

Cities in Bad Shape: Urban Geometry in India*

Mariaflavia Harari[†]

Massachusetts Institute of Technology

JOB MARKET PAPER

11 January 2015

Abstract

Cities are valuable to the extent they bring people (and jobs) together. To what extent is this value affected by the difficulty of commuting from various points in the city to others? While many factors can affect commuting length, this paper investigates one determining factor of urban commuting efficiency, previously highlighted by urban planners but overlooked by economists: city shape. A satellite-derived dataset of nighttime 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 that 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 each circle. With this instrument in hand, I investigate how city shape affects the location choices of consumers, 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 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 my advisers Esther Duflo, Ben Olken, and Dave Donaldson for their invaluable help throughout this project. I also thank Alex Bartik, Jie Bai, Nathaniel Baum-Snow, Alain Bertaud, Melissa Dell, John Firth, Ludovica Gazzè, Michael Greenstone, Gabriel Kreindler, Matthew Lowe, Rachael Meager, Yuhei Miyauchi, Hoai-Luu Nguyen, Paul Novosad, Arianna Ornaghi, Bimal Patel, Tommaso Porzio, Champaka Rajagopal, Otis Ried, Adam Sacarny, Albert Saiz, Ashish Shenoy, Chris Small, Kala Sridhar, William Wheaton, and participants at the MIT Development Lunch, MIT Applied Micro Lunch, and NEUDC for helpful comments and discussions at various stages of this work.

[†]harari@mit.edu. Website: <http://economics.mit.edu/grad/harari/research>.

1 Introduction

Urban transportation networks perform the key role of connecting people and jobs within cities. The ease of intra-urban commutes affects the range of jobs and services accessible in a city, and, potentially, the extent to which the benefits from proximity can be realized. Urban mobility is particularly challenging in the sprawling cities of the developing world, where infrastructure is often inadequate, and most inhabitants cannot afford individual means of transportation (Bertaud, 2004; Cervero, 2013). Commuting length is determined, in equilibrium, by the interaction of a wide range of factors including residential patterns, infrastructure, and regulation. This paper attempts to quantify the loss associated to lengthy urban commutes in the context of India by focusing on one determining factor of urban commuting efficiency, previously highlighted by urban planners but overlooked by economists: the spatial layout of cities. All else being equal, more compact urban geometries are characterized by shorter potential trips and more cost-effective transport networks, which, in turn, have the potential to affect productivity and welfare (Bertaud, 2004; Cervero, 2001). In this study, I exploit plausibly exogenous variation in city shape driven by geographic barriers and assess the welfare loss due to non-compact shape in a revealed preference Rosen-Roback framework.

India represents a promising setting to study the interactions of urban spatial structures and mobility. Like most developing countries, India is experiencing fast urban growth¹, accompanied by a significant physical expansion of urban footprints. This provides a unique opportunity to observe the shapes of cities as they evolve and expand over time. Moreover, unlike most other developing countries, India is characterized by an unusually large number of highly-populated cities, lending itself to an econometric approach based on a city-year panel. 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). Sprawl has been linked to a number of potentially distortive land use regulations, most notably vertical limits in the form of Floor Area Ratios (Bertaud, 2002a; Bertaud and Brueckner, 2005; Brueckner and Sridhar, 2012; Glaeser, 2011; Sridhar, 2010; World

¹According to the 2011 Census, the urban population amounts to 377 million, increasing from 285 million in 2001 and 217 million in 1991, representing respectively between 25 and 31 percent of the total. It is predicted that another 250 million will join the urban ranks by 2030 (Mc Kinsey, 2010).

Bank, 2013).²

This paper primarily explores the implications of city shape for the location choices of consumers across cities, and in particular, how much consumers value the potential commute lengths implied by different urban shapes. It is plausible that urban migrants consider the relative ease of commutes when evaluating the trade-offs associated to different locations. I investigate this in the framework of spatial equilibrium across cities à la Rosen-Roback, modelling “compact city shape”, i.e., an urban geometry conducive to shorter potential within-city trips, as an amenity. A complementary question relates to the interactions of city shape with policy, viewed both as a co-determinant of urban geometry - as in the case of land use regulations - and as a tool to counteract the effects of poor geometry on mobility - for instance, through infrastructural investment.

In terms of methodology, my empirical analysis is conducted mostly at the city-year level, and is based on an instrumental variables approach. I assemble an original panel dataset, covering more than 450 Indian cities, which includes detailed information on each city’s spatial properties and micro-geography, as well as economic outcomes from the Census and other data sources. In particular, urban footprints are retrieved by combining newly geo-referenced historic maps (1950) with a satellite-derived dataset of night-time lights (1992-2010). I then compute several geometric indicators for urban shape, used in urban planning as proxies for the patterns of within-city trips. I propose an instrument for city shape that combines geography with a mechanical model for city expansion. The underlying idea is that, as cities expand in space, they face different geographic constraints - steep terrain or water bodies - leading to departures from an ideal circular expansion path. 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 shape properties. I then proceed to instrument the shape 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

²Another 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).

impact of shape thus relies on comparing geography-driven changes in urban shape, over time, for each city. This instrument appears to have strong explanatory power, which is not limited to extremely constrained topographies (e.g., coastal or mountainous cities) in my sample.

With this instrument in hand, I investigate how city shape affects the spatial equilibrium across cities. My findings are broadly consistent with compact city shape being a consumption amenity. All else being equal, more compact cities grow faster. 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. The negative effects of deteriorating geometry on population appear to be mitigated by road infrastructure. This suggests that urban transit is indeed the main channel through which non-compact shape affects consumers, and supports the interpretation of city shape as a shifter of commuting costs. The loss associated with non-compact shape appears to be substantial: a one-standard deviation deterioration in city shape, corresponding to a 720 meter increase in the average within-city round-trip,³ entails a welfare loss equivalent to a 5% decrease in income. This is considerably larger than the direct monetary and opportunity cost associated to lengthier commutes. Less compact cities also appear to attract fewer low-income immigrants, as captured by the share of slum dwellers.

A question then arises concerning the role of policy: if city shape indeed has welfare implications, what are possible policy responses to existing geographic constraints? I find that more permissive building height restrictions, in the form of laxer Floor Area Ratios (FARs), result in cities that are less spread out in space and more compact than their topographies would predict. This is consistent with one of the most common arguments against restrictive FARs in India, namely that they cause sprawl by preventing cities from growing vertically (Glaeser, 2011).

This study contributes to the existing literature in a number of directions. First, it sheds light on the loss associated with lengthy urban commutes and on how this affects the spatial equilibrium across cities, in the context of a rapidly urbanizing developing country. Second, it provides evidence on the economic implications of urban geometry, a feature of urban form that has attracted little attention in economics, but which I find to have potentially important effects through commuting length. The third set of contributions is methodological. I employ

³As 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.

night-time satellite imagery as a new way to map urban expansion over time, following an approach which has been used in urban remote sensing, but not in economics. Furthermore, I develop a novel instrument for urban shape based on the interaction between mechanically predicted city growth and topography.

The remainder of the paper is organized as follows. Section 2 briefly reviews the existing literature. Section 3 outlines the conceptual framework. Section 4 documents my data sources and describes the geometric indicators I employ. 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 effects of city shape on the location choices of consumers. Section 7 provides results on various responses to city shape, including interactions between topography and policy. Section 8 concludes and discusses indications for future work.

2 Previous Literature

As pointed out by Van Ommeren and Fosgerau (2009), we know relatively little about the actual size of commuting costs. Studies in transport economics consider the time value of commuting (e.g., Small, K., et al., 2005) using stated and revealed preference data on commuter choices (e.g. whether to pay a toll for congestion-free express travel). A few studies have relied on the trade-off between house prices or wages and commuting distance, to quantify the marginal willingness to pay for commuting, in the framework of spatial equilibrium within cities. Zax (1991) exploits the trade-off between wages and the length of the commute using hedonic wage models. Tse and Chan (2003) consider housing prices as a function of distances to the Central Business District (CBD). My approach differs in a number of ways. I consider spatial equilibrium across, rather than within, cities, thus capturing consumers' valuations of average commute lengths at the city level. Moreover, I jointly consider wages and housing rents, and I exploit exogenous variation in the length of potential trips, rather than actual trips.

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. The link between city shape and urban mobility has been highlighted by the urban planning literature. Bertaud (2004) argues that contiguous, compact,

and predominantly monocentric urban morphologies can improve the welfare of city dwellers by providing better job access. Cervero (2001) emphasizes how compact, accessible cities are potentially more productive, through a combination of labor market pooling, savings in transporting inputs, and information spillovers, and shows a correlation between labor productivity and accessibility in a cross-section of US cities. My work elaborates on these insights, providing causal evidence of the economic implications of city compactness. Several studies have also investigated the implications of a polycentric urban structure for the job-housing balance and observed commutes (e.g., Giuliano and Small, 1993; Gordon et al., 1989). I address the link between city shape and polycentricity in Section 7.3.

A large empirical literature investigates urban sprawl (see Glaeser and Kahn, 2004), typically in the US context, suggesting longer commutes as one of the potential costs of sprawl (Bajari and Kahn, 2004). 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 of sprawl, however, the focus is on decentralization and density. On the other hand, 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.⁴

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 (2011) 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.

Finally, this paper 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, 2014). In particular, Morten and Oliveira (2014) also estimate a model of spatial equilibrium across cities and employ an instrumental variables approach.

⁴My 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. However, their analysis is based on a cross-section of US cities and does not address the endogeneity of city shape.

Differently from these studies, I do not look at the impact of roads connecting cities, but instead focus on urban geometry as a proxy for the trips within cities.

3 Conceptual Framework

In this study, I interpret city shape primarily as a shifter of commuting length: all else being equal, including city size, a city with a more compact geometry is characterized by shorter within-city trips.⁵ It is plausible that consumers incorporate considerations on the relative ease of commutes when evaluating the trade-offs associated with different locations. This might be even more relevant in the context of India, in which most migrants to urban areas cannot afford individual means of transport. A natural starting point in thinking how city shape affects the location choices of consumers across cities is the framework of spatial equilibrium à la Rosen-Roback (Rosen 1979; Roback, 1982). The basic underlying idea of this framework is that consumers and firms must be indifferent across cities, and wages and rents allocate people and firms to cities with different levels of amenities. Although Glaeser (2008) acknowledges that amenities presumably include also commuting times, I am not aware of studies explicitly looking at the link between commuting, urban geometry, and the spatial equilibrium across cities.⁶ 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.

I draw upon the Roback model in its simplest form (Roback, 1982), following 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

⁵In Section 4.2, I illustrate the geometric indicators I employ to measure compactness, which are all based on the relative distances between points within a footprint. In Section 6.5, I offer evidence substantiating the claim that city shape matters primarily through urban transit. In Section 8, I discuss briefly other possible second-order channels through which city shape might affect consumers.

⁶A large empirical literature has employed the Rosen-Roback framework to investigate the value of amenities. This literature, however, has almost exclusively focused on the US, and has not always adequately addressed the endogeneity of amenities (Gyourko, Kahn, and Tracy, 1999).

net income, i.e., labor income minus housing costs, and on a city-specific bundle of consumption amenities θ . Their optimization problem reads:

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

where p_h is the rental price of housing, and

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

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

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

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

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

The intuition for this condition is that consumers, in equilibrium, pay for amenities through lower wages (W) or through higher housing prices (p_h)⁷. 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 the marginal utility of income. Holding indirect utility \bar{v} constant, differentiating this expression with respect to some exogenous variable S - which could be (instrumented) city geometry - yields:

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

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 ,

⁷This 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.

and a local production amenity A . Their technology also requires traded capital K and a fixed supply of non-traded capital \bar{Z} . Firms solve the following profit maximization problem:

$$\max_{N,K} \{Y(N, K, \bar{Z}, A) - WN - K\} \quad (6)$$

where

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

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

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

To close the model, we need to specify the construction sector. Developers produce housing H , using land l and "building height" h . In each location there is a fixed supply of land \bar{L} , as a result of land use regulations.⁸ Denoting with p_l the price of land, their maximization problem reads:

$$\max_H \{p_h H - C(H)\} \quad (9)$$

where

$$H = l \cdot h \quad (10)$$

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

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 housing market equilibrium condition:

$$(\delta - 1) \log(H) = \log(p_h) - \log(c_0 \delta) - (\delta - 1) \log(N) + (\delta - 1) \log(\bar{L}) \quad (12)$$

⁸In 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).

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:

$$\log(N) = \frac{(\delta(1 - \alpha) + \alpha) \log(A) + (1 - \gamma) (\delta \log(\theta) + \alpha(\delta - 1) \log(\bar{L}))}{\delta(1 - \beta - \gamma) + \alpha\beta(\delta - 1)} + K_N \quad (13)$$

$$\log(W) = \frac{(\delta - 1)\alpha \log(A) - (1 - \beta - \gamma) (\delta \log(\theta) + \alpha(\delta - 1) \log(\bar{L}))}{\delta(1 - \beta - \gamma) + \alpha\beta(\delta - 1)} + K_W \quad (14)$$

$$\log(p_h) = \frac{(\delta - 1) (\log(A) + \beta \log(\theta) - (1 - \beta - \gamma) \log(\bar{L}))}{\delta(1 - \beta - \gamma) + \alpha\beta(\delta - 1)} + K_P. \quad (15)$$

These conditions translate into the following predictions:

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

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

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

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. Assume non-compact shape is purely a consumption disamenity, which decreases consumers' utility, all else being equal, but does not directly affect firms' productivity:

$$\frac{\partial \theta}{\partial S} < 0, \quad \frac{\partial A}{\partial S} = 0. \quad (19)$$

In this case we should observe the following reduced-form relationships:

$$\frac{dN}{dS} < 0, \frac{dW}{dS} > 0, \frac{dp_h}{dS} < 0 \quad (20)$$

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:

$$\frac{\partial \theta}{\partial S} < 0, \frac{\partial A}{\partial S} < 0. \quad (21)$$

This would imply the following:

$$\frac{dN}{dS} < 0, \frac{dW}{dS} \geq 0, \frac{dp_h}{dS} < 0 \quad (22)$$

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:

$$\log(A) = \kappa_A + \lambda_A S \quad (23)$$

$$\log(\theta) = \kappa_\theta + \lambda_\theta S. \quad (24)$$

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

$$\log(N) = \frac{[(\delta(1-\alpha) + \alpha)\lambda_A + (1-\gamma)\delta\lambda_\theta] \cdot S + (1-\gamma)\alpha(\delta-1)\log(\bar{L})}{\delta(1-\beta-\gamma) + \alpha\beta(\delta-1)} + K_N \quad (25)$$

$$\log(W) = \frac{[(\delta - 1)\alpha\lambda_A - (1 - \beta - \gamma)\delta\lambda_\theta] \cdot S - (1 - \beta - \gamma)\alpha(\delta - 1)\log(\bar{L})}{\delta(1 - \beta - \gamma) + \alpha\beta(\delta - 1)} + K_W \quad (26)$$

$$\log(p_h) = \frac{[(\delta - 1)\lambda_A + (\delta - 1)\beta\lambda_\theta] \cdot S - (1 - \beta - \gamma)(\delta - 1)\log(\bar{L})}{\delta(1 - \beta - \gamma) + \alpha\beta(\delta - 1)} + K_P. \quad (27)$$

For ease of exposition, define:

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

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

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

This allows us to rewrite (25), (26), (27) in a more parsimonious form:

$$\log(N) = B_N S + D_N \log(\bar{L}) + K_N \quad (31)$$

$$\log(W) = B_W S + D_W \log(\bar{L}) + K_W \quad (32)$$

$$\log(p_h) = B_P S + D_P \log(\bar{L}) + K_P. \quad (33)$$

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

$$\lambda_A = (1 - \beta - \gamma)B_N + (1 - \gamma)B_W \quad (34)$$

$$\lambda_\theta = \alpha B_P - B_W. \quad (35)$$

The welfare impact of a marginal increase in S can be quantified using equation (5), which states that in equilibrium a marginal change in $\log(\theta)$ needs to be compensated one-to-one by a change in $\log(W)$:

$$\frac{\partial \log(\theta)}{\partial S} = \lambda_\theta = \alpha \frac{\partial \log(p_h)}{\partial S} - \frac{\partial \log(W)}{\partial S} = \alpha B_P - B_W. \quad (36)$$

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 :

$$\widehat{\lambda}_A = (1 - \beta - \gamma)\widehat{B}_N + (1 - \gamma)\widehat{B}_W \quad (37)$$

$$\widehat{\lambda}_\theta = \alpha\widehat{B}_P - \widehat{B}_W. \quad (38)$$

This approach captures the overall net effect of S on the marginal city dweller without explicitly modelling the mechanism through which S enters the decisions of consumers. In Section 6.2, I provide empirical evidence suggesting that the urban transit channel is indeed involved, and in the concluding Section, I discuss some alternative, second-order channels through which city shape might affect consumers.

This simple model 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, and an overestimate for those with weak preferences. The empirical evidence presented in Section 6.2 indicates that there might indeed be sorting across cities with different geometries.

The model could be extended to allow for congestion or agglomeration in consumption (or production) by adding a new term, a function of city size N , to the utility of consumers (or to the production function of firms). In particular, if there are congestion externalities in consumption, the indirect utility of consumers will depend on city size as well. If shape is a 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 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. When I estimate the consumption amenity value of compact shape using equation (38), I will be capturing the equilibrium effect of shape, gross of congestion. If compact shape is a consumption amenity and more 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 (37), will

be an upper bound for λ_A .⁹

Reduced-form estimates for B_N, B_W, B_P are presented in Sections 6.1 and 6.2, whereas Section 6.3 provides estimates for parameters $\lambda_A, \lambda_\theta$. The next two Sections present the data sources and empirical strategy employed in the estimation.

4 Data Sources

I assemble an unbalanced panel of city-years, covering all Indian cities for which a footprint could be retrieved based on the methodology explained below.

4.1 Urban Footprints

I retrieve the boundaries of urban footprints 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 (Henderson, V., et al., 2012). However, the use of the DMSP/OLS dataset for delineating urban areas is quite common in urban remote sensing (Henderson, M., et al., 2003; Small, C., et al., 2005; Small, C., 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

⁹An additional source of exogenous variation in city size would be needed in order to compute the pure amenity value of compact shape, net of congestion / agglomeration.

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.¹⁰ 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.¹¹ In my full panel (including years 1950 and 1992-2010), the average city footprint occupies an area of approximately 63 square km.¹²

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.¹³ An extensive portion of the urban remote sensing literature compares the accuracy of this approach in mapping urban areas with that attainable with alternative satellite-based products, in particular day-time imagery (e.g., Henderson, M., et al., 2003; Small, C., 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

¹⁰Determining 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.

¹¹For 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).

¹²My results are robust to using alternative luminosity thresholds between 20 and 40. Results are available upon request.

¹³Using 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.

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 (Small, C., et al., 2005).¹⁴ This can only partially be limited by setting high luminosity thresholds. In my panel, urban footprints as reported for years 1992-2010 thus reflect a broad definition of urban agglomeration, which typically goes beyond the current administrative boundaries. This contrasts with urban boundaries reported in the US Army maps, which seem to reflect a more restrictive definition of urban areas (although no specific documentation is available). Throughout my analysis, I include year fixed effects, which amongst other things control for these differences in data sources, as well as for different calibrations of the night-time lights satellites.

By combining the US Army maps (1950s) with yearly maps obtained from the night-time lights dataset (1992-2010), I thus assemble an unbalanced¹⁵ 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 problems of non-classical measurement error, which could affect my data sources - for instance due to the well-known correlation between income and luminosity.

4.2 Shape Metrics

The indicators of city shape that I employ, based on those in Angel et al. (2009a, 2009b),¹⁶ are used in landscape ecology and urban studies to proxy for the length within-city trips and infer

¹⁴DMSP-OLS night-time imagery overestimates the actual extent of lit area on the ground, due to a combination of coarse spatial resolution, overlap between pixels, and minor geolocation errors (Small et al., 2005).

¹⁵The 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.

¹⁶I 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.

travel costs. They are all based on the distribution of points around the polygon's centroid¹⁷ 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 "tendrill-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 the center.

(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 mechanically correlated with polygon area. In order to separate the effect of geometry *per se* from that of city size, 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. An alternative approach 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). Conditional on footprint area, higher values of these indexes indicate longer within-city trips.

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

¹⁷The centroid of a polygon, or center of gravity, is the point that minimizes the sum of squared Euclidean distances between itself and each vertex.

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 trip to the center would be shorter by 4.5 km and the average trip within the city would be shorter by 6.2 km. The Indian Ministry of Urban Development (2008) estimates the average commute speed in million-plus cities to be of 12 km per hour in 2011, which is predicted to become 9 km by 2021. According to the 2011 estimated speed, the above differences in trip length translate to a difference in commute times of respectively 22.5 minutes (average trip to the center) and 31 minutes (average within-city trip). Although this is a very rough calculation, it is nevertheless revealing that city geometry indeed has potentially sizeable impacts on commute times.

4.3 Geography

Following Saiz (2011), 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.

4.4 Population and Other Census Data

City-level data for India is difficult to obtain (Greenstone and Hanna, forthcoming). 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.¹⁸

¹⁸Historic 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

Outcomes other than population are not consistently available for all Census years. 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.

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.¹⁹

4.5 Wages and Rents

For outcomes other than those available in the Census, 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 (forthcoming): 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. In these cases, I follow several matching approaches: considering only the main city for that district, and dropping the district entirely. I show results following both approaches. 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, 2010.²⁰ These are repeated cross-sections of plant-level data collected by the http://www.censusindia.gov.in/DigitalLibrary/Archive_home.aspx.

¹⁹In 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.

²⁰These are all the waves I could access as of June 2014.

Ministry of Programme Planning and Implementation of the Government of India. The ASI covers all registered manufacturing plants in India with more than fifty workers (one hundred if without power) and a random one-third sample of registered plants with more than ten workers (twenty if without power) but less than fifty (or one hundred) workers. As mentioned by Fernandes and Sharma (2012) amongst others, the ASI data are extremely noisy in some years, which introduces a further source of measurement error. The average individual yearly wage in this panel amounts to 94 thousand Rs at current prices.

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 (Dek, 2006). To cope with this problem, I also construct an alternative proxy for housing rents which focuses on the upper half of the distribution of rents per meter, which is less likely to include observations from rent-controlled housing.

4.6 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 current road network is constructed from the maps available on Openstreetmap,²¹ a collaborative mapping project that provides crowdsourced maps of the world. Openstreetmap data has been favorably compared with proprietary data sources and is continuously updated. As such, it reflects the current state of the street network. I consider the most recent night-time lights-based map of urban boundaries in my sample - corresponding to 2010 - and overlap it

²¹<http://www.openstreetmap.org/about>

with the street network map of India, provided by Openstreetmap. Information on the type of road - whether trunk, residential, secondary etc. - is also provided. Given the collaborative nature of Openstreetmap, there is a concern that the level of detail of such maps might be higher in larger cities, or in neighborhoods with more economic activity. Upon visual inspection, it appears that smaller roads are reported only in relatively bigger cities. To avoid this source of non-classical measurement error, I exclude small roads and focus on those denoted as "trunk", "primary", or "secondary". I then compute the total road length by considering street segments contained within the urban boundary. The average road density in my sample as of year 2010, obtained by dividing total city road length by footprint area, is 2.4 km per square km. Given the tendency of night-time lights to overestimate urban boundaries, this figure should be considered an underestimate of the actual road density.

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).

Data on the spatial distribution of employment in year 2005 is derived from the urban Directories of Establishments, pertaining to the 5th Economic Census. 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. Town or district identifiers are not provided to the general public. However, in year 2005, 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.²² 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

²²My results are robust to excluding firms whose address can only be approximately located by Google maps (available upon request).

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.

5 Empirical Strategy

The objective of my empirical analysis is to estimate the effects of city shape on a number of city-year level outcomes, most notably population, wages, and rents. My data has the structure of an unbalanced city-year panel. In every year, I observe the geometric properties of the footprint of each city, namely footprint area and the different shape metrics described above. The goal is to estimate the relationship between shape in a given city-year on a number of city-year level outcomes, conditional on city and year fixed effects. This strategy exploits variation in urban shape, which is both cross-sectional and temporal. City and year fixed effects account for time-invariant city characteristics and for country-level trends in population and other outcomes. The identification of the impacts of shape thus relies on comparing *changes* in urban shape, over time, for each city.²³

5.1 Instrumental Variable Construction

A major concern in estimating the relationship between city shape and city-level outcomes 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 policy.²⁴ Cities that experience faster population growth might be expanding in a more chaotic and unplanned fashion, generating a "leapfrog" pattern of development, which translates into less

²³As discussed below, some 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.

²⁴Regulation and infrastructural investment can affect urban shape both directly, through master plans, and indirectly (Bertaud, 2004). 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.

compact shapes. At the same time, cities that exhibit faster growth rates might be the object of more stringent regulations. For instance, restrictive Floor Area Ratios in India have been motivated by a perceived need to reduce urban densities (Sridhar, 2010). In order to address this endogeneity problem, I employ an IV approach, constructing an instrument for city shape which varies at the city-year level.

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 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 1951²⁵ (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 $t \in \{1992, 1993, \dots, 2010\}$ I consider concentrically larger radii $\hat{r}_{c,t}$ around the historic footprint, and construct corresponding potential footprints lying within these predicted radii (Figures 5c and 5d).

²⁵The 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.

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 "common rate" version, and a "city-specific" one.

Common rate: In this first 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. More formally, 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:

$$\log(area_{c,t}) = \theta_c + \gamma_t + \varepsilon_{c,t} \quad (39)$$

where θ_c and γ_t denote city and year fixed effects. From the regression above, I obtain $\widehat{area}_{c,t}$, the *predicted* area of city c in year t .

(ii) I compute $\widehat{r}_{c,t}$ as the radius of a circle with area $\widehat{area}_{c,t}$:

$$\widehat{r}_{c,t} = \sqrt{\frac{\widehat{area}_{c,t}}{\pi}}. \quad (40)$$

City-specific: In this alternative version of the model for city expansion, I make the rate of expansion of $\widehat{r}_{c,t}$ vary 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 $t \in \{1992, 1993, \dots, 2010\}$.

(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:

$$\log(area_{c,t}) = \alpha \cdot \log(\widehat{pop}_{c,t}) + \beta \cdot \log\left(\frac{pop_{c,1950}}{area_{c,1950}}\right) + \gamma_t + \varepsilon_{c,t} \quad (41)$$

from which I obtain an alternative version of $\widehat{area}_{c,t}$, and corresponding $\widehat{r}_{c,t} = \sqrt{\frac{\widehat{area}_{c,t}}{\pi}}$.

The interpretation of the circle with radius $\widehat{r}_{c,t}$ from figures 5c and 5d is thus the following: this is the area the city would occupy if it continued to grow as in 1871-1951, if its density remained the same as in 1951, and if the city could expand freely and symmetrically in all directions, in a fashion that optimizes the length of within-city trips.

This instrument seeks to isolate the variation in urban geometry induced by geography, excluding the variation which results from policy or other endogenous choices. Although resorting to geography arguably helps to address issues of policy endogeneity, there is nevertheless a concern that geography affects location choices directly, for instance, through the inherent amenity (or disamenity) value of water bodies, and not only through the constraints it posits on urban form. These concerns are mitigated by two features of my instrument. First, it is not purely cross-sectional but has time variation. This allows me to control, in the specification below, for time-invariant effects of geography through city fixed effects. Second, it captures a very specific feature of geography: whether it allows for compact development or not. My instrument is not based on the generic presence of topographic constraints, nor on the share of constrained over developable terrain. Rather, it measures the *geometry* of available land. In one of the robustness checks described below, I show that my results are unchanged when I exclude mountain and coastal cities, which would be the two most obvious examples of cities where geography might have a specific (dis)amenity value.

5.2 Estimating Equations

Consider a generic shape metric S - which could be any of the indexes discussed in Section 4.2. 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 (31)-(33), augmented with city and year fixed effects, take the following form:

$$\log(Y_{c,t}) = a \cdot S_{c,t} + b \cdot \log(area_{c,t}) + \mu_c + \rho_t + \eta_{c,t} \quad (42)$$

This equation contains two endogenous regressors: $S_{c,t}$ and $\log(\text{area}_{c,t})$. These are instrumented using respectively $\widetilde{S}_{c,t}$ and $\log(\widehat{\text{pop}}_{c,t})$ - 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:

$$S_{c,t} = \sigma \cdot \widetilde{S}_{c,t} + \delta \cdot \log(\widehat{\text{pop}}_{c,t}) + \omega_c + \varphi_t + \theta_{c,t} \quad (43)$$

and

$$\log(\text{area}_{c,t}) = \alpha \cdot \widetilde{S}_{c,t} + \beta \cdot \log(\widehat{\text{pop}}_{c,t}) + \lambda_c + \gamma_t + \varepsilon_{c,t}. \quad (44)$$

The counterpart of $\log(\text{area}_{c,t})$ in the conceptual framework is $\log(\bar{L})$, where \bar{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 employ, as my benchmark, an alternative, more parsimonious specification, 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 4.2 - and population density. This results in the following, more parsimonious single-instrument specification: define population density²⁶ as

$$d_{c,t} = \frac{\text{pop}_{c,t}}{\text{area}_{c,t}}$$

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

and denote the normalized version of shape metric S with nS . The estimating equation becomes

$$d_{c,t} = a \cdot nS_{c,t} + \mu_c + \rho_t + \eta_{c,t} \quad (45)$$

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

$$nS_{c,t} = \beta \cdot \widetilde{nS_{c,t}} + \lambda_c + \gamma_t + \varepsilon_{c,t}. \quad (46)$$

The same approach can be followed for other outcome variables representing quantities - such as road length. 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. My preferred approach is thus to employ the single-instrument specification, constructing the shape instrument using the "common rate" model for city expansion (see Section 5.1).

While population and road density are meaningful outcomes *per se*, it does not 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:

$$\log(Y_{c,t}) = a \cdot S_{c,t} + \mu_c + \rho_t + \eta_{c,t} \quad (47)$$

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

$$S_{c,t} = \sigma \cdot \widetilde{S_{c,t}} + \omega_c + \varphi_t + \theta_{c,t}. \quad (48)$$

All of the specifications discussed above include year and city fixed effects.²⁷ Although

²⁷This approach relies on comparing changes in shape that the same city undergoes over time. The implicit underlying assumption is a "parallel trends" one, which would be violated if outcomes in cities with different geometries followed differential trends. To mitigate this concern, in one of my robustness checks (Appendix Tables 1 and 2), I augment the specifications above with year fixed effects interacted with the city's initial shape at the beginning of the panel.

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 equations (42) to (48). In all specifications I employ robust standard errors clustered at the city level, to account for arbitrary serial correlation over time in cities.

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 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. Each observation is a city-year. Panels, A, B, C, and D each correspond to one of the four shape metrics discussed in Section 4.2: respectively, remoteness, spin, disconnection, and range.²⁸ Higher values of these indexes represent less compact shapes. Summary statistics are reported in Table 1. Column 1 reports the first-stage for normalized shape (eq. (46)), which is the explanatory variable used in the single-instrument specification. Recall that normalized shape is an area-invariant measure of shape obtained when normalizing a given shape metric by footprint radius. In this specification, the construction of the potential footprint is based on the common rate model for city expansion, outlined in Section 5.1.²⁹ Columns 2 and 3 report the first stage estimates for footprint shape (eq. (43)) and area (eq. (44)), which are relevant for the double-instrument specification. The

²⁸Recall 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.

²⁹The 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.

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 2A, which focuses on the remoteness index. As discussed in Section 4.2, 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. 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.

Column 1 reports the IV results from estimating the single-instrument specification (equation (45)), which links population density, measured in thousand inhabitants per square km, to (instrumented) normalized shape. The corresponding first stage is reported in column 1 of Table 2. Column 2 reports the IV results from estimating the double-instrument specification (equation (42)), which links population to city area and shape, separately instrumented for. The corresponding first stage is reported in columns 2 and 3 of Table 2. Column 3 reports the corresponding OLS estimates.

Recall that normalized shape metrics capture the departure of a city’s shape from an ideal circular shape and are invariant to city area, higher values implying longer trips. The IV estimates of the single-instrument specification indicate that less compact cities are associated with a decline in population density. The magnitudes of this effect are best understood in terms of standardized coefficients. Consider the remoteness index (panel A), representing the length, in km, of the average trip to the footprint’s centroid. A one-standard deviation increase in normalized remoteness (0.06) is associated with a decline in population density of 0.9 standard deviations.

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 precise 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 or in high-altitude areas. 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. Both the first-stage (columns 1, 2, 4, and 5) and the IV estimates of the effect of shape on population (columns 3 and 6) are minimally affected by excluding these cities. This shows that my instrument has explanatory power also in cities without extreme topographic constraints,³⁰ 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 1 and 2 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. A natural question then arises as to whether we can put a price on "good shape". As discussed in Section 3, the Rosen-Roback model

³⁰Recall 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.

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 I discuss in Section 4.5, 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. I cope with these cases following three different approaches: (i) ignore the issue and keep in the sample all cities; (ii) drop from the sample districts with more than one city, thus restricting my sample to cities that have a one-to-one correspondence with districts; (iii) include in the sample only the largest city in each district. Tables 5 and 6 report estimation results obtained from all three approaches.

[Insert Table 5]

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, 4, and 7 report the IV results from estimating the single-instrument specification (equation (47)), that does not explicitly control for city area. Columns 2, 5, and 8 report the IV results from estimating the double-instrument specification (equation (42)), 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, 7, and on the city-specific one in columns 2, 5, 8 – 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, 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 reduced-form relationship between area and wages is negative, as predicted by the conceptual framework (condition (17))

in Section 3).

[Insert Table 6]

Table 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 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 reduced-form relationship between city area and rents is also negative, consistent with the conceptual framework (condition (18) in Section 3).

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 3 (eq. (38)), 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 -0.518 . 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 -0.14 . Recall that a one standard deviation increase in disconnection for the average-sized city is about 360 meters. Interpreting this as commuting length, an interpretation substantiated by the evidence in Section 6.5, this suggests that an increase in one-way commutes of 360 meters entails a welfare loss equivalent to a 0.05 log points decrease in income.

It is interesting to compare this figure with the *actual* cost of an extra 360 meters in one's daily one-way commute, under different commuting options. 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 45% 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,³¹ the direct cost of increased commute length amounts to 1.3% of the yearly wage, or roughly one quarter of the welfare cost estimated above.

The estimated welfare loss from longer commutes appears to be large, relative to the immediate time and monetary costs of commuting. This is consistent with the interpretation that commuting is perceived as a particularly burdensome activity. The behavioral literature has come to similar conclusions, albeit in the context of developed countries. Stutzer and Frey (2008) find a large and robust negative relation between commuting time and subjective well-being, using German data. They 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. Based on diary-based surveys, designed to measure individuals' emotional experiences and time

³¹These 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.

use, Kahneman et al. (2004) find that commuting is the daily activity that generates the lowest level of "positive affect", as well as a relatively high level of "negative affect". These findings are also consistent with the literature on commuting stress and its impact on psycho-physical health (see Novaco and Gonzalez, 2009, for a review).

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 being a production amenity or disamenity. The effect of S on productivity is pinned down by equation (37) from Section 3:

$$\widehat{\lambda}_A = (1 - \beta - \gamma)\widehat{B}_N + (1 - \gamma)\widehat{B}_W$$

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 -0.098 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, and somewhat sensitive to the calibration of parameters β and γ . Overall, they suggest that city shape is not affecting firms' productivity, and that the cost of disconnected shape is borne mostly by consumers. It is possible that firms optimize against poor urban shape, in a way that consumers cannot.³²

6.5 Channels and Heterogeneous Effects

The results presented so far provide evidence that consumers have a preference for more compact cities. 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

³²Section 7.3 investigates how firms respond to city shape in their location choices *within* cities, by looking at the spatial distribution of employment.

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 (eq. (45)), 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. This issue is investigated explicitly in Section 7.2, in which I present some results on current road infrastructure and city shape. In Table 7, 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. 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 the number of motor vehicles in the state in 1989-1990, drawn from CIER (1990). In columns 4, 5, and 6, I interact shape 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 7, 8, and 9, 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 4, 5, and 6). The interaction between shape and motor vehicles availability appears to be significant when considering

the range index, which captures longer, cross-city trips. Overall, this can be interpreted as suggestive evidence that infrastructure mitigates the negative effects of poor 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 differential trends across cities with different incomes.

Housing Quality

[Insert Table 8]

A complementary question relates to who bears the cost of poor city geometry. 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 the poor can afford (Bertaud, 2004). In Table 8, I investigate the link between poverty and city shape by looking at two - admittedly very rough - proxies for city income: the total and the fraction of slum dwellers and a principal component wealth index based on physical characteristics of houses, including roof, wall, and floor material as well as the availability of running water and a toilet. This information is provided in the Census for a limited number of cities and years. 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. Moreover, less compact cities are characterized by houses of marginally better quality. In principle, this could arise from higher equilibrium rents in compact cities, which could be forcing more people into substandard housing. A more interesting interpretation, however, is that poorer migrants sort into cities with more compact shapes, possibly because of their lack of individual means of transport and consequent higher sensitivity to commute lengths. It is plausible that the marginal migrant to urban India would start out as slum dweller, which substantiates the sorting hypothesis.

7 Empirical Results: Endogenous Responses to City Shape

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, if any, and infrastructural investments best accommodate city growth? This section provides some evidence on the interactions between topography, city shape, and policy.³³

7.1 Floor Area Ratios

I start by considering the most controversial among land use regulations currently in place in urban India: Floor Area Ratios (FAR). As explained in Section 4.6, 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,³⁴ 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 9]

³³Due to data availability constraints, part of the analyses carried out in this section are based on comparatively small samples and/or rely on cross-sectional variation only. These results should therefore be interpreted cautiously.

³⁴Sridhar (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.

In Table 9, 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 re-proposed here, augmented with an interaction between each instrument and FAR levels. Each observation is a 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 seems to suggest 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 9, 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 cities 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).

My next question relates to the determinants of FARs; in particular, whether urban form considerations appear to be incorporated by policy makers in setting FARs. Panel C of Table 9 provides an attempt to investigate this issue, by regressing FAR values on urban form indicators

- shape and area - as measured in year 2005. The estimating equations are cross-sectional versions of equations (47) and (42), with FARs as a dependent variable. Each observation is a city in year 2005. Subject to the limitations of cross-sectional inference and small sample size, these results seem to indicate that larger cities have more restrictive FARs, both in the IV and in the OLS. This is consistent with one of the stated objectives of FARs: curbing densities in growing cities. At the same time, however, the coefficient on shape is positive, both in the IV and in the OLS. This could either indicate a deliberate willingness of policy makers to allow for taller buildings in areas with constrained topographies, or might reflect a historical legacy of taller architecture in more constrained cities.³⁵

7.2 Road Infrastructure

Next, I turn to another type of policy response to city shape: road infrastructure. In Section 6.3 I provide some suggestive evidence that infrastructure might mitigate the adverse effects of poor shape. In this sub-section, I am interested in the complementary question of whether current road infrastructure responds to city shape, and whether denser road networks are compensating for longer potential trips in less compact cities. Although the Census provides urban road density for a subset of cities and years, this data appears to be based on administrative boundaries, which are seldom updated. Instead, I compute the total length of the current road network as it appears on the maps available on Openstreetmap, overlapping the street network with my night-time lights-based urban footprint boundaries. More details can be found in Section 4.6. In Table 10, I investigate the cross-sectional relationship between current road network and city shape, measured in 2010 - the most recent year in my sample.

[Insert Table 10]

Each observation in Table 10 is a city in year 2010. Results are reported for the disconnection (columns 1 to 3) and range index (column 4 to 6). 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 might be more suitable to capture cross-city trips, possibly requiring highways. Panel A considers total road length, whereas panel B

³⁵Discussions with urban local bodies and developers suggest that FARs reflect, to some extent, the traditional architectural style of different cities.

considers road length per capita. Columns 1 and 4 report the IV results from estimating a cross-sectional version of the single-instrument specification (equation (45)). Columns 2 and 5 report the IV results from estimating a cross-sectional version of the double-instrument specification (equation (42)).

These results are somewhat mixed and not always consistent across specifications. However, the general pattern discernible from the IV results seems to be the following: non-compact cities have a road network that is less dense in absolute terms, but road density per person which is not statistically different from that of compact cities.³⁶

7.3 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.³⁷ 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 4.6. 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

³⁶The OLS coefficients display the opposite pattern: more disconnected cities have a denser road network in absolute, but not in per capita terms. This could be generated by the correlation between city size and non-compactness, discussed in Section 6.1. Growing cities are typically less compact, and also tend to have better infrastructure. The observed pattern can be due to the fact that the spurious correlation between non-compactness and population is stronger than that between non-compactness and infrastructure - possibly because in rapidly expanding cities, infrastructure responds to population growth with a delay.

³⁷The 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.

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.³⁸

[Insert Table 11]

In Table 11, I estimate the relationship between number of employment centers, city area, and shape, in a cross-section of footprints observed in 2005. Results are reported for the disconnection (columns 1 to 3) and remoteness index (columns 4 to 6). Recall that these two indexes represent, respectively, the average length of within-city trips, and the average length of trips leading to the footprint centroid. The latter index proxies for the commutes that would be prevalent if the city were predominantly monocentric. Columns 1 and 4 report the IV results from estimating a cross-sectional version of the single-instrument specification (equation (45)). The dependent variable is the number of subcenters per square km. Columns 2 and 5 report the IV results from estimating a cross-sectional version of the double-instrument specification (equation (42)). The dependent variable is the log number of employment subcenters. Columns 3 and 6 report the same specification, estimated by OLS.

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 have more employment subcenters (columns 2 and 5). Interestingly, conditional on city area, less compact cities do not appear to be more polycentric: if anything, longer trips induce a reduction in the number of subcenters. These results suggest that firms prefer to cluster centrally, 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 large welfare losses that poor shape is found to have on consumers (Section 6.4): if employment were as dispersed as population, actual commute trips should be relatively short, regardless of shape, and poor geometry would have negligible effects.

³⁸For 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).

8 Conclusion

In this paper, I use geography-driven variation in city shape to investigate the implications of intra-urban commute length for consumers, in the context of India. I find that poor urban connectivity 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, there is a growing concern that existing policies, in particular urban land use regulations, might directly or indirectly contribute to distorting urban form (Glaeser, 2011; Sridhar, 2010; World Bank, 2013). My findings can inform this policy debate on both fronts. Although this study focuses on geographic obstacles, which are 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.³⁹ 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.

This paper leaves a number of open questions for future research. First, my approach captures the overall effects of urban geometry, measured at the city level, but does not explicitly

³⁹Besides 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.

address the implications of geometry for the spatial equilibrium *within* cities. Although my empirical analysis highlights that more compact cities are generally preferred by the average consumer, it does not answer the question of who bears the costs of disconnected geometry within each city. More detailed, geo-referenced data at the sub-city level might help shed light on these issues. More work is also required to understand the response of firms to city shape. City shape has the potential to affect firms both directly, by increasing transportation costs, and indirectly, by distorting their location choices. These effects are likely to be heterogeneous across sectors, which would raise the complementary question of whether firms in different sectors sort into cities with different geometries.

Finally, it would be interesting to investigate other potential channels through which city shape affects consumers, besides commuting patterns. Throughout this project, I interpret city geometry mainly as a shifter of commuting costs, and I employ shape metrics specifically constructed to capture the implications of shape for transit. However, there could be other, second-order channels through which city shape might have economically-relevant effects. First, 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.⁴⁰ Second, the same topographic obstacles that make cities "disconnected" from an urban transit perspective, might also act as boundaries, promoting the separation of a city in different, disconnected neighborhoods and/or administrative units. This could have implications both in terms of political economy and segregation.⁴¹ Again, more disaggregated data at the sub-city level will be required to investigate these ramifications.

⁴⁰The optimal layout of such networks is founded on different engineering theories than those that found transport, and different shape indicators would be required to capture these aspects of shape. However, there might still be interactions with the shape metrics used in my analysis.

⁴¹Previous research, although in very different contexts, suggests this might be the case: for instance, Hoxby (2000) finds that the number of school districts in US cities is predicted by the number of natural barriers. Ananat (2011) finds that US cities subdivided by railroads into a larger number of physically defined neighborhoods became more segregated.

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] Ananat, E. O. (2011), "The Wrong Side(s) of the Tracks: The Causal Effects of Racial Segregation on Urban Poverty and Inequality", *American Economic Journal: Applied Economics*, 3(2), 34-66.
- [3] Anas, A., R. Arnott, and K. A. Small (1998), "Urban Spatial Structure", *Journal of Economic Literature*, 36 (3), 1426-1464.
- [4] Angel, S., J. Parent, and D. L. Civco (2009a), "Ten Compactness Properties of Circles: A Unified Theoretical Foundation for the Practical Measurement of Compactness", *The Canadian Geographer*, 54 (4), 441-461.
- [5] Angel, S., J. Parent, and D. L. Civco (2009b), "Shape Metrics", ESRI working paper.
- [6] Anselin, L. (1995), "Local Indicators of Spatial Association - LISA", *Geographical Analysis*, 27 (2), 93-115.
- [7] Bajari, P. and M. E. Kahn (2004), "The Private and Social Costs of Urban Sprawl: the Lot Size Versus Commuting Tradeoff ", working paper.
- [8] 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.
- [9] Baum-Snow, N., L. Brandt, V. Henderson, M. Turner, Q. Zhang (2013), "Roads, Railroads and Decentralization of Chinese Cities", working paper.
- [10] Baum-Snow, N. and M. Turner (2012), "Transportation and the Decentralization of Chinese Cities", working paper.
- [11] 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.
- [12] Bertaud, A. (2002a), "The Economic Impact of Land and Urban Planning Regulations in India", working paper, http://alainbertaud.com/wp-content/uploads/2013/06/AB_-India_-Urban_Land_Reform.pdf
- [13] Bertaud, A. (2002b), "Note on Transportation and Urban Spatial Structure", Washington, ABCDE conference, working paper.
- [14] Bertaud, A. (2004), "The Spatial Organization of Cities: Deliberate Outcome or Unforeseen Consequence?", working paper.
- [15] Bertaud, A. and J. K. Brueckner (2005), "Analyzing Building-Height Restrictions: Predicted Impacts and Welfare Costs", *Regional Science and Urban Economics*, Elsevier, 35 (2), 109-125.
- [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] Dev, S. (2006), "Rent Control Laws in India: A Critical Analysis", CCS Working Paper No. 158, Centre for Civil Society, New Delhi.
- [24] Fernandes, A. and G. Sharma (2012), "Determinants of Clusters in Indian Manufacturing: The Role of Infrastructure, Governance, Education, and Industrial Policy", IGC working paper.
- [25] Giuliano, G. and K. A. Small (1993), "Is the Journey to Work Explained by Urban Structure?", *Urban Studies*, 30 (9), 1485-1500.
- [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] Gordon, P., A. Kumar and H. Richardson (1989), "The Influence of Metropolitan Spatial Structure on Commuting Time", *Journal of Urban Economics*, 26 (2), 138-51.
- [30] Gyourko, J., M. Kahn, and J. Tracy (1999), "Quality of Life and Environmental Comparisons", in E.S. Mills and P. Cheshire (Eds.), *The Handbook of Regional and Urban Economics*, v.3, North Holland Press.
- [31] Greenstone, M., and R. Hanna (forthcoming), "Environmental Regulations, Air and Water Pollution, and Infant Mortality in India", *American Economic Review*.
- [32] 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.
- [33] Henderson, V., A. Storeygard, and D. N. Weil (2012). "Measuring Economic Growth from Outer Space", *American Economic Review*, 102 (2), 994-1028.

- [34] Hoxby, C. (2000), "Does Competition among Public Schools Benefit Students and Taxpayers?", *American Economic Review*, 90 (5), 1209–1238.
- [35] Indian Institute for Human Settlements (2013), "Urban India 2011: Evidence", working paper.
- [36] 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.
- [37] Kahneman, D., A. B. Krueger, D. A. Schkade, N. Schwarz, and A. Stone (2004), "A Survey Method for Characterizing Daily Life Experience: The Day Reconstruction Method", *Science*, 306 (5702), 1776-1780.
- [38] Mc Kinsey Global Institute (2010), "Globalisation and Urban Growth in the Developing World with Special Reference to Asian Countries".
- [39] McMillen, D. P. (2001), "Nonparametric Employment Subcenter Identification", *Journal of Urban Economics* 50 (3), 448-473.
- [40] 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.
- [41] Ministry of Urban Development (2008), *Study on Traffic and Transportation Policies and Strategies in Urban Areas in India*, Government of India: New Delhi.
- [42] Mitra, A. (1980), *Population and Area of Cities, Towns, and Urban Agglomerations, 1872-1971*, Bombay: Allied.
- [43] Mitric, S. and I. Chatterton (2005), "Towards a Discussion of Support to Urban Transport Development in India", Washington, DC: World Bank.
- [44] Morten, M. and J. Oliveira (2014), "Migration, Roads and Labor Market Integration: Evidence from a Planned Capital City", working paper.
- [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] Novaco, R. W. and O. Gonzalez (2009), "Commuting and well-being", in *Technology and well-being*, ed. Y. Amichai-Hamburger. Cambridge: Cambridge University Press.
- [47] Stutzer, A. and B. S. Frey (2008), "Stress that doesn't pay: The commuting paradox", *The Scandinavian Journal of Economics*, 110 (2), 339-366.
- [48] Roback, J. (1982), "Wages, Rents and the Quality of Life", *Journal of Political Economy*, 90 (6), 1257-1278.
- [49] Rosen, S. (1979), "Wage-Based Indexes of Urban Quality of Life", in Mieszkowski, P. and M. Straszheim (Eds.), *Current Issues in Urban Economics*, Baltimore, Johns Hopkins University Press.
- [50] 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.

- [51] Saiz, A. (2010), "The Geographic Determinants of Housing Supply", *Quarterly Journal of Economics*, 125 (3), 1253-1296.
- [52] 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.
- [53] 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.
- [54] 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.
- [55] Sridhar, K.S. (2010), "Impact of Land Use Regulations: Evidence From India's Cities", *Urban Studies* 47 (7), 1541-1569.
- [56] Storeygard, A. (2014), "Farther on Down the Road: Transport Costs, Trade and Urban Growth in Sub-Saharan Africa", working paper.
- [57] Tse, C.Y., A. W. H. Chan (2003), "Estimating the Commuting Cost and Commuting Time property Price Gradients", *Regional Science and Urban Economics*, 33 (6), 745-767.
- [58] 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.
- [59] U.S. Energy Information Administration (2014), Report on India, <http://www.eia.gov/countries/analysisbriefs/India/>.
- [60] Van Ommeren, J. and M. Fosgerau (2009), "Workers' Marginal Costs of Commuting", *Journal of Urban Economics*, 65 (1), 38-47.
- [61] World Bank (2013), "Urbanization Beyond Municipal Boundaries: Nurturing Metropolitan Economies and Connecting Peri-urban Areas in India", Washington D.C: The World Bank.
- [62] 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.⁴² 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:

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

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 \hat{y}_i the estimate of y for cell i , and with $\hat{\sigma}_i$ the corresponding standard error, candidate subcenters are grid cells such that $(y_i - \hat{y}_i) / \hat{\sigma}_i > 1.96$.

The second stage of the procedure selects those locations, among candidate subcenters, that have significant explanatory power in a semiparametric employment density function estimation.

⁴²The 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.

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:

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

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 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

	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)
	OLS	OLS	OLS
	Norm. shape of actual footprint	Shape of actual footprint, km	Log area of actual footprint, km ²
A. Shape Metric: Remoteness			
Norm. shape of potential footprint	0.0413*** (0.0141)		
Shape of potential footprint, km		0.352*** (0.0924)	0.0562*** (0.0191)
Log projected historic population		0.390** (0.160)	0.488*** (0.102)
Observations	6,276	6,276	6,276
B. Shape Metric: Spin			
Norm. shape of potential footprint	0.0249** (0.0103)		
Shape of potential footprint, km		0.585*** (0.213)	-7.43e-05 (0.000760)
Log projected historic population		2.470 (2.456)	0.571*** (0.101)
Observations	6,276	6,276	6,276
C. Shape Metric: Disconnection			
Norm. shape of potential footprint	0.0663*** (0.0241)		
Shape of potential footprint, km		1.392*** (0.229)	0.152*** (0.0457)
Log projected historic population		-1.180*** (0.271)	0.307*** (0.117)
Observations	6,276	6,276	6,276
D. Shape Metric: Range			
Norm. shape of potential footprint	0.0754** (0.0344)		
Shape of potential footprint, km		1.935*** (0.357)	0.0641*** (0.0212)
Log projected historic population		-4.199*** (0.977)	0.310*** (0.117)
Observations	6,276	6,276	6,276
Model for \hat{f}	common rate	city-specific	city-specific
City FE	YES	YES	YES
Year FE	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), (2), (3) report the results from estimating respectively equations (46), (43), (44). The dependent variables are normalized shape (dimensionless), shape, in km, and log area, in km², 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 (2), and on a city-specific one in column (3) - see Section 5.1. 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)
	IV	IV	OLS
	Population density	Log population	Log population
A. Shape Metric: Remoteness			
Norm. shape of actual footprint	-311.8*** (92.18)		
Shape of actual footprint, km		-0.137** (0.0550)	0.0338*** (0.0112)
Log area of actual footprint, km ²		0.785*** (0.182)	0.167*** (0.0318)
Observations	1,440	1,440	1,440
B. Shape Metric: Spin			
Norm. shape of actual footprint	-110.9*** (37.62)		
Shape of actual footprint, km		-0.00101 (0.000657)	0.000887*** (0.000288)
Log area of actual footprint, km ²		0.547*** (0.101)	0.197*** (0.0290)
Observations	1,440	1,440	1,440
C. Shape Metric: Disconnection			
Norm. shape of actual footprint	-254.6*** (80.01)		
Shape of actual footprint, km		-0.0991** (0.0386)	0.0249*** (0.00817)
Log area of actual footprint, km ²		0.782*** (0.176)	0.167*** (0.0318)
Observations	1,440	1,440	1,440
D. Shape Metric: Range			
Norm. shape of actual footprint	-84.14*** (27.70)		
Shape of actual footprint, km		-0.0284*** (0.0110)	0.00763*** (0.00236)
Log area of actual footprint, km ²		0.746*** (0.164)	0.171*** (0.0305)
Observations	1,440	1,440	1,440
Model for \hat{r}	common rate	city-specific	
City FE	YES	YES	YES
Year FE	YES	YES	YES

Notes: this table reports estimates of the relationship between city shape and population. Each observation is a city-year. Column (1) reports the results from estimating equation (45) (single-instrument specification). The dependent variable is population density, in thousands of inhabitants per km². The explanatory variable is normalized shape (dimensionless). Column (2) reports the results from estimating equation (42) (double-instrument specification). The dependent variable is log city population. The explanatory variables are log city area, in km², and city shape, in km. In columns (1) and (2), estimation is by IV. The instruments are discussed in Section 5.2, and the corresponding first-stage estimates are reported in Table 2. The construction of the potential footprint is based on a common rate model for city expansion in columns (1) and (2), and on a city-specific one in column (3) - see Section 5.1. Column (3) reports the same specification as column (2), 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

	(1)	(2)	(3)	(4)	(5)	(6)
	FS(1)	FS(2)	IV	FS(1)	FS(2)	IV
	Shape of actual footprint, km	Log area of actual footprint, km ²	Log population	Shape of actual footprint, km	Log area of actual footprint, km ²	Log population
Shape of potential footprint, km	1.365*** (0.228)	0.155*** (0.0462)		1.376*** (0.233)	0.158*** (0.0488)	
Log projected historic population	-1.201*** (0.267)	0.278** (0.118)		-1.213*** (0.276)	0.294** (0.123)	
Shape of actual footprint, km			-0.100** (0.0414)			-0.109** (0.0428)
Log area of actual footprint, km ²			0.776*** (0.190)			0.796*** (0.185)
Observations	6,006	6,006	1,373	5,996	5,996	1,385
Cities	433	433	433	440	440	440
Sample excludes	coastal cities	coastal cities	coastal cities	mountainous cities	mountainous cities	mountainous cities
City FE	YES	YES	YES	YES	YES	YES
Year FE	YES	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), (2) and (4), (5) are equivalent to columns (3), (4) in Table 1. They report OLS estimates of the first-stage relationship between city shape and area and the instruments discussed in Section 5.1. Columns (3) and (6) are equivalent to column (2) 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. Columns (1) to (3) exclude from the sample cities located within 5 km from the coast. Columns (4) to (6) 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)	(7)	(8)	(9)
	IV	IV	OLS	IV	IV	OLS	IV	IV	OLS
	<i>All districts</i>			<i>Only districts with one city</i>			<i>Only top city per district</i>		
	<i>Dependent variable: log wage</i>								
Shape of actual footprint, km	0.109*** (0.0275)	0.0381 (0.0386)	0.0538*** (0.0172)	0.0996*** (0.0336)	0.0626 (0.0536)	0.0586*** (0.0150)	0.112*** (0.0300)	0.0600* (0.0360)	0.0512*** (0.0181)
Log area of actual footprint, km ²		-0.409 (0.390)	-0.0478 (0.0356)		-0.167 (0.465)	-0.00936 (0.0516)		-0.337 (0.392)	-0.0308 (0.0478)
Observations	2,009	2,009	2,009	1,075	1,075	1,075	1,517	1,517	1,517
Model for \hat{f}	common rate	city-specific		common rate	city-specific		common rate	city-specific	
City FE	YES	YES	YES	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	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 km². Columns (1), (4), (7) report the results from estimating equation (47) (single-instrument specification) by IV. Columns (2), (5), (8) report the results from estimating equation (42) (double-instrument specification) by IV. Instruments are described in Section 5.2. Columns (3), (6), (9) 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), (4) and (7), and on a city-specific one in columns (2), (5), (8) – see Section 5.1. In columns (4), (5), (6) the sample is restricted to districts containing only one city. In columns (7), (8), (9), the sample is restricted to cities that are either the only ones or the top ones in their respective districts. 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)	(7)	(8)	(9)
	IV	IV	OLS	IV	IV	OLS	IV	IV	OLS
	<i>All districts</i>			<i>Only districts with one city</i>			<i>Only top city per district</i>		
A. Dependent variable: log yearly rent per square meter									
Shape of actual footprint, km	-0.303	-0.596	0.00204	-0.636	-0.518*	-0.00857	-0.198	-0.684	0.0107
	(0.758)	(0.396)	(0.0505)	(1.661)	(0.285)	(0.0736)	(0.806)	(0.538)	(0.0516)
Log area of actual footprint, km ²		-1.529	-0.0176		-0.919	-0.0632		-1.407	-0.0355
		(1.409)	(0.0847)		(0.870)	(0.108)		(1.582)	(0.0962)
Observations	895	895	895	476	476	476	711	711	711
B. Dependent variable: log yearly rent per square meter, upper 50%									
Shape of actual footprint, km	-0.274	-0.675	-0.00526	-0.595	-0.594*	-0.00736	-0.143	-0.790	0.00685
	(0.752)	(0.432)	(0.0577)	(1.626)	(0.324)	(0.0825)	(0.780)	(0.604)	(0.0590)
Log area of actual footprint, km ²		-1.649	-0.00333		-1.040	-0.0485		-1.586	-0.0141
		(1.542)	(0.0976)		(0.990)	(0.125)		(1.786)	(0.110)
Observations	895	895	895	476	476	476	711	711	711
Model for \hat{r}	common rate	city-specific		common rate	city-specific		common rate	city-specific	
City FE	YES	YES	YES	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	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 m² 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 m². The explanatory variables are city shape, in km, and log city area, in km². Columns (1), (4), (7) report the results from estimating equation (47) (single-instrument specification) by IV. Columns (2), (5), (8) report the results from estimating equation (42) (double-instrument specification) by IV. Instruments are described in Section 5.2. Columns (3), (6), (9) 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), (4) and (7), and on a city-specific one in columns (2), (5), (8) – see Section 5.1. In columns (4), (5), (6), the sample is restricted to districts containing only one city. In columns (7), (8), (9) the sample is restricted to cities that are either the only ones or the top ones in their respective districts. 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

<i>Dependent variable: population density</i>									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	IV	IV	IV	IV	IV	IV	IV	IV	IV
A. Shape Metric: Disconnection									
Norm. shape	-303.1**	-286.2*	-177.2**	-327.6***	-306.2***	-204.8***	-275.6***	-254.3***	-136.3***
	(154.0)	(149.4)	(78.11)	(103.8)	(98.46)	(54.16)	(93.29)	(89.68)	(45.02)
Norm. shape x state motor vehicles, 1990	0.000117	0.000113	0.000138						
	(0.000416)	(0.000394)	(0.000189)						
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	1,140	1,140	1,140	1,205	1,205	1,205	1,199	1,199	1,199
B. Shape Metric: Range									
Norm. shape	-119.8**	-113.8**	-59.04***	-97.25***	-91.21***	-54.40***	-84.00***	-78.69***	-35.86***
	(48.61)	(48.52)	(18.15)	(33.05)	(31.33)	(13.49)	(28.86)	(28.23)	(9.572)
Norm. shape x state motor vehicles, 1990	0.000200**	0.000196*	0.000101**						
	(9.88e-05)	(0.000102)	(4.18e-05)						
Norm. shape x urban road density, 1981				0.513**	0.487**	0.397***			
				(0.255)	(0.240)	(0.126)			
Norm. shape x state urban road density, 1991							2.75e-05	6.33e-05	0.000259**
							(0.000132)	(0.000125)	(0.000108)
Observations	1,140	1,140	1,140	1,205	1,205	1,205	1,199	1,199	1,199
Model for \hat{r}	common rate	common rate	city-specific	common rate	common rate	city-specific	common rate	common rate	city-specific
City FE	YES	YES	YES	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES	YES	YES	YES
Year FE x Banks in 1981	NO	YES	YES	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 (45) (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 km². The construction of the potential footprint is based on a common rate model for city expansion in columns (1), (2), (4), (5), (7), (8) and on a city-specific one in columns (3), (6) and (9) – see Section 5.1. 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 the number of motor vehicles in the state in 1989-1990 (source: Industrial Data Book 2002-03). In columns (4), (5) and (6) normalized shape is interacted with urban road density in the city, as reported in the 1981 Census. In columns (7), (8), and (9), 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), (6), (8) and (9) 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 Housing Quality

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	IV	IV	OLS	IV	IV	OLS	IV	IV	OLS
	<i>Log slum population</i>			<i>Log slum population share</i>			<i>Housing index</i>		
Shape of actual footprint, km	-0.0431 (0.0437)	-0.134** (0.0540)	-0.0353** (0.0167)	-0.0501*** (0.0185)	-0.113** (0.0561)	-0.0498*** (0.0169)	0.174* (0.0940)	0.195* (0.118)	-0.00698 (0.0312)
Log area of actual footprint, km ²		0.592 (0.492)	0.134 (0.0829)		0.264 (0.633)	0.0501 (0.0938)		0.438 (1.483)	0.141 (0.120)
Observations	1,067	1,067	1,067	1,067	1,067	1,067	822	822	822
Model for \hat{r}	common rate	city-specific		common rate	city-specific		common rate	city-specific	
City FE	YES	YES	YES	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES	YES	YES	YES

Notes: this table reports estimates of relationship between city shape and proxies for housing quality. Each observation is a city-year. Columns (1), (4), (7) report the results from estimating equation (47) (single-instrument specification) by IV. Columns (2), (5), (8) report the results from estimating equation (42) (double-instrument specification) by IV. Columns (3), (6), (9) 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. The construction of the potential footprint is based on a common rate model for city expansion in columns (1), (4), and (7), and on a city-specific one in columns (2), (5), (8) – 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)), and a housing quality principal component index (columns (7), (8), (9)). Data on slums is drawn from the 1991, 2001, and 2011 Census. The housing quality index is the city-year average of a principal component index based on the following characteristics of Census houses: roof, wall, and floor material, availability of running water and toilet in-premise. The source is the 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: Urban Form and Floor Area Ratios

A. First-stage impact of FARs on city shape						
	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	OLS	OLS	OLS	OLS	OLS
	Norm. shape of actual footprint	Shape of actual footprint, km	Log area of actual footprint, km ²	Norm. shape of actual footprint	Shape of actual footprint, km	Log area of actual footprint, km ²
Norm. shape of potential footprint	0.435*** (0.123)			0.337*** (0.117)		
Norm. shape of potential footprint x FAR	-0.0908** (0.0452)			-0.0571 (0.0371)		
Log projected historic population		2.998 (2.758)	1.984** (0.795)		1.322 (1.841)	1.086* (0.647)
Log projected historic population x FAR		-1.958* (1.023)	-0.686** (0.319)		-1.315* (0.779)	-0.342 (0.270)
Shape of potential footprint, km		0.156 (1.182)	-0.186 (0.233)		0.235 -0.671	0.0928 (0.187)
Shape of potential footprint, km x FAR		0.665 (0.487)	0.135 (0.106)		0.617** (0.279)	0.0244 (0.0810)
Observations	1,183	1,183	1,183	1,183	1,183	1,183
Model for \hat{f}	common rate	city-specific	city-specific	common rate	city-specific	city-specific
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
Shape of actual footprint, km		-0.0979 (0.124)	0.0509 (0.0312)		-0.271** (0.122)	0.0283 (0.0285)
Shape of actual footprint, km x FAR		0.00290 (0.0472)	-0.0160 (0.0129)		0.0688* (0.0371)	-0.00767 (0.0103)
Log area of actual footprint, km ²		0.683** (0.338)	0.0109 (0.128)		0.828*** (0.308)	-0.0443 (0.103)
Log area of actual footprint, km ² x FAR		0.0995 (0.116)	0.0665 (0.0471)		0.00166 (0.0886)	0.0981*** (0.0368)
Norm. shape of actual footprint	-97.49 (102.5)			-85.46 (108.1)		
Norm. shape of actual footprint x FAR	-10.91 (43.88)			-16.89 (45.12)		
Observations	252	252	252	252	252	252
Model for \hat{f}	common rate	city-specific		common rate	city-specific	
FAR	average	average	average	residential	residential	residential
City FE	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES

Notes: 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. 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 km². 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 (continued): Urban Form and Floor Area Ratios

C. FARs Determinants, 2005						
	(1)	(2)	(3)	(4)	(5)	(6)
	IV	IV	OLS	IV	IV	OLS
	Avg. FAR			Residential FAR		
Shape of actual footprint, km	0.000290 (0.00741)	0.0508 (0.0409)	0.021 (0.0140)	0.00515 (0.00914)	0.0715* (0.0401)	0.0439** (0.0199)
Log area of actual footprint, km ²		-0.190 (0.179)	-0.105* (0.0540)		-0.280 (0.176)	-0.177** (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 (47) (single-instrument specification), with log FARs as a dependent variable, estimated by IV. Columns (2) and (5) estimate a cross-sectional version of equation (42) (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.

Table 10: Impact of City Shape on Road Network, 2010

	(1)	(2)	(3)	(4)	(5)	(6)
	IV	IV	OLS	IV	IV	OLS
A. Total Roads						
	Road density, km/km ²	Log road length, km	Log road length, km	Road density, km/km ²	Log road length, km	Log road length, km
Norm. shape of actual footprint	9.131 (16.44)			-3.563 (3.816)		
Shape of actual footprint, km		-0.0769* (0.0429)	0.0231** (0.0111)		-0.0237** (0.0109)	0.00655** (0.00320)
Log area of actual footprint, km ²		1.292*** (0.147)	0.861*** (0.0415)		1.279*** (0.115)	0.871*** (0.0379)
Observations	429	429	429	429	429	429
B. Roads per capita						
	Road density per capita	Log road length per capita	per Log road length per capita	Road density per capita	Log road length per capita	Log road length per capita
Norm. shape of actual footprint	-0.0377 (0.0785)			-0.0152 (0.0211)		
Shape of actual footprint, km		0.0403 (0.0467)	-0.0677*** (0.0149)		0.00753 (0.0111)	-0.0190*** (0.00427)
Log area of actual footprint, km ²		-0.201 (0.178)	0.346*** (0.0597)		-0.143 (0.138)	0.317*** (0.0547)
Observations	429	429	429	429	429	429
Model for \hat{r}	common rate	city-specific		common rate	city-specific	
Shape metric	disconnection	disconnection	disconnection	range	range	range

Notes: This table estimates the cross-sectional relationship between city shape, city area, and road network as of year 2010. Panel A considers total road length and panel B considers road length per capita. Each observation is a city in year 2010. Columns (1) and (4) estimate a cross-sectional version of equation (45) (single-instrument specification), estimated by IV. In columns (1), (4), panel A, the dependent variable is road density in the footprint, in km per km². In columns (1), (4), panel B, the dependent variable is road density per capita, in in km per km² per thousand inhabitants. Columns (2) and (5) estimate a cross-sectional version of equation (42) (double-instrument specification), estimated by IV. Columns (3) and (6) present the same specification, estimated by OLS. In columns (2), (3), (5), (6), panel A, the dependent variable is log road length in the footprint, in km. In columns (2), (3), (5), (6), panel B, the dependent variable is log road length per capita in the footprint, in km per thousand inhabitants. The construction of the potential footprint is based on a city-specific model for city expansion in columns (1) and (4), and on a common rate one in columns (2) and (5) – see Section 5.1. Road length is computed based on current (2014) Openstreetmap maps of Indian cities, focusing on roads denoted as “trunk”, “primary”, or “secondary”. Disconnection is the average length of trips within the city footprint, in km. Range is the longest possible within-footprint trip, in km. City shape and area are calculated from maps constructed from the DMSP/OLS Night-time Lights dataset, in year 2010. Standard errors are clustered at the city level. *** p<0.01, ** p<0.05, * p<0.1.

Table 11: Impact of City Shape on the Number of Employment Subcenters, 2005

	(1)	(2)	(3)	(4)	(5)	(6)
	IV	IV	OLS	IV	IV	OLS
	Subcenters /km ²	Log subcenters	Log subcenters	Subcenters /km ²	Log subcenters	Log subcenters
Norm. shape of actual footprint	-0.371 (0.507)			-0.563 (2.632)		
Shape of actual footprint, km		-0.0639* (0.0379)	-0.0579*** (0.0154)		-0.0860 (0.0763)	-0.0808*** (0.0214)
Log area of actual footprint, km ²		0.611*** (0.125)	0.571*** (0.0568)		0.608*** (0.165)	0.574*** (0.0575)
Observations	188	188	188	188	188	188
Model for \hat{f}	common rate	city-specific		common rate	city-specific	
Shape metric	disconnection	disconnection	disconnection	remoteness	remoteness	remoteness

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 (columns (1) and (4)), and the log number of employment subcenters (columns (2), (3), (5), (6)). Columns (1) and (4) estimate a cross-sectional version of equation (45) (single-instrument specification), estimated by IV. Columns (2) and (5) estimate a cross-sectional version of equation (42) (double-instrument specification), estimated by IV. Columns (3) and (6) present 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. The shape metrics considered are the disconnection index (columns (1), (2), (3)) and remoteness index (columns (4), (5), (6)). Disconnection is the average length of trips within the city footprint, in km. Remoteness is the average length of trips to the centroid, 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.

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

	(1) IV	(2) IV
	Population density	Log population
Norm. shape of actual footprint	-348.5*** (118.0)	
Shape of actual footprint, km		-0.194*** (0.0596)
Log area of actual footprint, km ²		0.973*** (0.222)
Observations	1,440	1,440
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) and (2) in Table 3. Column (1) reports IV estimates of the relationship between normalized city shape and population density, in thousand inhabitants per km². 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)	(5)	(6)
	IV	IV	IV	IV	IV	IV
	<i>All districts</i>		<i>Only districts with one city</i>		<i>Only top city per district</i>	
A. Dependent variable: log wage						
Shape of actual footprint, km	0.0989*** (0.0272)	0.0116 (0.0420)	0.108** (0.0421)	0.0393 (0.0768)	0.102*** (0.0298)	0.0381 (0.0386)
Log area of actual footprint, km ²		-0.358 (0.395)		-0.164 (0.465)		-0.274 (0.379)
Observations	2,009	2,009	1,075	1,075	1,517	1,517
B. Dependent variable: log yearly rent per square meter						
Shape of actual footprint, km	-0.634 (1.732)	-0.739+ (0.487)	-1.815 (6.682)	-0.670* (0.359)	-0.454 (1.799)	-0.839 (0.664)
Log area of actual footprint, km ²		-1.301 (1.355)		-0.619 (0.807)		-1.092 (1.444)
Observations	895	895	476	476	711	711
Model for \hat{r}	common rate	city-specific	common rate	city-specific	common rate	city-specific
City FE	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES
Initial shape x year FE	YES	YES	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 m² in the city's district, in 2014 Rupees. Columns (1), (3), (5) report the results from estimating equation (47) (single-instrument specification) by IV. Columns (2), (4), (6) report the results from estimating equation (42) (double-instrument specification) by IV. Instruments are described in Section 5.2. The construction of the potential footprint is based on a common rate model for city expansion in columns (1), (3), (5), and on a city-specific one in columns (2), (4), (6) – see Section 5.1. In columns (3), (4), the sample is restricted to districts containing only one city. In columns (5), (6), the sample is restricted to cities that are either the only ones or the top ones in their respective districts. 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.

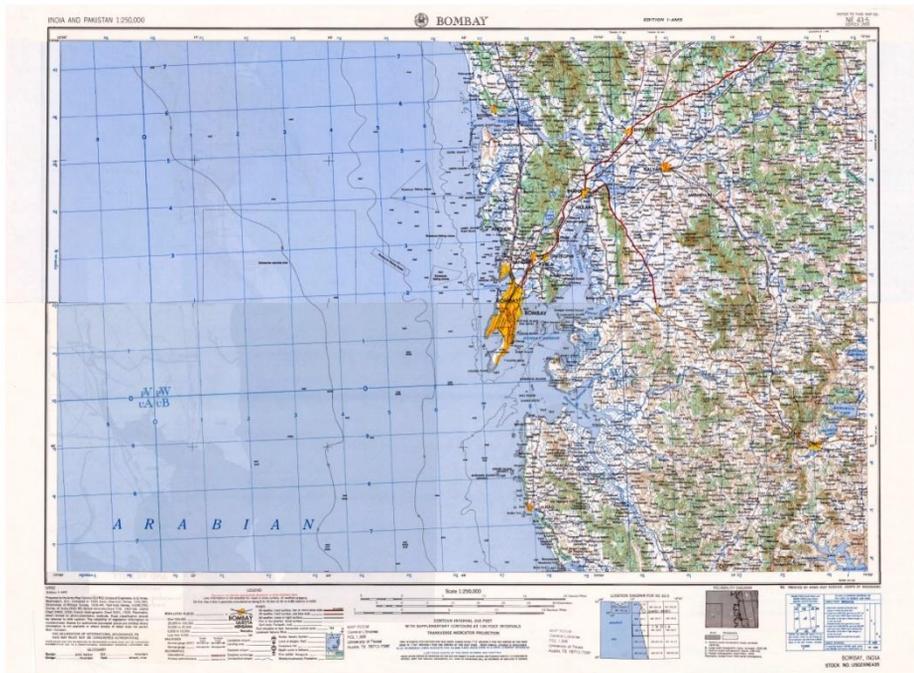


Figure 1A

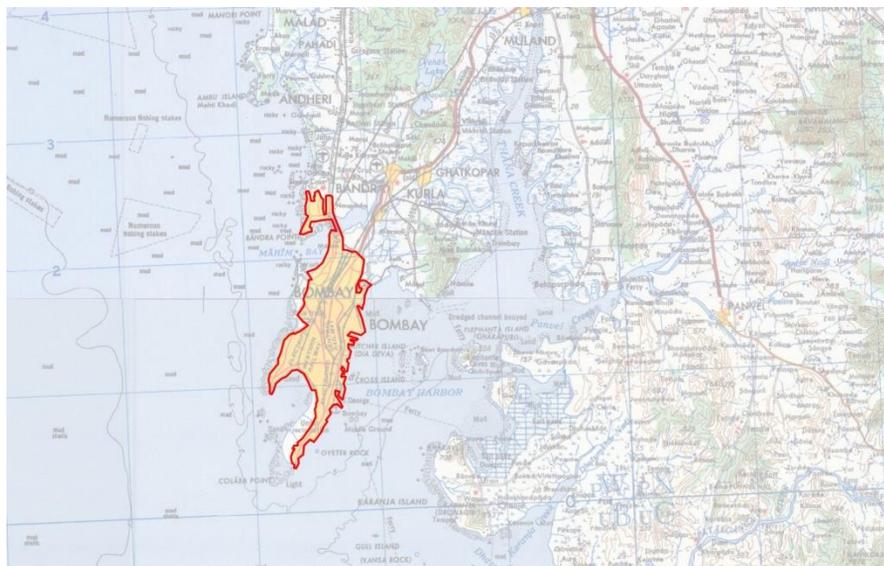


Figure 1B

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

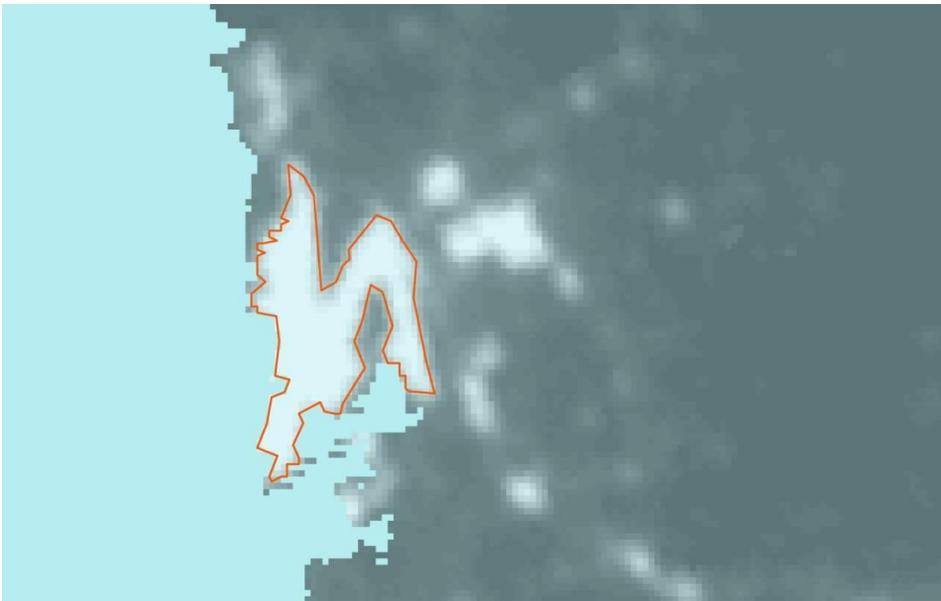
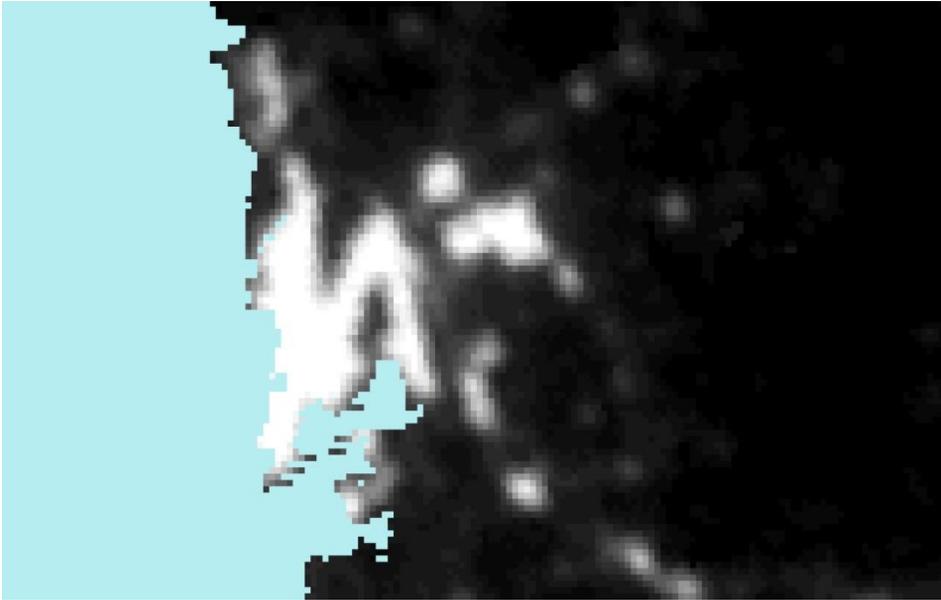
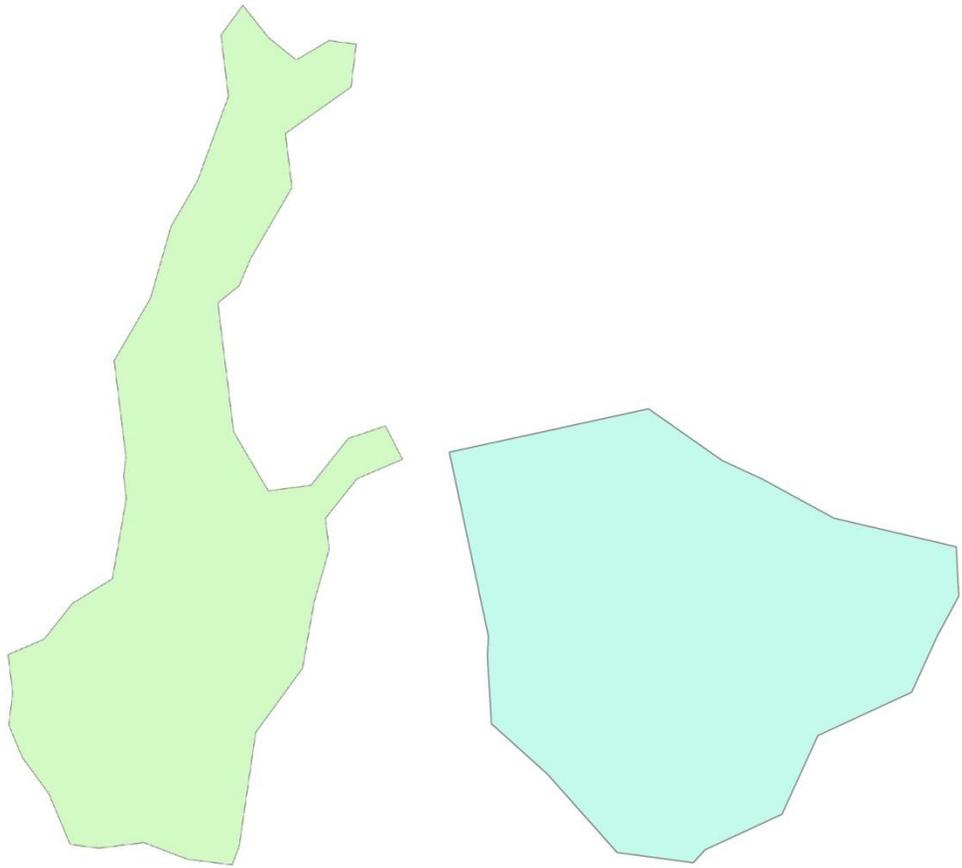
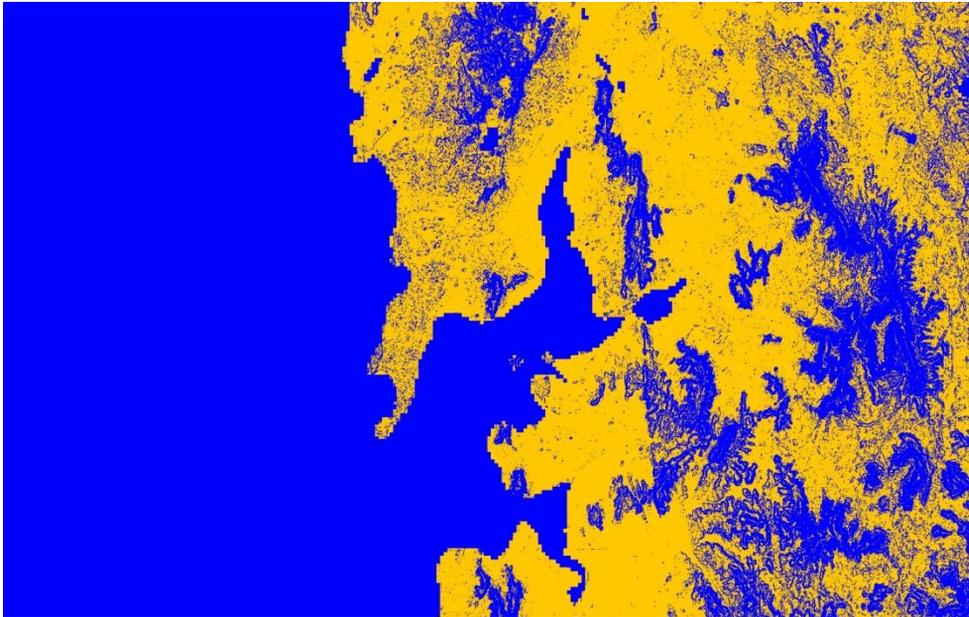


Figure 2
DMS/OLS nighttime lights, year 1992, luminosity threshold : 40.



Shape metric	Kolkata		Bengaluru	
		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

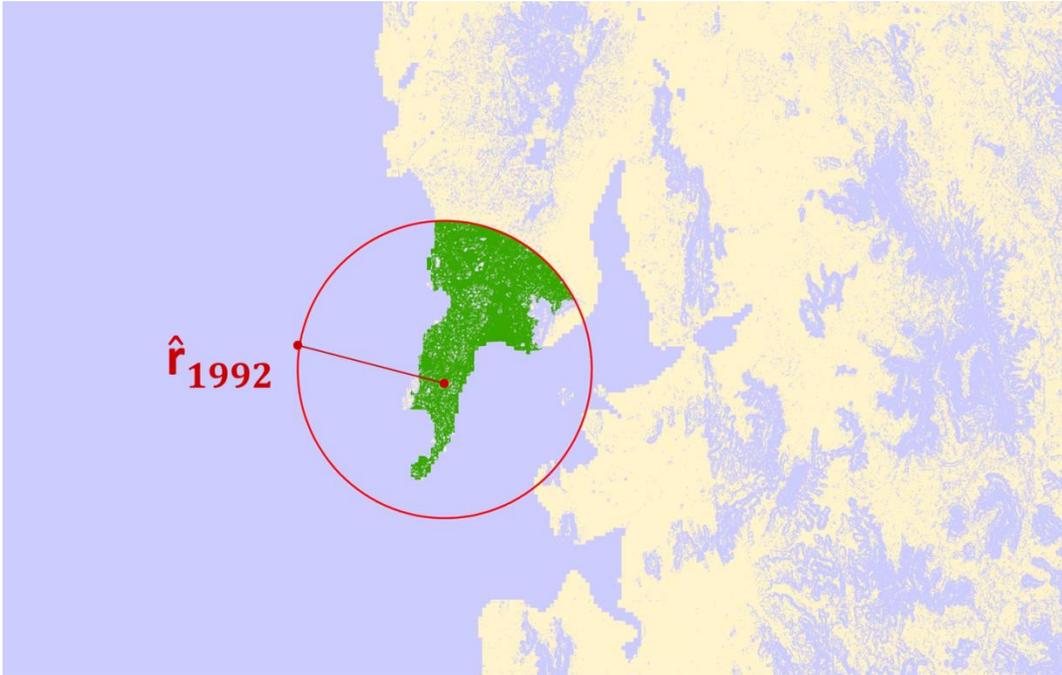


Figure 5c



Figure 5d

Figure 5
Instrument construction