Cities in Bad Shape: Urban Geometry in India*

Mariaflavia Harari†

September 2016

Abstract

The spatial layout of cities is an important determinant of urban commuting efficiency, previously highlighted by urban planners but overlooked by economists. This paper investigates the economic implications of urban geometry in the context of India. A satellite-derived dataset of night-time lights is combined with historic maps to retrieve the geometric properties of urban footprints in India over time. I propose an instrument for urban shape, which combines geography with a mechanical model for city expansion: in essence, cities are predicted to expand in circles of increasing sizes, and actual city shape is predicted by obstacles within this circle. With this instrument in hand, I investigate how city shape affects the location choices of consumers and firms, in a spatial equilibrium framework à la Roback-Rosen. Cities with more compact shapes are characterized by larger population, lower wages, and higher housing rents, consistent with compact shape being a kind of consumption amenity. The implied welfare cost of deteriorating city shape is estimated to be sizeable. I also attempt to shed light on policy responses to deteriorating shape. The adverse effects of unfavorable topography appear to be exacerbated by building height restrictions, and mitigated by road infrastructure.

*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, Tommaso Porzio, Champaka Rajagopal, Otis Reid, Albert Saiz, Chris Small, Kala Sridhar, Matthew Turner, Maisy Wong and seminar participants at MIT, NEUDC, UIUC, Columbia SIPA, LSE, Zurich, Wharton, the World Bank, Johns Hopkins (Carey), the Minneapolis FED, the IGC Cities Program, the NBER Summer Institute, the Meeting of the Urban Economics Association, NYU, CEMFI, UPF, Stockholm School of Economics, Stockholm University, the IEB Urban Economics Conference and the Barcelona Summer Forum for helpful comments and discussions.

†The Wharton School, University of Pennsylvania. harari@wharton.upenn.edu

1 Introduction

Most urban economics assumes implicitly that cities are circular. Real-world cities, however, often depart significantly from this assumption. Geographic or regulatory constraints prevent cities from expanding radially in all directions and can result in asymmetric or fragmented development patterns. While the economics literature has devoted very little attention to this particular feature of urban form, city shape as 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, have the potential to affect productivity and welfare (Bertaud, 2004; Cervero, 2001). This is particularly relevant for cities in the developing world, where most inhabitants cannot afford individual means of transportation (Bertaud, 2004). Cities in developing countries currently host 1.9 billion residents (around 74% 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 impact of urban geometry in the context of India, exploiting plausibly exogenous variation in city shape driven by geographic barriers. More specifically, 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. I investigate this in the framework of spatial equilibrium across cities à la Roback-Rosen. By examining the impact of city shape on population, wages, and housing rents, I attempt to quantify the loss from deteriorating geometry in a revealed preference setting.

As the country with the world's second largest urban population (UN, 2015), India represents a relevant setting for researching urban expansion. However, systematic data on Indian cities and their spatial structures is not readily available. In order to investigate these issues empirically, I assemble a 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 data sources. I trace the footprints of Indian cities at different points in time by combining newly geo-referenced historic maps (1950) with satellite imagery of night-time lights (1992-2010). The latter approach overcomes the lack of high-frequency land use data. For each city-year, I then compute quantitative indicators for urban geometry, used in urban planning as proxies for the patterns of within-city trips. 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. One of the contributions of this paper thus relates to the measurement of the properties of urban footprints over time.

A second 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 a given 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 use a mechanical model for city expansion to predict the area that a city should occupy in a given year; in its simplest version, such a model postulates a common growth rate for all cities. Second, I consider the largest contiguous patch of developable land within this predicted radius ("potential footprint") and compute its geometric properties. I then proceed to 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 document that city shape, a previously overlooked feature of urban form, can have substantial economic implications. I first investigate whether households and firms value compact city shape when making location choices across cities. Guided by a simple model of spatial equilibrium across cities, I examine the aggregate responses of population, wages and housing rents, measured at the city level, to changes in shape. My findings suggest that consumers value city compactness as a kind of "consumption amenity". 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, possibly, 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,[1] entails a loss equivalent to a 4% decrease in income. On the other hand, firms do not appear to require such a compensation: in equilibrium, compactness has negligible impacts on the productivity of firms. 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. These 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 patterns within cities. Evidence from the street addresses of firms suggests 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.

In the second part of the paper I consider the role of policy: if city shape indeed has welfare implications, what are possible policy responses to deteriorating shape? On the one hand, I consider infrastructural investment as a policy tool to counteract the effects of poor geometry. I find that

[1] As a reference, the average city in my sample has an area of 62.6 square km, and an average, one-way within-city trip of 3.3 km.

the negative effects of deteriorating geometry on population are mitigated by road infrastructure, supporting 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 a co-determinant of city shape. I find that more permissive vertical limits, 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.

The remainder of the paper is organized as follows. Section 2 provides some background on urbanization in India and reviews the existing literature. Section 3 documents my data sources and describes the geometric indicators I employ. Section 4 outlines the conceptual framework. Section 5 presents my empirical strategy and describes in detail how my instrument is constructed. The empirical evidence is presented in the following two Sections. Section 6 discusses my main results, which pertain to the implications of city shape for the spatial equilibrium across cities. Section 7 provides results on responses to city shape, including interactions between topography and policy. Section 8 concludes and discusses indications for future work.

2 Background and previous literature

India represents a promising setting to study urban spatial structures for a number of reasons. First, as most developing countries, India is experiencing fast urban growth. According to the 2011 Census, the urban population amounts to 377 million, increasing from 285 million in 2001 and 217 million in 1991, representing between 25 and 31 percent of the total. Although the pace of urbanization is slower than in other Asian countries, it is accelerating, and it is predicted that another 250 million will join the urban ranks by 2030 (Mc Kinsey, 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 and expand over time.

Secondly, unlike most other developing countries, India is characterized by a diffused urbanization pattern and an unusually large number of highly-populated cities. This lends itself to an econometric approach based on a city-year panel.

The challenges posed by rapid urban expansion on urban form and mobility have been gaining increasing importance in India's policy debate, which makes it particularly relevant to investigate these matters from an economics perspective. Limited urban mobility and lengthy commutes are often cited among the perceived harms of rapid urbanization (e.g., Mitric and Chatterton, 2005; World Bank, 2013), and providing effective urban public transit systems has been consistently identified as a key policy recommendation for the near future (Mc Kinsey, 2013). There is also a growing concern that existing land use regulations might directly or indirectly 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, 2002a;

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).

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 for transit are left mostly unexplored. A large 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 somewhat related to the notion of "compactness" that I investigate, in most analyses the focus is on decentralization and density, neglecting differences in city geometry. I focus on a different set of spatial properties of urban footprints: conditional on the overall amount of land used, I look at geometric properties aimed at proxying the pattern of within-city trips, and view density as an outcome variable. In this respect, my work is more closely 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 attempt to address the endogeneity of city shape in an instrumental variables framework.

The geometry of cities has attracted the attention of the quantitative geography and urban planning literature, from which I borrow indicators of city shape (Angel et al., 2009). Urban planners emphasize the link between city shape, average intra-urban trip length and accessibility, claiming that contiguous, compact and predominantly 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 views city structure as the outcome of fractal processes and 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 urban areas over time. More specifically, they analyze changes in the extent of sprawl in US cities between 1992 and 1996. The data I employ comes mostly from night-time, as opposed to day-time, imagery, and covers a longer time span (1992-2010). Saiz (2010) also looks at 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 contributes to a growing literature on road infrastructure and urban growth in developing countries (Baum-Snow and Turner, 2012; Baum-Snow et al., 2013; Morten and Oliveira, 2014; Storeygard, forthcoming). Differently from these studies, I do not look at the impact of roads connecting cities, but instead focus on trips within cities, as proxied by urban geometry.

Finally, the geometry of land parcels in a rural, as opposed to urban context, has received some attention in the law and economics literature. Libecap and Luek (2011) have explored the economic implications of different land demarcation regimes, showing that land values are higher when the system in place is one generating regular-shaped parcels.

3 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.

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. 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, such as that displayed in Figure 1A. I geo-referenced each of these sheets and manually traced the reported perimeter of urban areas, which are clearly demarcated (Figure 1B).

The second source is derived from the DMSP/OLS Night-time Lights dataset. This dataset is based on 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). Night-time lights have been employed in economics typically for purposes other than urban mapping (V. Henderson et al., 2012). However, the use of the DMSP-OLS dataset for delineating urban areas is quite common in urban remote sensing (M. Henderson et al., 2003; C. Small et al., 2005; C. 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, illustrated in Figure 2, 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.[2] 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.[3] In my full panel (including years 1950 and 1992-2010), the average city footprint occupies an area of approximately 63 square km.[4]

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 that allow us to detect changes in urban areas over time, due to their yearly temporal frequency. Finally, unlike multi-spectral satellite imagery such as Landsat- or MODIS- based products, which in principle would be available for different points in time, night-time lights do not require any sophisticated manual pre-processing.[5] 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., M. Henderson et al., 2003; C. 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.

It is well known that urban maps based on night-time lights will tend to inflate urban boundaries, due to "blooming" effects (C. Small et al., 2005).[6] 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

[2]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.

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

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

[5]Using multi-spectral imagery to map urban areas requires a manual classification process, which relies extensively on alternative sources, mostly aerial photographs, to cross-validate the spectral recognition, and is subject to human bias.

[6]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 (C. Small et al., 2005).

dataset (1992-2010), I thus assemble an unbalanced[7] panel of urban footprints. The criteria 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. Throughout my analysis, I instrument all the geometric properties of urban footprints, including both area and shape. This IV approach addresses issues of measurement error, which could affect my data sources - for instance due to the well-known correlation between income and luminosity.

3.2 Shape Metrics

The indicators of city shape that I employ (Angel et al., 2009a, 2009b),[8] are used in landscape ecology and urban studies to proxy for the length within-city trips and infer travel costs. They are based on the distribution of points around the polygon's centroid[9] or within the polygon, and are measured in kilometers. Summary statistics for the indicators below are reported in Table 1.

(i) The *remoteness* index is the average distance between all interior points and the centroid. It can be considered 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 centroid. This is similar to the remoteness index, but gives more weight to the polygon's extremities, corresponding to the periphery of the footprint. This index is more capable of identifying footprints that have "tendril-like" projections, often perceived as an indicator of sprawl.

(iii) The *disconnection* index captures the average distance between all pairs of interior points. It can be considered a proxy for commutes within the city, without restricting one's attention to those to or from to the center. As discussed below, 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 separate 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 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 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. Conditional on footprint area, higher values of these indexes indicate longer within-city trips.

[7]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.

[8]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.

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

Figure 3 provides a visual example of how these metrics map to the shape of urban footprints. Among cities with a population over one million, I consider those with respectively the "best" and the "worst" geometry based on the indicators described above, namely Bengaluru and Kolkata (formerly known as Bangalore and Calcutta). The figure reports the footprints of the two cities as of year 2005, where Bengaluru's footprint has been rescaled so that they have the same area. The figure also reports the above shape metrics computed for these two footprints. The difference in the remoteness index between Kolkata and (rescaled) Bengaluru is 4.5 km; the difference in the disconnection index is 6.2 km. The interpretation is the following: if Kolkata had the same compact shape that Bengaluru has, the average potential trip to the center would be shorter by 4.5 km and the average potential trip within the city would be shorter by 6.2 km.

It is worth emphasizing that these metrics should be viewed as proxies for the length of *potential* intra-urban trips as driven by the city's layout, and they abstract from the actual distribution of or 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's location choices within each city. As I discuss in Section 3.5, detailed data on actual commuting patterns and on the distribution of households and jobs within cities is, in general, very difficult to obtain for India. To have a rough sense of the mapping between city shape and realized commuting length, I draw upon recently released Census data on distance from residence to place to work, available at the district-urban level. The 2011 Census reports the number of urban workers in each district residing at different reported distances from their residence to their place of work, by coarse bins.[10] 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).[11]As expected, this correlation is positive.

3.3 Geography

For the purposes of constructing the instrument, I code geographic constraints to urban expansion as follows. Following Saiz (2010), I consider as "undevelopable" terrain that is either occupied by a water body, or characterized by a slope above 15%. 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.

[10]These figures are from table B28, "Other workers by distance from residence to place of work". The Census definition of "other workers" refers to those employed outside agriculture. The distance bins are 0-1 km, 2-5 km, 6-10 km, 11-20 km, 21-30 km, 31-50 or above 51 km.

[11]The matching between cities and districts is not one to one (see Section 3.4). The correlation reported above is robust to different approaches for matching cities to districts. If I exclude districts with more than one city or focus on the top city in each district, the correlation above becomes respectively 0.2 (p-value 0.005) or 0.204 (p-value 002). Results employing the remoteness index are very similar.

3.4 Outcome data: population, wages, rents

The main outcome variables that I consider are population, wages and housing rents, derived from a variety of sources.

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. As explained in Section 5.1, historic 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.[12] 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.[13]

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.

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

[12] Historic 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.

[13] In order to assemble a consistent panel of city population totals over the years I also take into account changes in the definitions of "urban agglomerations" and "outgrowths" across Census waves.

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.[14] This may affect the interpretation of my results, to the extent that this sector is systematically over- or underrepresented in cities with worse shapes. I 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. I construct a rough proxy for the rental price of housing drawing upon 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). 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, which is *a priori* less likely to include observations from rent-controlled housing.

3.5 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 (CIER)'s Industrial Databooks.

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. The average FAR in this sample is 2.3, a very restrictive figure compared to international standards. For a detailed discussion of FARs in India, see Sridhar (2010) and Bertaud and Brueckner (2005).

[14] 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.

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.[15]

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.[16] Although limited by their cross-sectional nature, these data provide an opportunity to study the spatial distribution of employment within cities, and in particular to investigate polycentricity.

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. Details can be found in the Appendix.

As this data description shows, retrieving and assembling together city-level data for Indian cities is not a straightforward exercise, and I face considerable data constraints. The main limitation is that, for most outcomes, only city-level averages can be observed, at very little information is available at a more disaggregated level. In particular, I do not observe population densities and location choices within cities,[17] nor actual commuting patterns, on which data has never been collected in a reliable and systematic way (see Mohan, 2013, for a discussion of the limited sources of data available and their severe limitations).

[15] The definition of sectors, identified by NIC codes, varies over Economic Census waves. I define sectors based on a coarse, 1-digit NIC code disaggregation so as to maintain consistency across waves.

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

[17] The Census collects ward-level population data for 2011 and 2001, but reports no information is reported on the location of wards themselves, which at the moment prevents an accurate spatial analysis. Unfortunately, I am also unable to infer within-city density patterns through the DMSP/OLS dataset, which does not appear to display enough variation in luminosity within urban areas.

4 Conceptual Framework

With these data constraints in mind, I frame the empirical question of the value of city shape drawing upon 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 levels of amenities. I hypothesize that households and firms 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. In this model, wages and housing rents equalize equilibrium utility across cities, striking the balance between the location preferences of consumers and firms. This modeling approach is attractive because it allows me to shed light on the economic value of city shape by observing the aggregate responses of population, wages and housing rents, measured at the city-year level, to changes in urban shape.

I follow the exposition of the model by Glaeser (2008). Households consume a composite good C and housing H. They supply inelastically one unit of labor receiving a city-specific wage W. Their utility depends on net income, i.e., labor income minus housing costs, and on a city-specific bundle of consumption amenities θ. 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:[18]

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

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 this condition is that consumers, in equilibrium, pay for amenities through lower wages (W) or through higher housing prices (p_h).[19]The extent to which wages net of housing costs rise with an amenity is a measure of the extent to which that amenity decreases utility, relative to

[18]The notion of spatial equilibrium across cities presumes that consumers are choosing across a number of different locations. The pattern of migration to urban areas observed in India is compatible with this element of choice: according to the 2001 Census, about 38 percent of rural to urban internal migrants move to a location outside their district of origin, supporting the interpretation that they are effectively choosing a city rather than simply moving to the closest available urban location.

[19]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.

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, 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 single internationally traded good Y, using labor N, and a local production amenity A. Their technology also requires traded capital K and a fixed supply of non-traded capital $\overline{Z}$.[20] 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 location there is a fixed supply of land $\overline{L}$, as a result of land use regulations.[21] Denoting with p_l the price of land, their maximization problem reads:

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

where

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

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

The construction sector operates optimally, with construction profits equalized across cities. By combining the housing supply equation, resulting from the developers' maximization problem, with the housing demand equation, resulting from the consumers' problem, we obtain the following

[20] This ensures constant returns to scale at the firm level and decreasing returns at the city level, which, in turn, is required in order to have a finite city size. This assumption could be dropped by assuming, for instance, decreasing returns in the production of housing (Glaeser, 2008).

[21] 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 quality of regulation, city growth, and geographic constraints. In my empirical analysis, when city area is explicitly controlled for, it is instrumented using historic population, thus abstracting from these issues (see Section 5.2, double-instrument specification).

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 (12), this model can be solved for the three unknowns N, W, and p_h, representing, respectively, population, wages, and housing prices, as functions of the model parameters, and in particular, 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:

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

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

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

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

Consider now an indicator of urban geometry S, higher values of S denoting "worse" shapes, in the sense of shapes conducive to longer commute trips. 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 are unaffected - possibly because of better access to transportation technology, or because of being centrally located within a city.[22] In this case we should observe the following reduced-form relationships:

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

A city with poorer shape should have, *ceteris paribus,* a smaller population, higher wages, and lower house rents. The intuition is that consumers prefer to live in 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 are borne by households as well as firms. For example, longer travel distances may increase a firm's transportation costs. 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. The reason for the ambiguous sign of $\frac{dW}{dS}$ is that now both firms and consumers want to locate in compact cities. With respect to the previous case, there now is an additional force that tends to bid wages up in compact cities: competition among firms for locating in low-S 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 we should observe $\frac{dW}{dS} > 0$.

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

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

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

Plugging (23) and (24) into (13), (14), and (15) yields

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

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

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

[22] Note that the model is agnostic about the specific channels through which shape impacts consumers and firms. Empirically, pinning down these channels directly would require more disaggregated data than what is available for India, and the empirical analysis will focus on aggregate city-level outcomes. Some evidence on mechanisms can be inferred from heterogeneous effects, and is discussed in Section 6.5 below.

where

$$\begin{aligned} B_N &:= \frac{(\delta(1-\alpha)+\alpha)\lambda_A + (1-\gamma)\delta\lambda_\theta}{\delta(1-\beta-\gamma)+\alpha\beta(\delta-1)},\ D_N := \frac{(1-\gamma)\alpha(\delta-1)}{\delta(1-\beta-\gamma)+\alpha\beta(\delta-1)} \quad (28)\\ B_W &:= \frac{(\delta-1)\alpha\lambda_A - (1-\beta-\gamma)\delta\lambda_\theta}{\delta(1-\beta-\gamma)+\alpha\beta(\delta-1)},\ D_W := \frac{-(1-\beta-\gamma)\alpha(\delta-1)}{\delta(1-\beta-\gamma)+\alpha\beta(\delta-1)} \quad (29)\\ B_P &:= \frac{(\delta-1)\lambda_A + (\delta-1)\beta\lambda_\theta}{\delta(1-\beta-\gamma)+\alpha\beta(\delta-1)},\ D_P := \frac{-(1-\beta-\gamma)(\delta-1)}{\delta(1-\beta-\gamma)+\alpha\beta(\delta-1)}. \quad (30)\end{aligned}$$

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

$$\begin{aligned} \lambda_A &= (1-\beta-\gamma)B_N + (1-\gamma)B_W \quad (31)\\ \lambda_\theta &= \alpha B_P - B_W. \quad (32)\end{aligned}$$

Parameter λ_θ captures, in log points, the loss from a marginal increase in S. Parameter λ_A captures the impact of a marginal increase in S on city-specific productivity. 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 β, γ, α, can be used to back out λ_A and λ_θ:

$$\begin{aligned} \widehat{\lambda_A} &= (1-\beta-\gamma)\widehat{B_N} + (1-\gamma)\widehat{B_W} \quad (33)\\ \widehat{\lambda_\theta} &= \alpha\widehat{B_P} - \widehat{B_W}. \quad (34)\end{aligned}$$

This approach captures the overall, net effect of S, in equilibrium, on the marginal city dweller, without explicitly modeling the mechanism through which S enters the decisions of consumers or firms. In Section 6.5, I provide empirical evidence suggesting that the urban transit channel is indeed involved, and in the concluding Section, I briefly discuss some alternative, second-order channels through which city shape might affect consumers.

This simple model makes a number of simplifying assumptions, that I discuss below. First, it does not explicitly address heterogeneity across consumers in tastes and skills. However, we expect that people will sort themselves into locations based on their preferences. The estimated differences in wages and rents across cities will thus be 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. While a richer model would allow to capture these important nuances and interactions, the scope of my empirical analysis is limited by the lack of disaggregated data. [23]

Second, the model could be extended to allow for congestion or agglomeration in consumption (or production). The utility of consumers, and the production function of firms could be augmented with a term that depends on city size N. In particular, in the presence of congestion externalities in consumption, the indirect utility of consumers will depend on city size as well. If shape is a

[23] Some indirect evidence of sorting is discussed in Section 6.5, in which I examine slum populations across cities with different geometries.

kind of consumption amenity, it will affect the utility of consumers both directly, through λ_θ, and indirectly, through its effect on city size, captured by $\frac{\partial \log(N)}{\partial S}$. If more compact cities have have larger populations, they will also be more congested; this congestion effect, in equilibrium, will tend to reduce the positive impact of compact shape of utility. The implication would be that, when I estimate the consumption amenity value of compact shape using equation (34), I would be capturing the equilibrium effect of shape, gross of congestion. In sum, if compact shape is a consumption amenity and compact cities are larger, then $\widehat{\lambda_\theta}$ 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 $\frac{\partial \log(N)}{\partial S}$. If compact cities have larger populations, this will tend to make them more productive through agglomeration; this effect will amplify the direct productivity impact of compactness. In this case, my estimate of the production amenity value of compact shape, obtained from equation (33), will be an *upper* bound for λ_A. Unfortunately, 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.

Reduced-form estimates for B_N, B_W, B_P are presented in Sections 6.2 and 6.3, whereas Section 6.4 provides estimates for parameters $\lambda_A, \lambda_\theta$. The next two Sections present the data sources and empirical strategy employed in the estimation.

5 Empirical Strategy

The conceptual framework outlined above suggests that the aggregate, city-level responses of population, wages and housing rents to city shape are informative of whether consumers and firms value city compactness as a production / consumption amenity in equilibrium. 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 that S is an indicator for city shape. The specification of interest is:

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

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 geographic conditions, city growth, and deliberate policy choices concerning land use and infrastructure. Urban shape is affected by land use regulations both directly, through master plans, and indirectly - for instance, land use regulations can encourage land consolidation, resulting in a more compact, as opposed to fragmented development pattern. Similarly, investments in road infrastructure can encourage urban growth along transport corridors, generating distinctive geometric patterns of development. Such policy choices are likely to be jointly determined with the outcome variables at hand. To see how this could bias my estimate 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 into 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 - compatible with compactness being an amenity. 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, biasing the estimates in the opposite direction. 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 into 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.[24] Intuitively, my identification relies on plausibly exogenous changes in shape that a given city undergoes over time, as a result of encountering topographic obstacles along its expansion path. More specifically, I construct an instrument for city shape that isolates the variation in urban shape driven by topographic obstacles 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, while Section 5.3 discusses in more depth the identification strategy and possilble threats to identification.

5.1 Instrumental Variable Construction

My instrument is constructed combining geography with a mechanical model for city expansion in time. The underlying idea is 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., 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

[24] As discussed below, a subset of the outcomes analyzed in Section 7 are available only for a cross-section of cities, in which case the comparison is simply across cities.

the shape properties of the *potential* footprint and use this as an instrument 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[25] (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 historic footprint, and construct corresponding potential footprints lying within these predicted radii (Figures 5c and 5d).

To complete the description of the instrument, I need to specify how $\widehat{r}_{c,t}$ is determined. 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: In this first version of the model for city expansion, I postulate that the rate of expansion of $\widehat{r}_{c,t}$ varies across cities, depending on their historic (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 - not projected - 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:

$$\log(area_{c,t}) = \alpha \cdot \log(\widehat{pop_{c,t}}) + \beta \cdot \log\left(\frac{pop_{c,1950}}{area_{c,1950}}\right) + \gamma_t + \varepsilon_{c,t} \tag{36}$$

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

[25] The US Army 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.

same as in 1951, and if the city could expand 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 $\widehat{r_{c,t}} = \sqrt{\frac{\widehat{area_{c,t}}}{\pi}}$.

As I will discuss below, 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 historic projected population, a variable that may be correlated with present-day outcomes.

5.2 Estimating Equations

Consider a generic shape metric S - which could be any of the indexes discussed in Section 3.2. 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

Consider outcome variable $Y \in (N, W, p_H)$ and let $area_{c,t}$ be the area of the urban footprint. The empirical counterparts of equations (25) – (27), augmented with city and year fixed effects, take the following form:

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

This equation contains two endogenous regressors: $S_{c,t}$ and $\log(area_{c,t})$. These are instrumented using respectively $\widetilde{S_{c,t}}$ and $\log(\widehat{pop_{c,t}})$ - the same projected historic 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 historic 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 of this strategy is that it requires not only an instrument for shape, but also one for area. Moreover, there is a concern that historic population might be correlated with current outcomes, leading to possible violations of the exclusion restrictions. This motivates me to complement this specification with an alternative, more parsimonious one, that does not explicitly include city area in the regression, and therefore does not require including projected historic population among the instruments.

Single-Instrument Specification

When focusing on population as an outcome variable, a natural way to do this is to normalize both the dependent and independent variables by city area, considering respectively the normalized shape metric - see Section 3.2 - and population density. I thus estimate the following, more parsimonious single-instrument specification: define population density as[26]

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

and denote the normalized version of shape metric S with nS. I then estimate

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

which contains 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. Although it does not allow the effects of shape and area to be separately identified, this approach is less demanding. In particular, it does not require using projected historic 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

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

seem as natural to normalize factor prices - wages and rents - by city area. For these other 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 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 employ robust standard errors clustered at the city level, to account for arbitrary serial correlation over time in cities.

5.3 Discussion

Let us now take a step back and reconsider the identification strategy as a whole.

As highlighted in the introduction to Section 5, the challenge in the estimation of the effects of city shape is that the latter is jointly determined with the outcomes of interest. One of the reasons why it may be the case is that city shape is partly the result of deliberate policy choices. My instrument addresses this, insofar as it is based on the variation in city shape induced by geography and mechanically predicted city growth, excluding, by construction, the variation resulting from policy choices.

Another reason for simultaneity is due to unobserved factors that may be systematically correlated both with city shape and with the outcomes of interest. My identification strategy helps address this concern in two ways. On the one hand, city fixed effects control for time-invariant city characteristics - for example, the fact of being a coastal city, or a state capital. On the other hand, city shape is instrumented using a time-varying function of the "potential footprint"'s geometry, that is arguably orthogonal to most time-varying confounding factors - such as rule of law, or changes in local politics.

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. These threats are discussed below.

A possible channel is the inherent amenity or disamenity value of geography. The topographic constraints that affect city shape, such as coasts and slopes, may also make cities more or less attractive for households and/or firms. This concern is mitigated by two features of my identification strategy. First, city fixed effects control for the time-invariant effects of geography. The inherent amenity value of geography thus poses a threat to identification only to the extent that this value is time varying. An example of this would be if coasts have a productive amenity value that declines over time. Second, the construction of the instrument relies on 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 not the generic presence nor on the magnitude of topographic constraints. Rather, the instrument captures the geometry of developable terrain, once topographic obstacles are excluded. Changes in this geometry over time are dictated primarily by the relative position of newly encountered topographic obstacles relative to previously encountered ones. 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 illustrated in one of the robustness checks discusses in Section 6.2, in which I show that my results are unchanged when I exclude mountain and coastal cities. The latter would be the two most obvious examples of cities where geography could have a specific (dis)amenity value.

It is nevertheless important to discuss the direction of the bias, if indeed the instrument were capturing some inherent amenity value of geography. Suppose that topographic obstacles that make cities less compact are also inherent consumption amenities – as it may be the case for coasts or lakes. This would bias my results *against* finding a consumption amenity value of compact shape. This scenario does not seem particularly plausible in a developing country setting, given that landscape amenities are likely to be a luxury good. On the other hand, suppose that topographic obstacles making cities less compact are also making cities more productive - for example, because they lead to waterway configurations that are favorable to trade (Bleakley and Lin, 2012). This would lead me to *underestimate* the productive amenity value of compact shape.

Another way in which geographic obstacles may directly affect outcomes such as population, housing rents or wages is by limiting the availability of developable land and hence affecting 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. *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. Rather, it is based on the relative position of individual topographic constraints. Again, the robustness check in Table 4 substantiates empirically this point, by showing that the results are not driven by cities that are particularly land constrained due to being coastal or mountainous. Moreover, the data indicates that Indian cities are particularly land-constrained by topographic obstacles. When I examine empirically the relationship between a city's area and the shape of the "potential footprint" (Table 2, columns 3 and 6), I find that, all 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 don't seem to prevent development, but rather induce cities to grow into less compact and potentially more land consuming shapes. If "bad shapes", as instrumented by topography, were nevertheless associated with land scarcity and 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 disamenity value of bad shape. As discussed below, 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.

Another concern is that cities with different "potential footprints" may be on different trends for the outcomes considered. To address this, I conduct a robustness check in which I augment the specifications above with year fixed effects interacted with the city's initial shape at the beginning of the panel (Appendix Tables 1 and 2). This more conservative specification allows cities that have different initial geometries to be on different trends.

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. The predictions of the conceptual framework suggest that, if city shape is valued as a kind of consumption amenity by consumers, cities with longer trip patterns should be characterized by lower population, higher wages and lower rents.

6.1 First Stage

[Insert Table 2]

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.[27] 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.[28] Higher values of these indexes represent less compact shapes. Summary statistics are reported in Table 1. Columns 1 and 4 report the first-stage for normalized shape (equation (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,

[27] Table 2 reports first stage estimates from the full sample. Most outcome variables considered in the subsequent analysis are observed in a subsample of cities and years, resulting in a smaller sample.

[28] 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.

the construction of the potential footprint is based on the common rate model for city expansion, outlined in Section 5.1.[29] Columns 2, 3 and 5, 6 report the first stage estimates for footprint shape (equation (40)) and area (equation (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 historic 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.

Let us consider first Table 2, panel A, which focuses on the remoteness index. As discussed in Section 3.2, this index captures the length of the average trip to the footprint's centroid, and can be considered a proxy for the average commute to the CBD. The remoteness of the *potential* footprint is a highly significant predictor of the remoteness index computed for the *actual* footprint, both in the normalized (column 1) and non-normalized version (column 2). Similarly, in column 3, projected historic population predicts footprint area. As expected, the city-specific, double-instrument specification is a 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. Column 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

[Insert Table 3]

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.

Columns 1 and 4 report the IV results from estimating the single-instrument specification (equation (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 columns 1 and 4 of Table 2. Columns 2 and 5 report the IV results from estimating the double-instrument specification (equation (39), which links population to city area and shape, separately instrumented for. The corresponding first stage estimates are reported respectively in columns 2, 3 and 5, 6 of Table 2. Columns 3 and 6 report the corresponding OLS estimates.

[29]The normalized shape instrument can, in principle, be constructed also using the city-specific model for urban expansion (see Section 5.1). Results of the corresponding first-stage are not reported in the table for brevity, but are qualitatively similar to those in column 3 and are available upon request.

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. The magnitudes for this specification are best understood in standardized terms. A one-standard deviation increase in normalized remoteness (0.06) is associated with a decline in population density of almost one standard deviation.

Interestingly, the OLS relationship between population and shape, conditional on area (column 3) 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. The intuition for this is the following: 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 (column 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 that a one-standard deviation increase in normalized remoteness (0.06), 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% 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, column 1) is driven both by a decrease in population (Table 3, column 2) and by an increase in footprint area (Table 2, column 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 those which more closely proxy 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 restring 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 patters in polycentric as well as monocentric cities. Unless otherwise specified, in the rest of the tables "shape" will indicate the disconnection index.

[Insert Table 4]

As a robustness check, in Table 4, 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 % 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.[30] Both the first-stage (columns 1, 3, and 4) and the IV estimates of the effect of shape on population density (column 2) and population (column 5) are minimally affected by excluding these cities. This shows that the instrument has explanatory power also in cities without extreme topographic constraints,[31] and that my IV results are not driven by a very specific subset of compliers.

Another robustness check is provided in Appendix Table 1. I re-estimate the IV impact of shape on density and population (columns 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.[32]A question then arises as to whether we can put a price on "good shape". As discussed in Section 4, the Rosen-Roback model provides a framework for doing so, by showing how urban amenities are capitalized in wages and rents. In particular, the model predicts that cities with better consumption amenities should be characterized by higher rents and lower wages.

Results on wages and rents are reported in Tables 5 and 6 respectively. As discussed in Section 3.4, the measures of wages and rents that I employ are subject to significant measurement error. Both are urban district-level averages, derived respectively from the Annual Survey of Industries and the National Sample Survey Consumer Expenditure Schedule. The matching between cities and districts is not one-to-one. In particular, there are numerous instances of districts that include more than one city. For robustness, in tables 5 and 6 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).

[Insert Table 5]

[30]In the panel fixed effects estimation singleton observations are dropped - this is reflected in the sample size reported in the tables.

[31]Recall that my instrument - the shape of the "potential" footprint - is not based on the severity of topographic constraints nor on the total share of land lost to such constraints, but is mostly driven by the relative *position* of constrained pixels.

[32]In the framework of spatial equilibrium across cities, differences in population growth rates across cities with different geometries are interpreted as the result of utility-equalizing migration. In principle, these observed differences could also be driven by differences in organic growth across cities. However, it is unclear how city shape would affect fertility or mortality rates. Exploring possible relationships between city shape and health is left for future research.

In Table 5, 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. Columns 1 and 4 report the IV results from estimating the single-instrument specification (equation (44)), that does not explicitly control for city area. Columns 2 and 5 report the IV results from estimating the double-instrument specification (equation (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 columns 1, 4, and on the city-specific one in columns 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 2, 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 (condition (17) in Section 4).

This "amenity" interpretation is subject to an important 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., 2007). For instance, low-income, low-skill workers could be disproportionately locating in cities with more compact shapes, that are also friendlier to commuters with limited individual transport options.[33] In this case, the finding that compact cities have lower wages may partly reflect systematic differences is workers' productivity, rather than "amenity" effects. These concerns could be alleviated controlling for workers' characteristics, which unfortunately are not available in the ASI data. [34]

[Insert Table 6]

Tables 6 reports the same set of specifications for house 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 averaging only the upper half of the distribution of urban housing rents in each district. This addresses the concern that reported rents are a downward-biased estimate of market rents due

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

[34]While I cannot observe workers' skills directly, in Appendix Table 4 I attempt to test whether the industry mix changes in response to changes in city shape, using data on urban establishments from the Economic Census (see Section 3.5). More specifically, I examine the relationship between city shape and the shares of workers in different sectors, as captured by coarse National Industry Code definitions. These results are not conclusive due to high standard errors, but, in general, do not support the interpretation that compact cities attract particular sectors.

to rent control policies. These estimates appear noisy or only borderline significant, with p values between 0.10 and 0.15. However, a consistent pattern emerges: the impact of disconnected shape on rents is negative in the IV and close to zero, or possibly positive, in the OLS. Appendix Table 2, 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 conceptual framework (condition (18) in Section 4).

6.4 Interpreting Estimates through the Lens of the Model

Tables 3, 5, and 6 provide estimates for the reduced-form relationship between city shape and, respectively, log population, wages, and rents, conditional on city area. Although results for wages and rents should be interpreted with caution due to the data limitations discussed above, the signs of the estimated coefficients are consistent with the interpretation that consumers view compact shape as an amenity. In this sub-section, I use these reduced-form estimates to back out the implied welfare loss associated with poor city geometry, according to the model outlined in Section 3. I focus here on the disconnection index, representing the average potential commute within the footprint. All monetary values are expressed in 2014 Rupees.

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 (equation (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 5 and 6A 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. Recall that a one standard deviation increase in disconnection for the average-sized city is 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 decrease in income.

In order to evaluate this magnitude, it would be interesting to compare this figure to estimates of the value of other amenities. Unfortunately, however, the literature on urban amenities in developing countries is limited and I am not aware of estimates available for India. As a reference, however, one could compare this figure with the *actual* cost of an extra 360 meters in one's daily one-way commute. Postulating one round-trip per day (720 meters), 5 days per week, this amounts to 225 extra km per year. 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. 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% of the yearly wage. This figure is roughly 53% of the welfare cost I estimate. Assuming instead that trips occur by car, postulating a speed of 25 km per hour, a fuel efficiency of 5 liters per 100 km, and fuel prices of 77 Rs. per liter,[35] the direct cost of increased commute length amounts to 1.3% of the yearly wage, or roughly 30% of the welfare cost estimated above.

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.[36] 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 would be their location choices within cities, both in terms of residence and employment. In the empirics, lack of systematic data on commuting and on residential patterns within cities prevents a detailed investigation of these patterns for consumers.

To complete the exercise, let us now 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 equation (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 equation (7). My most conservative point estimates for $\widehat{B_N}$ and $\widehat{B_W}$ are -0.099 and 0.038, from Tables 3C and 5 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.

[35]These figures are based respectively on: Ministry of Urban Development, 2008; U.S. Energy Information Administration, 2014; http://www.shell.com/ind/products-services/on-the-road/fuels/fuel-pricing.html accessed in August 2014.

[36]The behavioral literature has come to similar conclusions, 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.

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 times are indeed the main channel through which urban shape matters, road infrastructure should mitigate the adverse effects of poor geometry. By the same argument, all else being equal, consumers with individual means of transport should be less affected by city shape.

[Insert Table 7]

In Table 7, I attempt to investigate these issues interacting city shape with a number of indicators for infrastructure. For ease of interpretation, I focus on the single-instrument specification (equation (42)), which links normalized shape to population density, measured in thousand inhabitants per square km. For brevity, I report only IV estimates, using both the common rate and the city-specific model for city expansion. The shape indicators considered are the disconnection (Panel A) and range index (Panel B). Recall that these two indexes represent, respectively, the average and maximum length of within-city trips. While disconnection is a general indicator for city shape, the range index 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.[37] I partly address the endogeneity of city infrastructure by employing state-level proxies. Another concern is that the effect of infrastructure might be confounded by differential trends across cities with different incomes: cities that start out with 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 for a subset of cities.

In columns 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 columns 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 (columns 1, 2, and 3). Overall, this can be interpreted as suggestive evidence that infrastructure mitigates the negative effects of poor

[37] Lack of reliable, time-varying data on the road network of Indian cities make it difficult to investigate the relationship between city shape and the properties of a city's road network. I attempt to investigate the cross-sectional correlation between instrumented city shape and current road density drawing upon Openstreetmap data, but results do not reveal statistically significant patterns, possibly due to the low quality of the roads data.

geometry. Estimates are qualitatively similar when I include year fixed effects interacted with number of banks, as a time-varying proxy for city income. This suggests that the interaction terms are capturing indeed the role of infrastructure, and not simply the effect of being a higher-income city to start with.

Slums

[Insert Table 8]

A complementary question relates to who bears the costs of "bad" shape. Emphasizing the link between city shape, transit, and poverty, Bertaud (2004) claims that compact cities are, in principle, more favorable to the poor because they reduce distance, particularly in countries where they cannot afford individual means of transportation or where the large size of the city precludes walking as a means of getting to jobs. At the same time, however, if compact cities are also more expensive, this would tend to reduce the housing floor space that the poor can afford (Bertaud, 2004), potentially pricing low-income households out of the formal market. In Table 8, I investigate the link between city shape and slum prevalence using data from the more recent Census waves. For a limited number of cities and years, the Census provides information on slum population totals per city.[38] The dependent variables are log slum population (columns 1, 2 and 3) and log share of slum population (columns 4, 5 and 6) in a given city-year. As in previous tables, I provide IV estimates obtained with both the single- and the double-instrument specification, as well as OLS estimates for the double-instrument specification. 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 individual means of transport and consequent higher sensitivity to commute lengths.[39] This sorting interpretation would suggest that poorer households bear a disproportionate share of the costs of "bad" city geometry.

7 Empirical Results: Endogenous Responses to City Shape

In this Section I examine private and regulatory responses to deteriorating city geometry.

[38]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 (Census of India, 2011).

[39]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. This, of course, does not address the more general concern of sorting effects among the formally employed, an issue that is difficult to tackle without observing worker's characteristics.

7.1 Polycentricity

After having considered possible policy responses to deteriorating shape, in this Section I consider a private kind of response: firms' location choices within cities. As cities grow into larger and more disconnected footprints, resulting in lengthy commutes to the historic core, businesses might choose to locate further apart from each other, and/or possibly form new business districts elsewhere.[40] I attempt to shed light on this hypothesis by analyzing the spatial distribution of productive establishments listed in the Urban Directories of the 2005 Economic Census. This data source is described in greater detail in Section 3.5. The literature has proposed a number of methodologies to detect employment sub-centers within cities. I employ the two-stage, non-parametric approach developed by McMillen (2001), detailed in the Appendix. This procedure 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. 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. I compute the number of employment subcenters for each city in year 2005. This figure ranges from 1, for cities that appear to be purely monocentric, to 9, in large cities such as Delhi and Mumbai.[41] It should be noted that the sample size is quite small, due to inconsistencies in reported addresses and difficulty in geocoding them.

[Insert Table 9]

In Table 9, I estimate the relationship between number of employment centers, city area, and shape, in a cross-section of footprints observed in 2005. Column 1 reports the IV results from estimating a cross-sectional version of the single-instrument specification (equation (42)). The dependent variable is the number of subcenters per square km. Column 2 reports the IV results from estimating a cross-sectional version of the double-instrument specification (equation (39)). The dependent variable is the log number of employment subcenters. Column 3 reports the same specification, estimated by OLS.

Due to the small sample size and the fact that only cross-sectional variation is exploited, the results are quite noisy and suffer from weak instruments. Nevertheless, subject to these limitations, a number of interesting qualitative patterns are detected. Consistent with most theories of endogenous subcenter formation, and with the results obtained in the US context by McMillen and Smith (2003), larger cities tend to be have more employment subcenters (column 2). Interestingly, conditional on city area, less compact cities do not appear to be more polycentric: if anything, longer potential trips reduce the number of subcenters. Although these results are to be taken cautiously, they lend suggestive support to the following interpretation: as cities grow into more disconnected shapes,

[40] The literature on polycentricity and endogenous subcenter formation is reviewed, amongst others, by Anas et al. (1998), McMillen and Smith (2003), and Agarwal et al. (2012). Such models emphasize the firms' trade-off between a centripetal agglomeration force and the lower wages that accompany shorter commutes in peripheral locations.

[41] For the purposes of implementing the subcenter detection procedure, the CBD is defined as the centroid of the 1950 footprint. Results are robust to defining the CBD as the current centroid (available upon request).

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 for consumers, but has negligible impacts on firms (Section 6.4). If employment were as dispersed as population, commute trips 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 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 city shape affects the spatial equilibrium across cities, and, in particular, that deteriorating urban geometry entails welfare losses for consumers. The next question concerns the role of policy: given that most cities cannot expand radially due to their topographies, 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 (FAR).

As explained in Section 3.5, 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. Previous work has linked the restrictive FARs in place 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 restrictive FARs in Bengaluru carry a welfare cost ranging between 1.5 and 4.5%.

Information on FAR values across Indian cities is very hard to obtain. My data on FARs is drawn from Sridhar (2010), who collects a cross-section of the maximum allowed FAR levels as of the mid-2000s, for about 50 cities,[42] disaggregating by residential and non-residential FAR. Based on discussions with urban local bodies and developers, it appears that FARs are very resilient, and have rarely been updated. While the data collected by Sridhar reflects FARs as they were in the mid-2000s, they are likely to be a reasonable proxy for FAR values in place throughout the sample period.

[Insert Table 10]

In Table 10, panel A, I explore the interaction between topography and FARs in determining city shape and area. The three first-stage equations presented in Table 2 are reproposed here, augmented with an interaction between each instrument and FAR levels. Each observation is a

[42] Sridhar (2010) collects data for about 100 cities, but many of those cities are part of larger urban agglomerations, and do not have appear as individual footprints in my panel. Moreover, some are too small to be detected by night-time lights. This reduces the effective number of observations which I can use in my panel to 55. My analysis is thus subject to significant power limitations.

city-year. In columns 1, 2, and 3, I consider the average of residential and non-residential FARs, whereas in columns 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 columns 1 and 4 is negative, and significant in column 1, indicating that cities with higher FARs have a more compact shape than their topography would predict. The interaction between projected population and FARs appears to have a negative impact on city area (columns 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 (columns 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.

In Table 10, panel B, I investigate the impact of FARs interacted with city shape on population and density. Again, each observation is a city-year. The specifications proposed here are equivalent to those in Table 3, augmented with interactions between the explanatory variables and FARs. The corresponding interacted first-stage equations are proposed in panel A. Results are mixed, possibly due to small sample size. However, the results from the interacted version of the double-instrument specification (columns 2 and 5) suggest that laxer FARs mitigate the negative impact of non-compactness on population: the interaction between instrumented shape and FARs is positive, and significant in the case of residential FARs. An interpretation for this result is that long potential commutes matter less in cites which allow taller buildings, since this allows more consumers to live in central locations. This result, however, is not confirmed in the single-instrument specification (columns 1 and 4).

While these regressions take FARs as given, a question might arise on the determinants of FARs - in particular, whether urban form considerations appear to be incorporated by policy makers in setting FARs. Appendix Table 3 reports cross-sectional regressions of FAR values on urban form indicators - shape and area - as measured in year 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, or may simply reflect a historical legacy of taller architecture in more constrained cities.[43]

[43]Perhaps surprisingly, there is no evidence of FARs being driven by the geology of cities, including their earthquake proneness. Results are available upon request.

8 Conclusion

In this paper, I examine the economic implications of city shape in the context of India, exploiting variation in urban form driven by topography. I find that an urban layout conducive to longer within-city trips has sizeable welfare costs, and that city compactness affects the spatial equilibrium across cities. Urban mobility thus impacts not only the quality of life in cities, but also influences rural to urban migration patterns, by affecting city choice.

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, the policy debate has focused on the perceived harms of haphazard urban expansion, including limited urban mobility and lengthy commutes (World Bank, 2013). On the other hand, existing policies, especially land use regulations, have been indicated as a potential source of significant distortions in urban form (Glaeser, 2011; Sridhar, 2010; World Bank, 2013). My findings can 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 wide range of policy options to improve urban mobility and prevent the deterioration in connectivity that fast city growth entails. Urban mobility can be enhanced through direct interventions in the transportation sector, such as investments in infrastructure and public transit. Indeed, I find evidence that road infrastructure might mitigate the impact of disconnected city shape. My study also suggests that urban connectivity can be indirectly improved through another channel: promoting more compact development. This can be encouraged through master plans and land use regulations. Bertaud (2002a) reviews a number of urban planning practices and land use regulations, currently in place in Indian cities, that tend to "push" urban development towards the periphery.[44] I find that restrictive Floor Area Ratios, the most controversial of such regulations, result in less compact footprints, suggesting that city shape can indeed be affected by regulation and is not purely driven by geography. 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 such as FARs.

This paper leaves a number of questions open for future research, to be addressed as data availability improves. First, it would be important to understand the implications of geometry for the spatial equilibrium *within* cities, in terms of commuting and location choices. Second, it would be interesting to uncover heterogeneous effects of city shape, and gain a deeper understanding of who bears the costs of disconnected geometry. Finally, there is a range of potential mechanisms through which city shape affects consumers, each of which could be addressed individually. Throughout this paper, I employ shape metrics specifically constructed to capture the implications of shape for transit. However, there could be other channels through which city shape matters from an economic

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

standpoint. As noted by Bertaud (2002b), geometry affects not only transportation but all kinds of urban utilities delivered through spatial networks, including those collecting and distributing electricity, water, and sewerage. Moreover, a given urban layout may promote the separation of a city in different, disconnected neighborhoods and/or administrative units. This 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, C.L. Redfearn (2012), "Strangers in our midst: the usefulness of exploring polycentricity", Annals of Regional Science, 48 (2), 433-450.

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

[3] 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.

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

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

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

[7] 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, Q. Zhang (2013), "Roads, Railroads and Decentralization of Chinese Cities", working paper.

[9] Baum-Snow, N. and M. Turner (2012), "Transportation and the Decentralization of Chinese Cities", working paper.

[10] 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.

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

[12] Bertaud, A. (2002b), "Note on Transportation and Urban Spatial Structure",Washington, ABCDE conference, working paper.

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

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

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

[16] 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.

[17] 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.

[18] Carroll, M., J. Townshend, C. DiMiceli, P. Noojipady, R. Sohlberg (2009), "A New Global Raster Water Mask at 250 Meter Resolution", International Journal of Digital Earth, 2(4).

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

[20] Cervero, R. (2013), "Linking urban transport and land use in developing countries", The Journal of Transport and Land Use, 6 (1), 7-24.

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

[22] 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).

[23] Combes, P-P., Duranton, G., Gobillon, L. (2007), "Spatial wage disparities: Sorting matters!", Journal of Urban Economics, 63(2), 723-742.

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

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

[26] Glaeser, E. (2011), *Triumph of the city: how our greatest invention makes us richer, smarter, greener, healthier, and happier.* New York: Penguin Press.

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

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

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

[30] 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.

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

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

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

[34] Langer, A. and C. Winston (2008), "Toward A Comprehensive Assessment of Road Pricing Accounting for Land Use", Brookings-Wharton Papers on Urban Affairs, 127-175.

[35] Libecap, G. D. and D. Lueck (2011), "The Demarcation of Land and the Role of Coordinating Property Institutions", Journal of Political Economy, 119-33, 426-467.

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

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

[38] McMillen, D. P. and S. C. Smith (2003), "The number of subcenters in large urban areas", Journal of Urban Economics, 53 (3), pp. 321-338.

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

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

[41] Mitric, S. and I. Chatterton (2005), "Towards a Discussion of Support to Urban Transport Development in India", Washington, DC: World Bank.

[42] Mohan, D. (2013), "Moving around in Indian Cities", Economic and Political Weekly, 48, 40-48.

[43] Morten, M. and J. Oliveira (2014), "Migration, Roads and Labor Market Integration: Evidence from a Planned Capital City", working paper.

[44] Mulalic, I., J.N. Van Ommeren, and N. Pilegaard (2014), "Wages and Commuting: Quasi-natural Experiments' Evidence from Firms that Relocate", The Economic Journal, 124, 1086–1105.

[45] 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.

[46] Stutzer, A. and B. S. Frey (2008), "Stress that doesn't pay: The commuting paradox", The Scandinavian Journal of Economics, 110 (2), 339-366.

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

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

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

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

[51] 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.

[52] Small, C., 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.

[53] Small, K. A., C. Winston, and J. Yan (2005), "Uncovering the Distribution of Motorists' preferences for Travel Time and Reliability", Econometrica, 73 (4), 1367-1382.

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

[55] Storeygard, A. (forthcoming), "Farther on Down the Road: Transport Costs, Trade and Urban Growth in Sub-Saharan Africa", Review of Economic Studies.

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

[57] 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.

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

[59] Van Ommeren, J. and M. Fosgerau (2009), "Workers' Marginal Costs of Commuting", Journal of Urban Economics, 65 (1), 38-47.

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

[61] Zax, J. (1991), "Compensation for commutes in labor and housing markets", Journal of Urban Economics, 30 (2), 192–207.

Appendix

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.[45] In order to define the CBD using a uniform criterion for all cities, I consider the centroid of the 1950 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:

$$y_i = f(x_i^N, x_i^E) + \varepsilon_i \tag{A1}$$

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

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

$$y_i = g(DCBD_i) + \sum_{j=1}^{S} \delta_j^1 (D_{ji})^{-1} + \delta_j^2(-D_{ji}) + u_i \tag{A2}$$

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

[45] 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.

of interest δ_j^1 and δ_j^2. (A2) 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 (A2) is that, with a large number of candidate subcenters, the distance variables D_{ij} can be highly multicollinear. To cope with this problem, a stepwise procedure is used to select which subcenter distance variables to include in the regression. In the first step, all distance variables are included. At each step, the variable corresponding to the lowest t statistic is dropped from the regression, and the process is repeated until all subcenter distance variables in the regression have a positive coefficient, significant at the 20% level. The final list of subcenters includes the sites with positive coefficients on either $(D_j)^{-1}$ or $(-D_j)$.

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
Avg. yearly rent per m^2, 2014 Rs.	895	603.27	324.81	104.52	3821.59

Table 2: First Stage

	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	OLS	OLS	OLS	OLS	OLS
	Norm. shape of actual footprint	Shape of actual footprint km^2	Log area of actual footprint km^2	Norm. shape of actual footprint	Shape of actual footprint km^2	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^2		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^2		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. Columns (1) and (4) report the results from estimating equation (43). Columns (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 columns (1) and (4), and on a city-specific one in column (2), (3), (5), (6) - see Section 5.1. Angrist-Pischke F statistics are reported in columns (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)	(2)	(3)	(4)	(5)	(6)
	IV	IV	OLS	IV	IV	OLS
	Population density	Log population	Population density	Population density	Log population	Population density
	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^2		-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^2		-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 $\hat{r}$	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 the relationship between city shape and population. Each observation is a city-year. Column (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). Column (2) reports 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 columns (1), (4) and (2), (5) 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 columns (1) and (4) and on a city-specific one in column (2) (5) - see Section 5.1. Column (3), (6) report the same specification as column (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: 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 $\hat{r}$	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. Columns (1) and (3), (4) are equivalent to columns 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. Columns (2) and (5) are equivalent to columns (1), (2) and (4), (5) 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 exclude 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 4 (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 $\hat{r}$	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. Columns (1) and (3), (4) are equivalent to columns 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. Columns (2) and (5) are equivalent to columns (1), (2) and (4), (5) 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 exclude 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 5: Impact of City Shape on Wages

	(1)	(2)	(3)	(4)	(5)	(6)
	IV	IV	OLS	IV	IV	OLS
	All districts			*Only districts with one city*		
	Dependent variable: log wage					
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. Columns (1) and (4) report the results from estimating equation (47) (single-instrument specification) by IV. Columns (2) and (5) report the results from estimating equation (42) (double-instrument specification) by IV. Instruments are described in Section 5.2. Angrist-Pischke F statistics for the shape variable are reported. Columns (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 columns (1) and (4), and on a city-specific one in columns (2) and (5) - see Section 5.1. In columns (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 6: Impact of City Shape on Housing Rents

	(1)	(2)	(3)	(4)	(5)	(6)
	IV	IV	OLS	IV	IV	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 $\hat{r}$	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 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 km². Columns (1) and (4) report the results from estimating equation (47) (single-instrument specification) by IV. Columns (2) and (5) report the results from estimating equation (42) (double-instrument specification) by IV. Instruments are described in Section 5.2. Angrist-Pischke F statistics for the shape variable are reported. Columns (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 columns (1) and (4), and on a city-specific one in columns (2) and (5) - see Section 5.1. In columns (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 7: 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	1.671**	1.607**	1.231***			
	(0.831)	(0.797)	(0.430)			
Norm. shape x urban road density, 1981				0.000109	0.000214	0.000757**
				(0.000388)	(0.000368)	(0.000315)
Norm. shape x state urban road density, 1991	-99.33***	-94.23***	-52.63***	-83.63***	-78.21***	-35.37***
	(34.36)	(33.23)	(12.09)	(28.61)	(27.96)	(9.469)
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. Each observation is a city-year. All columns 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 columns (1), (2), (4), (5) and on a city-specific one in columns (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 columns (1) , (2) and (3), normalized shape is interacted with urban road density in the city, as reported in the 1981 Census. In columns (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 columns (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 8: Impact of City Shape on Slum Population

	(1)	(2)	(3)	(4)	(5)	(6)
	IV	IV	OLS	IV	IV	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. Columns (1) and (4) report the results from estimating equation (44) (single-instrument specification) by IV. Columns (2) and (5) report the results from estimating equation (39) (double-instrument specification) by IV. Columns (3) and (6) report the same specification, estimated by OLS. The explanatory variables are shape, in km, and log city area, in km². 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 columns (1) and (4) and on a city-specific one in columns (2) and (5) -- see Section 5.1. The dependent variables are the following: log slum population in the city (columns (1), (2), (3)), log of the share of slum to total population in the city (columns (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 9: Impact of City Shape on the Number of Employment Subcenters, 2005

	(1)	(2)	(3)
	IV	IV	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 km² (column (1)) and the log number of employment subcenters (columns (2) and (3)). Column (1) reports estimates from a cross-sectional version of equation (45) (single-instrument specification), estimated by IV. Column (2) reports estimates from a cross-sectional version of equation (42) (double-instrument specification), estimated by IV. Column (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 column (1), and on a city-specific one in in column (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 is calculated from maps constructed from the DMSP/OLS Night-time Lights dataset, in year 2005.

Table 10: Urban Form and Floor Area Ratios

	(1)	(2)	(3)	(4)	(5)	(6)
	A. First-stage impact of FARs on city shape					
	OLS	OLS	OLS	OLS	OLS	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 r̂	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: each observation is a city-year. Panel A of this table reports estimates of the first-stage relationship between Floor Area Ratios, city shape, and area. Columns (1), (2), (3) and (4), (5), (6), panel A, report the same specifications reported in Table 2, 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. Columns (1), (2), (3) consider the average of residential and non-residential FARs, while columns (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 10 (continued): Urban Form and Floor Area Ratios

B. Impact of city shape and FARs on population						
	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 x FAR	-10.91			-16.89		
	(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 x FAR		0.00290	-0.0160		0.0688*	-0.00767
		(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 x FAR		0.0995	0.0665		0.00166	0.0981***
		(0.116)	(0.0471)		(0.0886)	(0.0368)
Observations	252	252	252	252	252	252
Model for $\hat{r}$	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: each observation is a city-year. Panel B reports estimates of the relationship between Floor Area Ratios, shape, and population. Columns (1), (2), (3) and (4), (5), (6), panel B, report the same specifications reported in Table 3, 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. Columns (1), (2), (3) consider the average of residential and non-residential FARs, while columns (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.

Appendix Table 1: Impact of Shape on Population, Robustness to Initial Shape x Year Fixed Effects

	(1)	(2)
	IV	IV
	Population density	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
Number of id	432	432
F stat shape	9.115	26.699
Model for $\hat{r}$	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. Columns (1) and (2) are equivalent to columns (1), (4) and (2), (5) in Table 3. Column (1) reports IV estimates of the relationship between normalized city shape and population density, in thousand inhabitants per km^2. Column (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 column (1) and on a city-specific one in column (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.

Appendix Table 2: 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 $\hat{r}$	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. Columns (1), (3) report the results from estimating equation (44) (single-instrument specification) by IV. Columns (2), (4) report the results from estimating equation (39) (double-instrument specification) by IV. Instruments are described in Section 5.2. The F statistic reported is the Angrist-Pischke F statistic for the shape variable. The construction of the potential footprint is based on a common rate model for city expansion in columns (1), (3), and on a city-specific one in columns (2), (4) -- see Section 5.1. In columns (3), (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.

Appendix Table 3: Urban Form and Floor Area Ratios

	(1)	(2)	(3)	(4)	(5)	(6)
	FARs Determinants, 2005					
	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 $\hat{r}$	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. Columns (1) and (4) estimate a cross-sectional version of equation (44) (single-instrument specification), with log FARs as a dependent variable, estimated by IV. Columns (2) and (5) estimate a cross-sectional version of equation (39) (double-instrument specification), with log FARs as a dependent variable, estimated by IV. Columns (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. Columns (1), (2), (3) consider the average of residential and non-residential FARs, while columns (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.

Appendix Table 4: Impact of City Shape on Sectoral Shares

	(1)	(2)	(3)	(4)	(5)
	IV	IV	IV	IV	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 $\hat{r}$	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 (col. 1), services (col. 2), transport and storage (col. 3), retail (col. 4), manufacturing (col. 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.

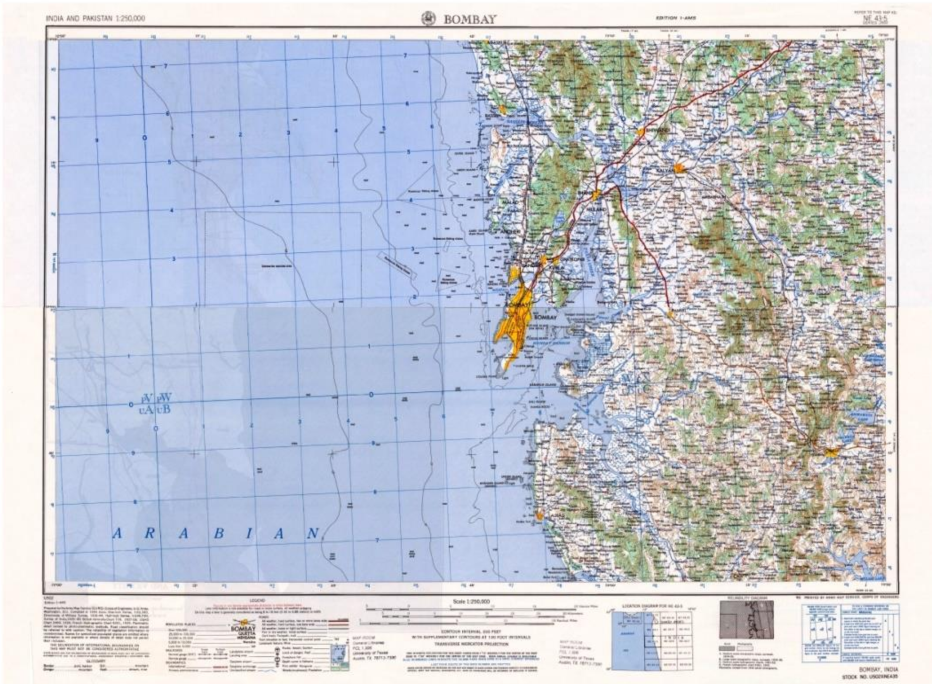

Figure 1A

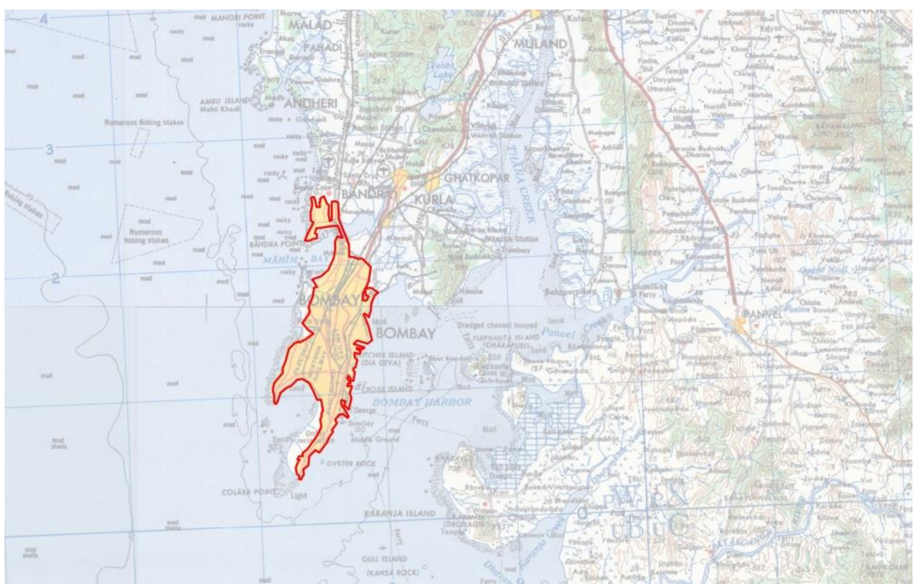

Figure 1B

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

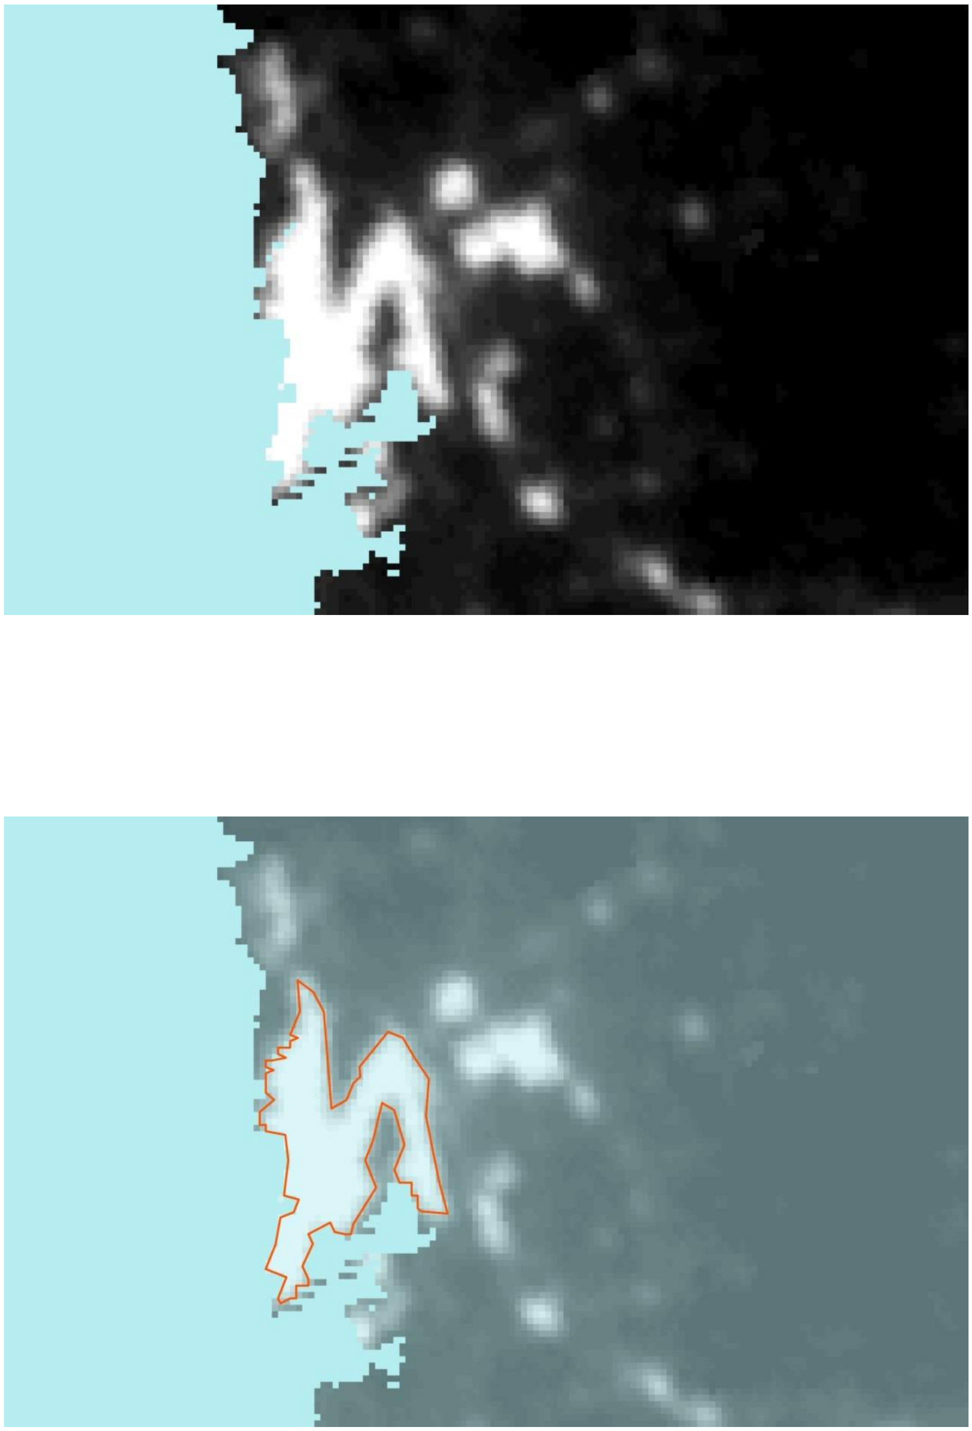

Figure 2
DMS/OLS nighttime lights, year 1992, luminosity threshold : 40.

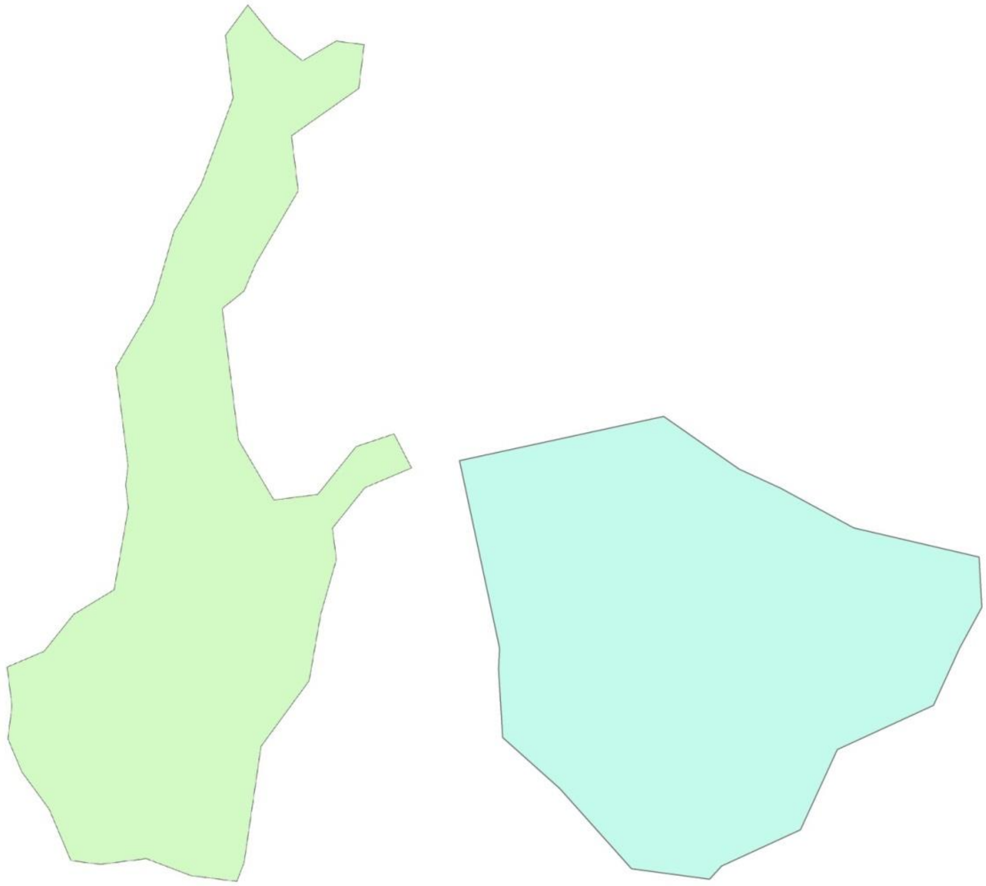

	Kolkata		Bengaluru	
Shape metric		Normalized		Normalized
remoteness, km	14.8	0.99	10.3	0.69
spin, km^2	288.4	1.29	120.9	0.54
disconnection, km	20.2	1.35	14	0.94
range, km	62.5	4.18	36.6	2.45

Figure 3
Shape metrics: an example

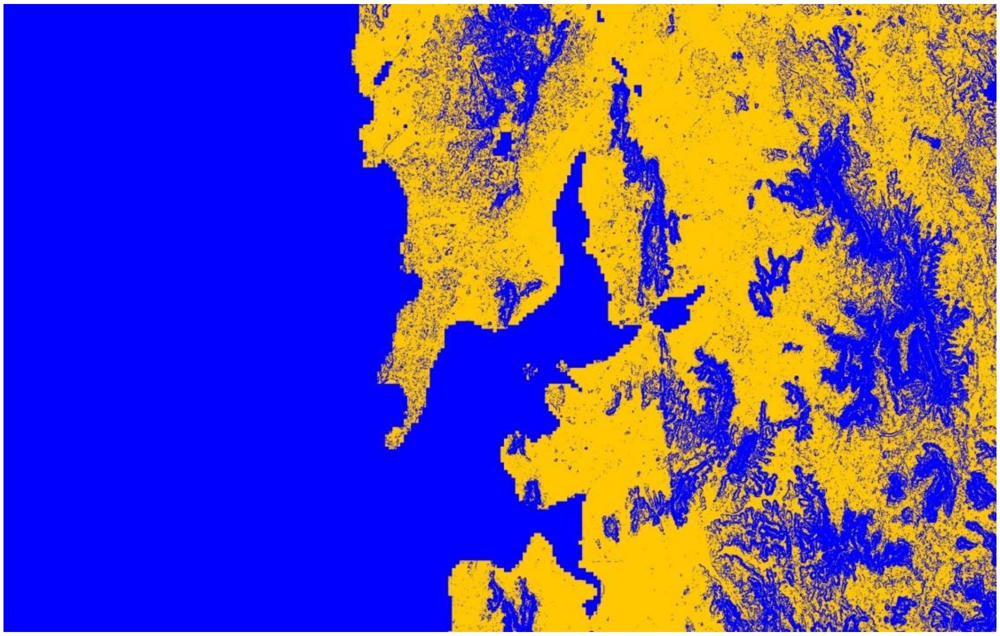

constrained
developable

Figure 4
Developable vs. constrained land

Figure 5a

Figure 5b

Figure 5
Instrument construction

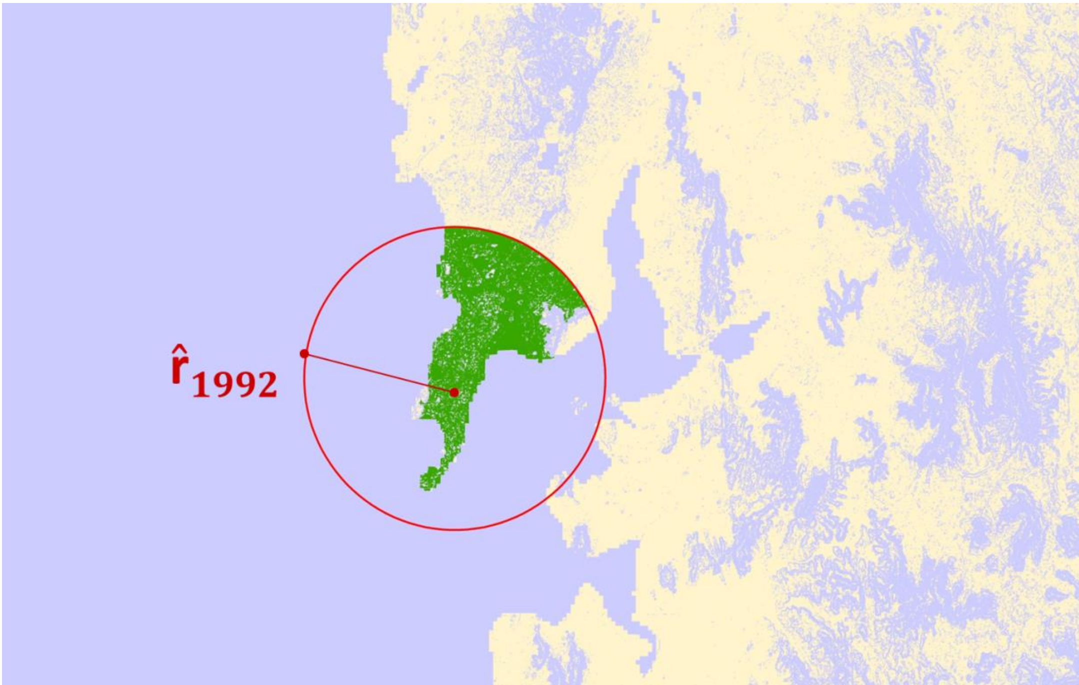

Figure 5c

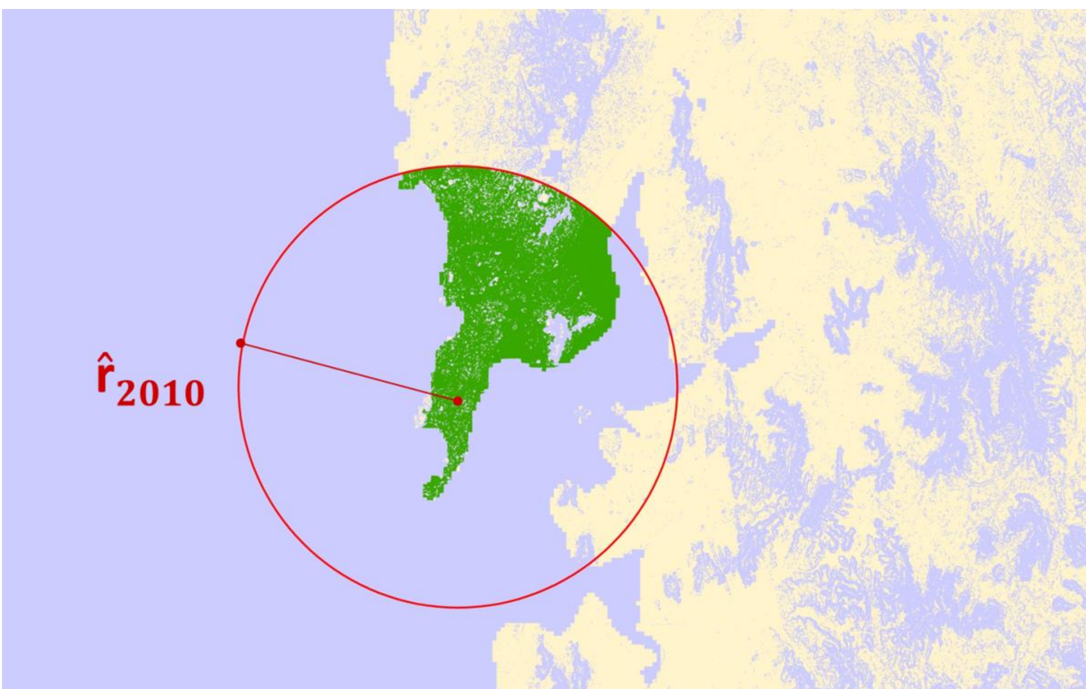

Figure 5d

Figure 5
Instrument construction