) consider the average of residential and non-residential FARs, while cols. (4), (5), (6) only consider residential FARs. Shape is captured by the disconnection index, which measures the average length of trips within the city footprint, in km. City shape and area are calculated from maps constructed from the DMSP/OLS Night-time Lights dataset (1992-2010) and U.S. Army maps (1951). Population is from the Census (1951, 1991, 2001, 2011). Population density is measured in thousands of inhabitants per km^2. All specifications include city and year fixed effects. Standard errors are clustered at the city level. *** p<0.01,** p<0.05,* p<0.1.

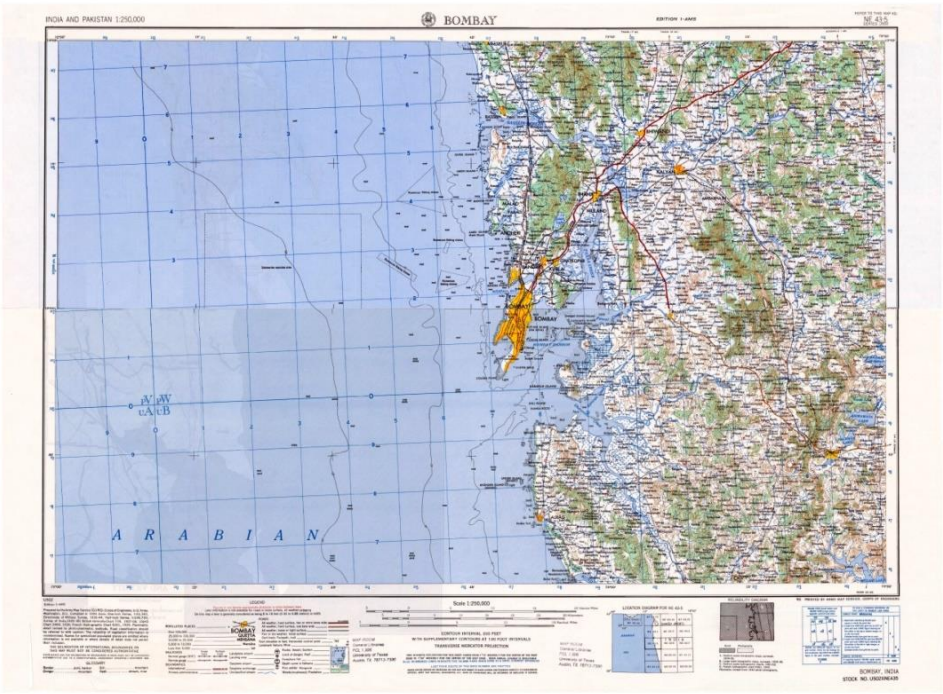

Figure 1A

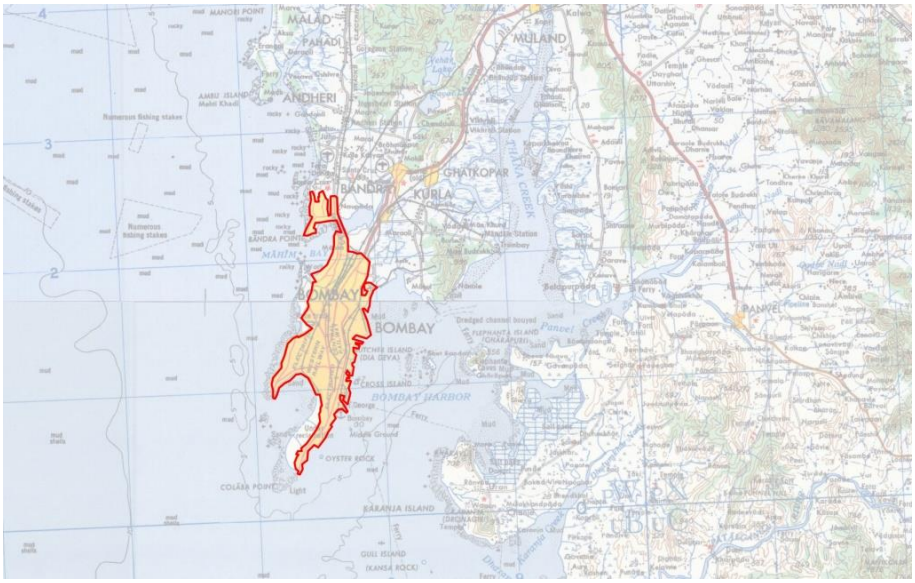

Figure 1B

Figure 1
U.S. Army India and Pakistan Topographic Maps

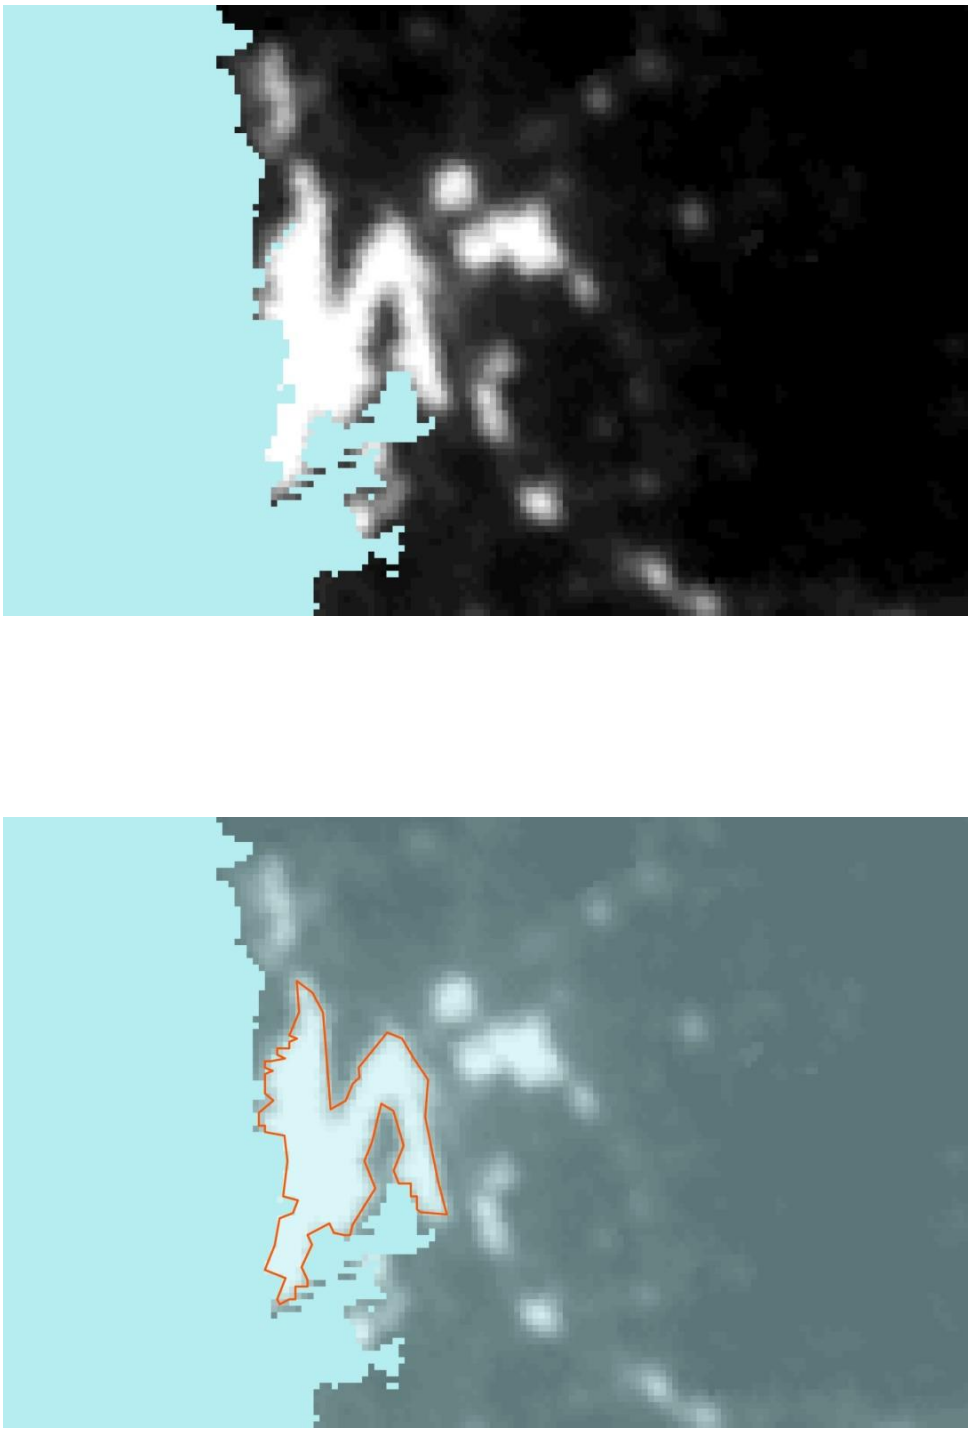

Figure 2
DMS/OLS nighttime lights, year 1992, luminosity threshold : 40.

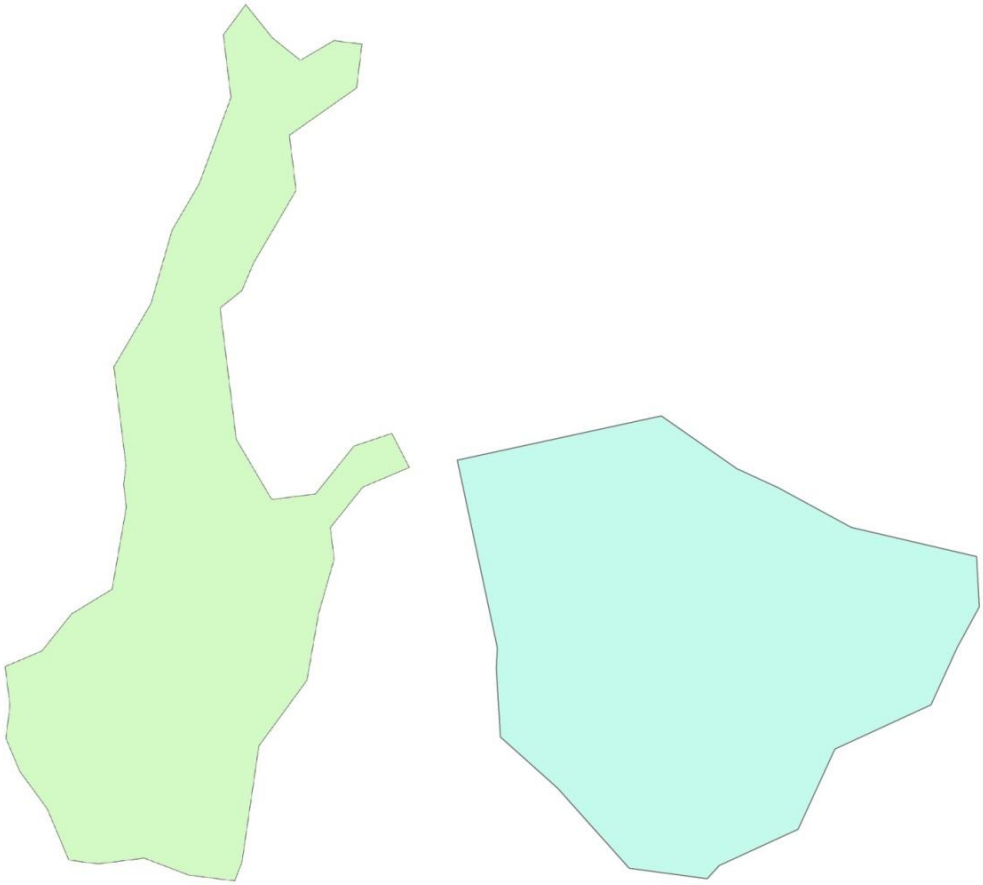

	Kolkata		Bengaluru	
Shape metric		Normalized		Normalized
remoteness, km	14.8	0.99	10.3	0.69
spin, km^2	288.4	1.29	120.9	0.54
disconnection, km	20.2	1.35	14	0.94
range, km	62.5	4.18	36.6	2.45

Figure 3
Shape metrics: an example

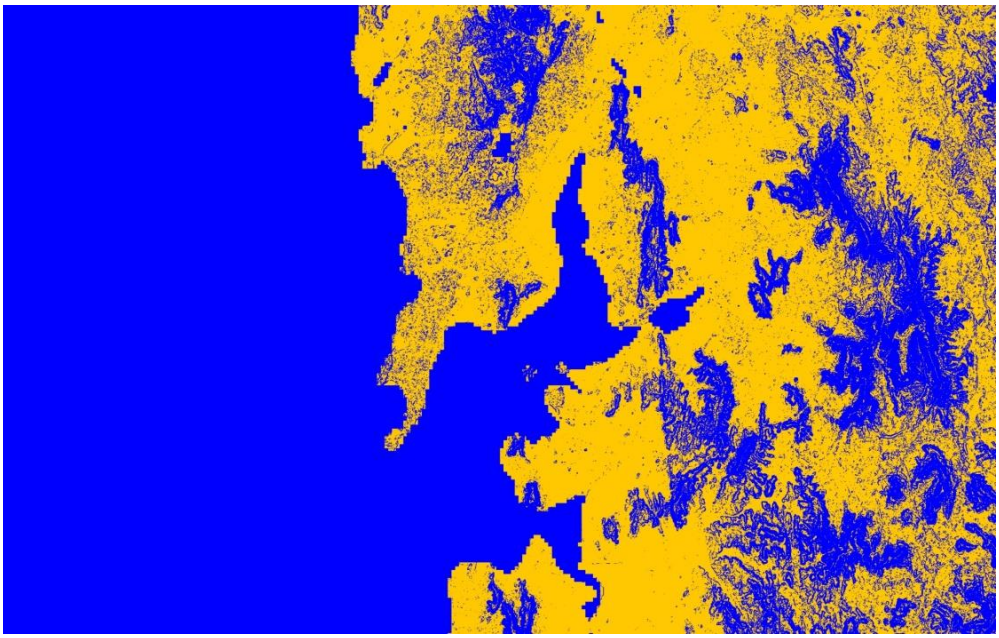

Figure 4
Developable vs. constrained land

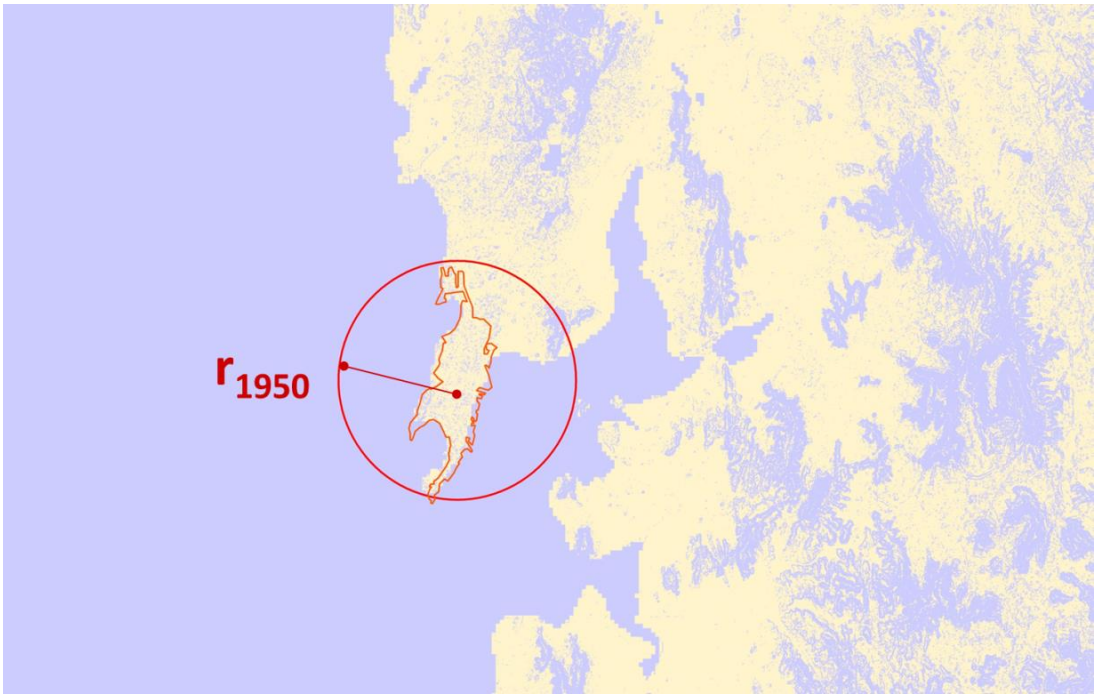

Figure 5a

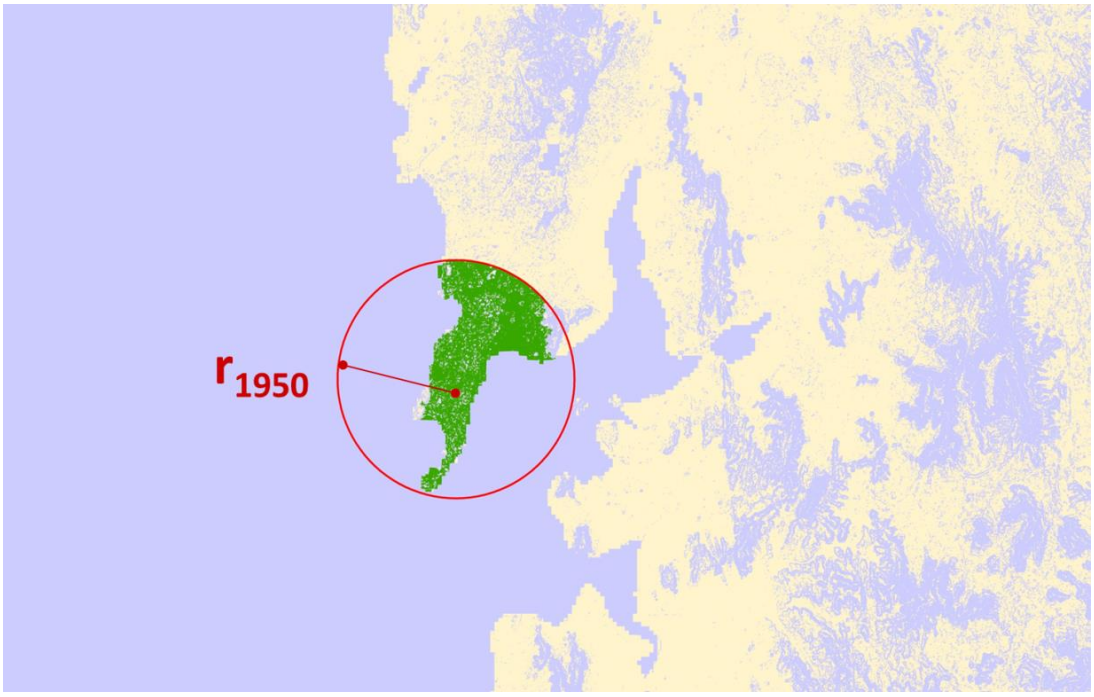

Figure 5b

Figure 5
Instrument construction

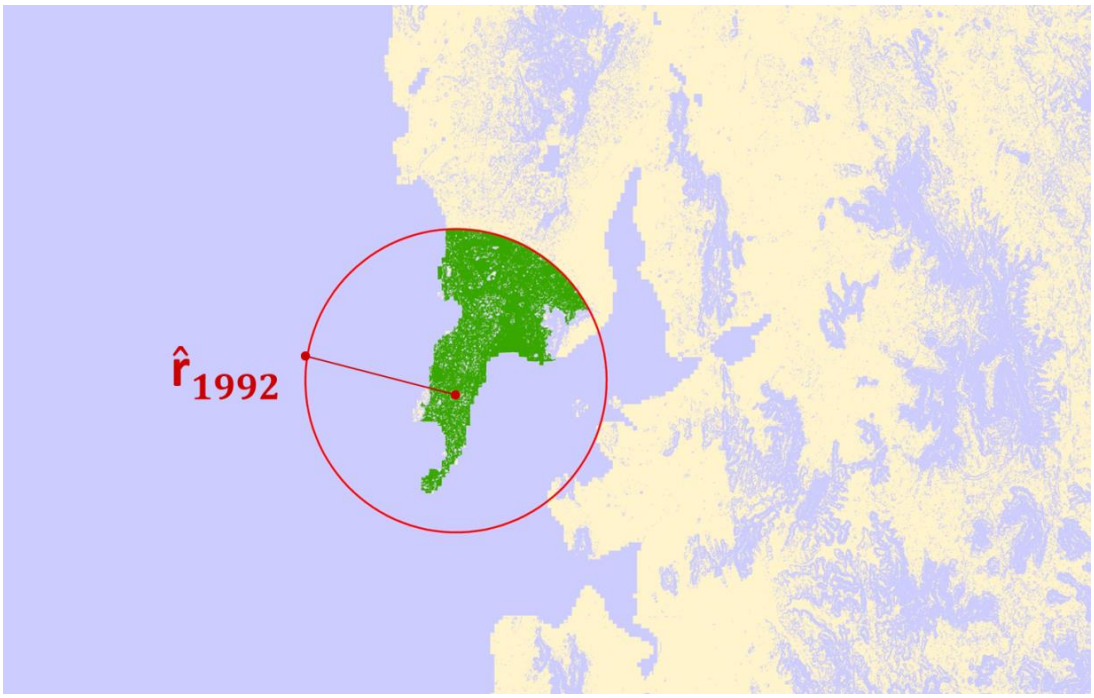

Figure 5c

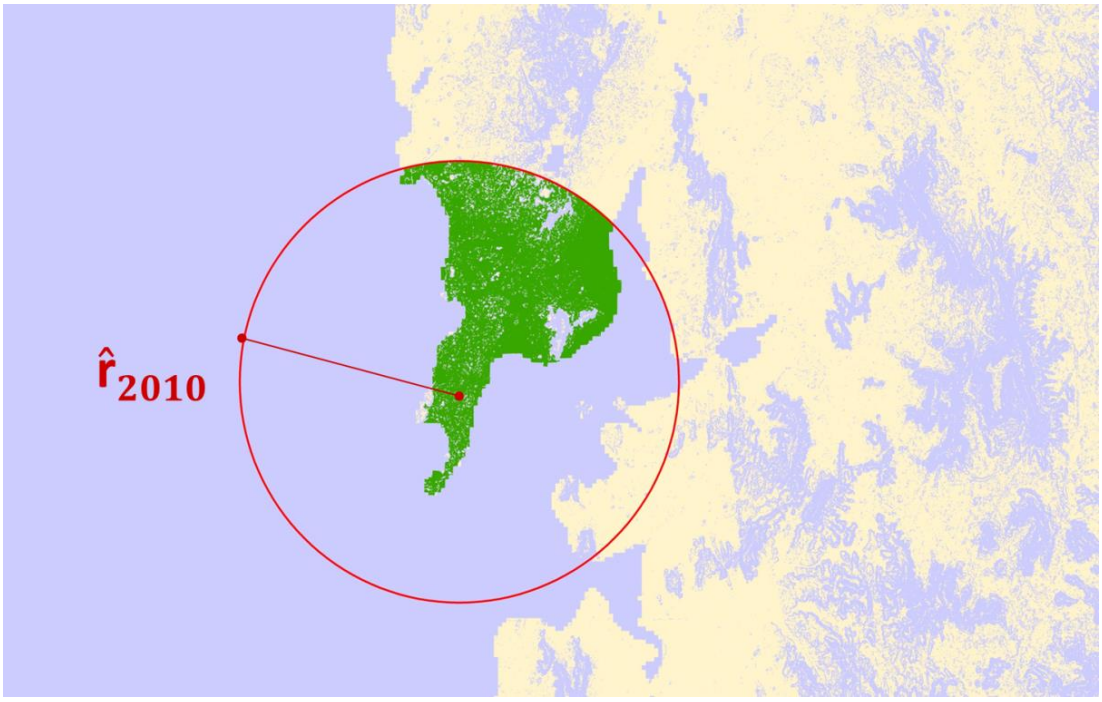

Figure 5d

Figure 5
Instrument construction

Appendix

A. Data Sources

My empirical analysis is based on a newly assembled, unbalanced panel of city-year data, covering all Indian cities for which a footprint could be retrieved based on the methodology explained below. For each city-year in the panel, I collect data on the geometric properties of the footprint, the city's topography, and various economic aggregate outcomes - in particular, population, average wages and average housing rents.

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 am attributing to the footprints observed in these maps the year 1951, corresponding to the closest Census year.

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 km). 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 basic 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.[41] 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.[42] In my full panel (including years 1951 and 1992-2010), the average city footprint occupies

[41] 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.

[42] For years covered by both sources (1990, 1995, 2000), my maps also appear consistent with those from the GRUMP -

an area of approximately 63 square km.[43]

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.[44]

It is well known that urban maps based on night-time lights will tend to inflate urban boundaries, due to "blooming" effects (Small et al., 2005).[45] 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 agglomeration, 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 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.

By combining the US Army maps (1950s) with yearly maps obtained from the night-time lights dataset (1992-2010), I thus assemble an unbalanced panel of urban footprints. The resulting panel dataset 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 later years. The number of cities in the panel ranges from 352 to 457, depending on the year considered.

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.

Shape Metrics

The indicators of city shape that I employ (Angel et al., 2009a, 2009b),[46] are used in landscape ecology and urban studies as proxies for the length within-city trips. They are based on the distribution of points around the polygon's centroid or within the polygon, and are measured in kilometers. The centroid of a polygon, or center of gravity, is the point that minimizes the sum of squared Euclidean distances

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).

[43] My results are robust to using alternative luminosity thresholds between 20 and 40. Results are available upon request.

[44] 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.

[45] 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).

[46] I am thankful to Vit Paszto for help with the ArcGis shape metrics routines. I have renamed some of the shape metrics for ease of exposition.

between itself and each vertex.

(i) The *remoteness* index is the average distance between all interior points and the centroid. It can be considered as a proxy for the average length of commutes to the urban 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. To give a concrete example, this index would have particularly high values for footprints that have "tendril-like" projections.

(iii) The *disconnection* index captures the average distance between all pairs of interior points. It can be considered as a proxy for commutes within the city, without restricting one's attention to those to or from to the center. Throughout the paper, I will employ this as my benchmark indicator.

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

All these measures are correlated mechanically 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. Alternatively, it is possible to normalize each of these indexes, 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, obtaining what I define *normalized remoteness*, *normalized spin,* etc. One way to interpret these normalized metrics is as deviations of a polygon's shape from that of a circle, the shape that minimizes all the indexes above.

Geography

For the purposes of constructing the 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 percent. I draw upon the highest resolution sources available: 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 illustrates this classification for the Mumbai area.

Outcome Data: Population

City-level data for India is difficult to obtain (Greenstone and Hanna, 2014). 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.[47] As explained in Section 5.1, historical population (1871-1941) is used to construct one of the two versions of my instrument, whereas population drawn from more recent waves (1951, 1991, 2001, and 2011) is used as an outcome variable.

[47] Historical population totals were taken 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 http://www.censusindia.gov.in/DigitalLibrary/Archive_home.aspx.

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.

Besides population, the Census provides a number of other city-level variables, which, however, are not consistently available for all Census years and for all cities. I draw data on urban road length in 1991 from the 1991 Town Directory. In recent Census waves (1991, 2001, 2011), data on slum population and physical characteristics of houses are available for a subset of cities.

Outcome Data: Wages, 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): 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 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. For robustness, I also report results obtained focusing on districts that contain one city only. 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), waves 1990, 1994, 1995, 1997, 1998, 2009, 2010. 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. The average individual yearly wage in this panel amounts to 94 thousand Rs at current prices.

A drawback of the ASI data is that it covers the formal manufacturing sector only.[48] This may affect the interpretation of my results, to the extent that this sector is systematically over- or underrep-

[48] 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.

resented in cities with worse shapes. I provide some suggestive evidence on the relationship between city shape and the local industry mix using data from the Economic Census, a description of which is provided in Section 3.5 below. 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 (Appendix Table A4).

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. I focus on rounds 62 (2005-2006), 63 (2006-2007), and 64 (2007-2008), since they are the only ones for which the urban data is representative at the district level and which report total dwelling floor area as well. I use this information to construct a measure of rent per square meter. The average yearly total rent paid in this sample amounts to about 25 thousand Rs., whereas the average yearly rent per square meter is 603 Rs., at current prices. 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.

Other Data

Data on state-level infrastructure is taken from the Ministry of Road Transport and Highways, Govt. of India and from the Centre for Industrial and Economic Research's Industrial Databooks (CIER, 1990).

Data on the maximum permitted Floor Area Ratios for a small cross-section of Indian cities (55 cities in my sample) 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. For a detailed discussion of FARs in India, see Sridhar (2010) and Bertaud and Brueckner (2005).

Data on the the industry mix of cities is derived from rounds 3, 4 and 5 of the Economic Census, collected in 1990, 1998 and 2005 respectively. The Economic Census is a complete enumeration of all productive establishments, with the exception of those involved in crop production, conducted by the Indian Ministry of Statistics and Programme Implementation. For each establishment, the Census reports sector (according to the National Industry Code classification) and number of workers. The Economic Census provides state and district identifiers, but town identifiers are not provided to the general public. In order to approximately identify cities within each district, I rank cities by total number of workers, and compare this ranking with that obtainable in the population Census that is closest in time - 1991, 2001 or 2011. Matching cities by their rank within each district allows me to create a tentative crosswalk between the economic and the population Census.[49]

[49]The definition of sectors, identified by NIC codes, varies over Economic Census waves. I define sectors based on a

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.[50]

I use these data to compute the number of employment subcenters in each city, following the two-stage, non-parametric approach described in McMillen (2001). Of the various methodologies proposed in the literature, this appears to be the most suitable for my context, given that it does not require a detailed knowledge of each study area, and 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 C of this Appendix.

B. Spatial Equilibrium within the City and Topographic Constraints

The conceptual framework adopted in this paper is one of spatial equilibrium across cities, in which consumers have the same utility across cities, and an irregular city layout is postulated to be a potential disamenity. Empirically, I find cities with irregular layouts to be characterized by lower populations and lower rents, which I interpret as due to a disamenity effect from non-compact city shapes. The question may however arise on how these findings reconcile with the predictions of an alternative modeling approach based on the notion of spatial equilibrium within, as well as across cities. In this section, I consider a simple model of spatial equilibrium within a city that has an irregular layout due to topographic or regulatory constraints. Individuals in this city are indifferent across locations within the city, as well as across cities. Under this alternative model, I show 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. My empirical findings are consistent with these predictions.

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: individuals are indifferent between locating in the given city or anywhere else. Each individual earns a wage w and consumes L units of land; for simplicity, I consider both to be 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

coarse, 1-digit NIC code disaggregation so as to maintain consistency across waves.

[50] My results are robust to excluding firms whose address can only be approximately located by Google Maps (available upon request).

costs 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:

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

which yields first-order condition

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

known as the Alonso-Muth condition, that holds for any d in the city.

The rent function in the city is thus

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

with rents declining with distance in a way that offsets the increase in transportation costs. Utility is equalized at any distance from the CBD. Note that, in deriving this well-known result, no particular assumptions have been made on the spatial layout of the 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. In order to convey the intuitions of the model in the simplest way and with closed-form solutions, I consider a linear city, in which people live on one dimension along a line. 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 $0 < \alpha_1 < \alpha_2$. This layout is illustrated in Figure B.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 B.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. 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). This simple, one-dimensional setting simplifies the algebra considerably, but the intuitions and qualitative predictions of the model carry over to a 2-dimensional city, for which the unconstrained, distance-minimizing city structure is a circle centered on the CBD (as in the standard Alonso-Mills-Muth framework). The model's intuitions also apply to a city with multiple constraints.

Let us now 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

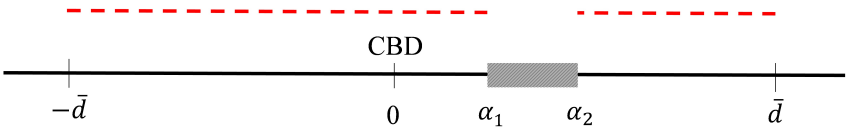

Figure B.1: Population in a linear city with a constraint

that N is sufficiently large relative to the size of the topographic obstacle,[51] the population in the constrained city will distribute itself around the CBD as in Figure B.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 $NL=\bar{d}+\bar{d}-(\alpha_2-\alpha_1)$ which implies:

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

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 from the center to the edge of the city at distance $\overline{d}$. The next step is pinning down 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 (B.3) one can obtain $r(0)=\underline{r}+\frac{\tau}{L}\bar{d}$ which implies

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

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

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

This is the equilibrium rent gradient for a given N. 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 (B.6) into the utility function, this condition becomes:

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

Condition (B.7) pins down N and allows us to perform some comparative statics. As the size of the constraint $(\alpha_2-\alpha_1)$ increases, the city's population N declines. Intuitively, a city with topographic

[51] Specifically, NL has to be greater than $(\alpha_1+\alpha_2)$. When NL is smaller, the population will distribute itself in a way that never reaches the constraint and there will be no occupied locations past distance α_1 on the constrained side of the city. In equilibrium, the condition $NL>(\alpha_1+\alpha_2)$ will be met provided that the city pays a high enough wage relative to transportation costs.

constraints is one in which, for a given maximal distance from the CBD, there are fewer locations available, and in equilibrium it will host a smaller population. This prediction is borne in my data.

Next, 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 $\underline{C}$:

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

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

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

As expected, the larger the constraint, as captured by $(\alpha_2 - \alpha_1)$, the smaller the constrained city's population, relative to the unconstrained one.

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

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

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

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

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 B.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 B.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

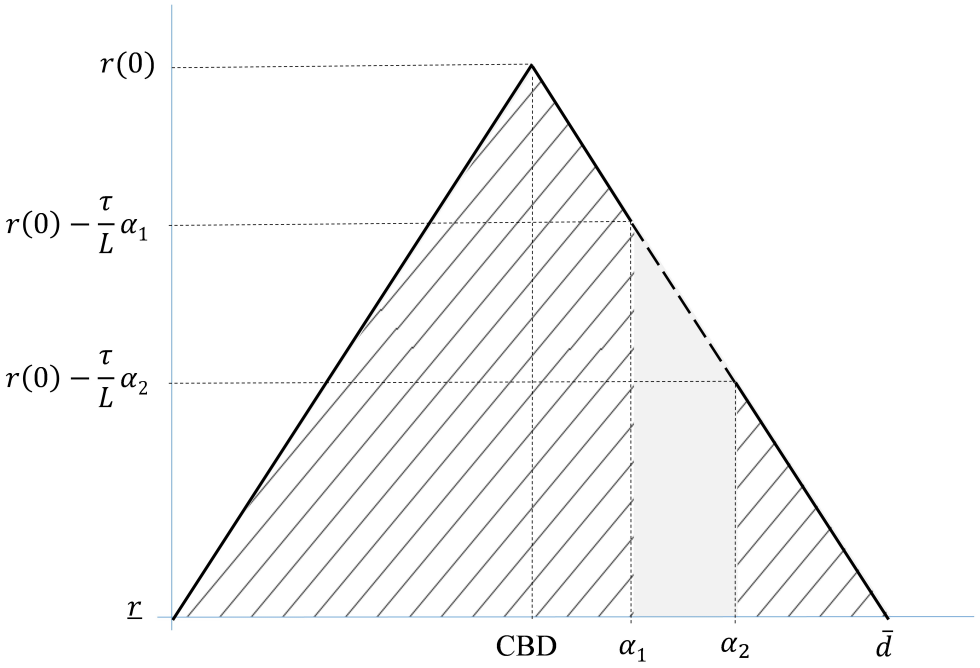

Figure B.2: Rents distribution in a linear city with a constraint

done by calculating the areas of the relevant triangles and rectangles in Figure B.2. Total rents in the unconstrained city, denoted as R_U, can be calculated as:

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

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

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

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 B.2, which we can denote as A:

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

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

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

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

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

Plugging in the expressions for N_U, R_U and A, after a few algebraic passages this inequality simplifies to:

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

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. This simple example illustrates that the relationship between irregular city layouts and average rents in a framework of spatial equilibrium within cities is theoretically ambiguous; whether rents are higher or lower in constrained cities is thus ultimately an empirical question.

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.

C. Nonparametric Employment Subcenter Identification (McMillen, 2001)

In order to compute the number of employment subcenters in each city, 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 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 km. While this is arbitrary, this approach is not particularly sensitive to the size of the unit considered. I calculate a proxy for employment density in each cell, by considering establishments located in that cell and summing their reported number of employees.[52]

[52]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.

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:

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

using locally weighted regression, employing a tricube kernel and a 50 percent 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:

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

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. (C.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 (C.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 percent level. The final list of subcenters includes the sites with positive coefficients on either $(D_j)^{-1}$ or $(-D_j)$.

References

[1] Alonso, W. (1964), Location and land use, Cambridge: Harvard University Press.Anas, A., R. Arnott, and K. A. Small (1998), "Urban Spatial Structure", Journal of Economic Literature, 36 (3), 1426-1464.

[2] 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.

[3] Brueckner, J. K. (1987), "The Structure of Urban Equilibria: A Unified Treatment of the Muth-Mills Model", in *The Handbook of Regional and Urban Economics*, 2, ed. E.S. Mills, Amsterdam: North Holland Press.

[4] 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.

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

[6] 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).

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

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

[9] 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.

[10] Mills, E.S. (1967) "An aggregative model of resource allocation in a metropolitan area", American Economic Review, 57, 197-210.

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

[12] 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.

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

[14] 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.

[15] 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.

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

Table A1: First Stage and Impact of City Shape on Population

Robustness to Excluding Cities with Extreme Topographies

	A. Excluding coastal cities				
	(1)	(2)	(3)	(4)	(5)
	FS	IV	FS(1)	FS(2)	IV
	Norm. shape of actual footprint	Population density	Shape of actual footprint, km	Log area of actual footprint, km^2	Log population
Norm. shape of potential footprint	0.0670***				
	(0.0249)				
Norm. shape of actual footprint		-241.6***			
		(76.98)			
Shape of potential footprint, km			1.352***	0.156***	
			(0.220)	(0.0462)	
Log projected historic population			-1.182***	0.277**	
			(0.260)	(0.117)	
Shape of actual footprint, km					-0.100**
					(0.0414)
Log area of actual footprint, km^2					0.776***
					(0.190)
Observations	5,917	1,266	5,917	5,917	1,266
Model for $\hat{f}$	common rate	common rate	city-specific	city-specific	city-specific
F stat shape	7.844	11.810	67.476	67.476	67.476
F stat area			13.031	13.031	13.031
Cities	410	410	410	410	410
City FE	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES

Notes: this table presents a robustness check to Tables 2 and 3, excluding cities with "extreme" topographies from the sample. Each observation is a city-year. Cols. (1) and (3), (4) are equivalent to cols. (4), (5), (6), panel C, in Table 2. They report OLS estimates of the first-stage relationship between city shape and area and the instruments discussed in Section 5.1. Cols. (2) and (5) are equivalent to cols. (4) and (5), panel C, in Table 3, and report IV estimates of the impact of shape on log city population. Shape is captured by the disconnection index, which measures the average length of trips within the city footprint, in km. The construction of the potential footprint is based on a city-specific model for city expansion - see Section 5.1. Angrist-Pischke F statistics for the shape and area variables are reported. Panel A excludes from the sample cities located within 5 km from the coast. Panel B excludes from the sample cities with an elevation above 600 m. City shape and area are calculated from maps constructed from the DMSP/OLS Night-time Lights dataset (1992-2010) and U.S. Army maps (1951). Population is drawn from the Census of India (1951, 1991, 2001, 2011). Elevation is from the ASTER dataset. All specifications include city and year fixed effects. Standard errors are clustered at the city level. *** p<0.01,** p<0.05,* p<0.1.

Table A1 (Continued): First Stage and Impact of City Shape on Population

Robustness to Excluding Cities with Extreme Topographies

	B. Excluding mountainous cities				
	(1)	(2)	(3)	(4)	(5)
	FS	IV	FS(1)	FS(2)	IV
	Norm. shape of actual footprint	Population density	Shape of actual footprint, km	Log area of actual footprint, km^2	Log population
Norm. shape of potential footprint	0.0662***				
	(0.0248)				
Norm. shape of actual footprint		-258.6***			
		(83.07)			
Shape of potential footprint, km			1.363***	0.159***	
			(0.225)	(0.0488)	
Log projected historic population			-1.193***	0.292**	
			(0.269)	(0.123)	
Shape of actual footprint, km					-0.109**
					(0.0428)
Log area of actual footprint, km^2					0.796***
					(0.185)
Observations	5,905	1,278	5,905	5,905	1,278
Model for f̂	common rate	common rate	city-specific	city-specific	city-specific
F stat shape	7.672	11.437	71.927	71.927	71.927
F stat area			13.899	13.899	13.899
Cities	414	414	414	414	414
City FE	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES

Notes: this table presents a robustness check to Tables 2 and 3, excluding cities with "extreme" topographies from the sample. Each observation is a city-year. Cols. (1) and (3), (4) are equivalent to cols. (4), (5), (6), panel C, in Table 2. They report OLS estimates of the first-stage relationship between city shape and area and the instruments discussed in Section 5.1. Cols. (2) and (5) are equivalent to cols. (4) and (5), panel C, in Table 3, and report IV estimates of the impact of shape on log city population. Shape is captured by the disconnection index, which measures the average length of trips within the city footprint, in km. The construction of the potential footprint is based on a city-specific model for city expansion - see Section 5.1. Angrist-Pischke F statistics for the shape and area variables are reported. Panel A excludes from the sample cities located within 5 km from the coast. Panel B excludes from the sample cities with an elevation above 600 m. City shape and area are calculated from maps constructed from the DMSP/OLS Night-time Lights dataset (1992-2010) and U.S. Army maps (1951). Population is drawn from the Census of India (1951, 1991, 2001, 2011). Elevation is from the ASTER dataset. All specifications include city and year fixed effects. Standard errors are clustered at the city level. *** p<0.01,** p<0.05,* p<0.1.

Table A2: Impact of Shape on Population, Robustness to Initial Shape x Year Fixed Effects

	(1) IV Population density	(2) IV Log population
Norm. shape of actual footprint	-348.4***	
	(118.0)	
Shape of actual footprint, km		-0.194***
		(0.0596)
Log area of actual footprint, km^2		0.973***
		(0.222)
Observations	1,329	1,329
F stat shape	9.115	26.699
Model for f̂	common rate	city-specific
City FE	YES	YES
Year FE	YES	YES
Initial shape x year FE	YES	YES

Notes: this table presents a robustness check to Table 3, augmenting the specification with year fixed effects interacted with initial shape. Each observation is a city-year. Cols. (1) and (2) are equivalent to cols. (4) and (5), panel C, in Table 3. Col. (1) reports IV estimates of the relationship between normalized city shape and population density, in thousand inhabitants per km^2. Col. (2) reports IV estimates of the relationship between city shape and area, and log population. Shape is captured by the disconnection index, which measures the average length of trips within the city footprint, in km. The construction of the potential footprint is based on a common rate model for city expansion in col. (1) and on a city-specific one in col. (2) – see Section 5.1. City shape and area are calculated from maps constructed from the DMSP/OLS Night-time Lights dataset (1992-2010) and U.S. Army maps (1951). Population is drawn from the Census of India (1951, 1991, 2001, 2011). All specifications include city and year fixed effects, as well as year fixed effects interacted with the city's disconnection index measured in the initial year of the panel (1951). Standard errors are clustered at the city level. *** p<0.01,** p<0.05,* p<0.1.

Table A3: Impact of Shape on Wages and Rents, Robustness to Initial Shape x Year Fixed Effects

	(1)	(2)	(3)	(4)
	IV	IV	IV	IV
	All districts		*Only districts with one city*	
	A. Dependent variable: log wage			
Shape of actual footprint, km	0.0989***	0.0116	0.108**	0.0393
	(0.0272)	(0.0420)	(0.0421)	(0.0768)
Log area of actual footprint, km^2		-0.358		-0.164
		(0.395)		(0.465)
F stat shape	6.689	11.480	4.158	5.518
Observations	1,943	1,943	1,021	1,021
	B. Dependent variable: log yearly rent per square meter			
Shape of actual footprint, km	-0.0305	-0.74	-0.0319	-0.669*
	(0.242)	(0.482)	(0.363)	(0.356)
Log area of actual footprint, km^2		-1.245		-0.562
		(1.348)		(0.802)
Observations	841	841	445	445
F stat shape	5.650	6.457	2.974	7.314
Model for $\hat{f}$	common rate	city-specific	common rate	city-specific
City FE	YES	YES	YES	YES
Year FE	YES	YES	YES	YES
Initial shape x year FE	YES	YES	YES	YES

Notes: this table presents a robustness check to Tables 5 and 6, augmenting the specifications with year fixed effects interacted with initial shape. Each observation is a city-year. In panel A the dependent variable is the log urban average of individual yearly wages in the city's district, in thousand 2014 Rupees. In panel B the dependent variable is the log urban average of housing rents per m^2 in the city's district, in 2014 Rupees. Cols. (1), (3) report the results from estimating equation (44) (single-instrument specification) by IV. Cols. (2), (4) report the results from estimating equation (39) (double-instrument specification) by IV. Instruments are described in Section 5.2. Angrist-Pischke F statistics for the shape variable are reported. The construction of the potential footprint is based on a common rate model for city expansion in cols. (1), (3), and on a city-specific one in cols. (2), (4) - see Section 5.1. In cols. (3) and (4), the sample is restricted to districts containing only one city. Shape is captured by the disconnection index, which measures the average length of trips within the city footprint, in km. City shape and area are calculated from maps constructed from the DMSP/OLS Night-time Lights dataset (1992-2010). Wages are from the Annual Survey of Industries, waves 1990, 1994, 1995, 1997, 1998, 2009, 2010. Housing rents are from the NSS Household Consumer Expenditure Survey, rounds 62 (2005-2006), 63 (2006-2007) and 64 (2007-2008). All specifications include city and year fixed effects, as well as year fixed effects interacted with the city's disconnection index measured in the initial year of the panel (1951). Standard errors are clustered at the city level. *** p<0.01,** p<0.05,* p<0.1.

Table A4: Impact of City Shape on Sectoral Shares

	(1) IV	(2) IV	(3) IV	(4) IV	(5) IV
	Social Services	Services	Transport and Storage	Retail	Manufacturing
Shape of actual footprint, km	0.0219	0.0681	0.111	0.0570	-0.0957
	(0.0504)	(0.103)	(0.100)	(0.0457)	(0.0674)
Observations	684	684	684	684	684
F stat shape	23.69	23.69	23.69	23.69	23.69
Mean of share	0.29	0.06	0.05	0.34	0.24
Model for $\hat{f}$	common rate	common rate	common rate	common rate	common rate
City FE	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES

Notes: This table investigates the relationship between city shape and the share of firms in different sectors, weighted by number of workers. Each observation is a city-year. The dependent variable is the log share of workers in different sectors, derived from the 3rd, 4th and 5th Economic Census. Sectors are defined based on National Industry Classification 1-digit codes: social services (1), services (2), transport and storage (3), retail (4), manufacturing (5). Sectors with negligible (<0.007) shares (electricity and gas, construction, mining) are not reported. "Services" sector includes: financial, insurance, real estate and business services. "Social Services" sector includes: community, social and personal services. See Section 3.5 for a description of the data. Cols. (1) to (5) report the results from estimating equation (44) (single-instrument specification) by IV. Instruments are described in Section 5.1. F statistics for the shape variable are reported. The construction of the potential footprint is based on a common rate model for city expansion. Shape is captured by the disconnection index, which measures the average length of trips within the city footprint, in km. City shape and area are calculated from maps constructed from the DMSP/OLS Night-time Lights dataset (1992-2010). All specifications include city and year fixed effects. Standard errors are clustered at the city level. *** p<0.01,** p<0.05,* p<0.1.

Table A5: Impact of City Shape and FARs on Population

	(1)	(2)	(3)	(4)	(5)	(6)
	IV	IV	OLS	IV	IV	OLS
	Population Density	Log population	Log population	Population Density	Log population	Log population
Norm. shape of actual footprint	-97.49			-85.46		
	(102.5)			(108.1)		
Norm. shape of actual footprint	-10.91			-16.89		
x FAR	(43.88)			(45.12)		
Shape of actual footprint, km		-0.0979	0.0509		-0.271**	0.0283
		(0.124)	(0.0312)		(0.122)	(0.0285)
Shape of actual footprint, km		0.00290	-0.0160		0.0688*	-0.00767
x FAR		(0.0472)	(0.0129)		(0.0371)	(0.0103)
Log area of actual footprint, km^2		0.683**	0.0109		0.828***	-0.0443
		(0.338)	(0.128)		(0.308)	(0.103)
Log area of actual footprint, km^2		0.0995	0.0665		0.00166	0.0981***
x FAR		(0.116)	(0.0471)		(0.0886)	(0.0368)
Observations	252	252	252	252	252	252
Model for $\hat{f}$	common rate	city-specific		common rate	city-specific	
FAR	average	average	average	residential	residential	residential
City FE	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES

Notes:This table reports estimates of the relationship between Floor Area Ratios, shape, and population. Each observation is a city-year. Cols. (1), (2), (3) and (4), (5), (6) report the same specifications reported in Table 3, panel C, cols. (4), (5), (6), with the addition of interactions between each explanatory variable and FARs. FARs are drawn from Sridhar (2010) and correspond to the maximum allowed Floor Area Ratios in each city 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. Cols. (1), (2), (3) consider the average of residential and non-residential FARs, while cols. (4), (5), (6) only consider residential FARs. Shape is captured by the disconnection index, which measures the average length of trips within the city footprint, in km. City shape and area are calculated from maps constructed from the DMSP/OLS Night-time Lights dataset (1992-2010) and U.S. Army maps (1951). Population is from the Census (1951, 1991, 2001, 2011). Population density is measured in thousands of inhabitants per km^2. All specifications include city and year fixed effects. Standard errors are clustered at the city level. *** p<0.01,** p<0.05,* p<0.1.

Table A6: FARs Determinants

	(1)	(2)	(3)	(4)	(5)	(6)
	IV	IV	OLS	IV	IV	OLS
	Avg. FAR			***Residential FAR***		
Shape of actual footprint, km	0.000290	0.0508	0.021	0.00515	0.0715*	0.0439**
	(0.00741)	(0.0409)	(0.0140)	(0.00914)	(0.0401)	(0.0199)
Log area of actual footprint, km^2		-0.190	-0.105*		-0.280	-0.177**
		(0.179)	(0.0540)		-0.176	(0.0876)
Observations	55	55	55	55	55	55
Model for f̂	common rate	city-specific		common rate	city-specific	

Notes: This table investigates the cross-sectional relationship between city shape, city area, and Floor Area Ratios as of year 2005. Each observation is a city in year 2005. Cols. (1) and (4) estimate a cross-sectional version of equation (44) (single-instrument specification), with log FARs as a dependent variable, estimated by IV. Cols. (2) and (5) estimate a cross-sectional version of equation (39) (double-instrument specification), with log FARs as a dependent variable, estimated by IV. Cols. (3) and (6) present the same specification, estimated by OLS. FARs are drawn from Sridhar (2010) and correspond to the maximum allowed Floor Area Ratios in each city 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. Cols. (1), (2), (3) consider the average of residential and non-residential FARs, while cols. (4), (5), (6) only consider residential FARs. Shape is captured by the disconnection index, which measures the average length of trips within the city footprint, in km. City shape and area are calculated from maps constructed from the DMSP/OLS Night-time Lights dataset, in year 2005. Standard errors are clustered at the city level. *** p<0.01,** p<0.05,* p<0.1.