IN SEARCH OF A SPATIAL EQUILIBRIUM IN THE DEVELOPING WORLD

Douglas Gollin
Martina Kirchberger
David Lagakos

In Search of a Spatial Equilibrium in the Developing World
Douglas Gollin, Martina Kirchberger, and David Lagakos
NBER Working Paper No. 23916
October 2017
JEL No. E0,O11,O18

ABSTRACT

In most developing countries, there is a large gap in average consumption per capita between urban and rural areas. One appealing interpretation of this gap is that it reflects a spatial equilibrium, in which the higher consumption levels of urban areas are offset by lower non-monetary amenities. In this paper, we draw on new high-resolution evidence to document how non-monetary amenities vary across space within 20 developing countries. We focus on measures of health, public goods, crime and pollution. These vary substantially across locations within countries and can be carefully measured with highly comparable data. We find that in almost all countries, and for almost all measures, the quality of these amenities is non-decreasing in population density. In addition, net internal migration flows are directed toward denser areas in every country. These findings are hard to reconcile with a spatial equilibrium. Instead, they suggest that developing countries are undergoing a reallocation of workers to densely populated areas, consistent with many models of structural change but inconsistent with models that assume a simple static spatial equilibrium.

Douglas Gollin
University of Oxford
Department of International Development
Queen Elizabeth House
3 Mansfield Road
Oxford OX1 3TB
United Kingdom
Douglas.Gollin@qeh.ox.ac.uk

Martina Kirchberger
Department of Economics
Room 3012, Arts Building
Trinity College Dublin
Dublin 2 Ireland
kirchbem@tcd.ie

David Lagakos
Department of Economics, 0508
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093
and NBER
lagakos@ucsd.edu

1. Introduction

Economists have long recognized that there are large income differences across space within countries. A body of recent evidence documents that real urban-rural consumption gaps are particularly large in developing countries (Ferré, Ferreira, and Lanjouw, 2012; Young, 2014), as, similarly, are gaps in average labor productivity between non-agricultural and agricultural workers (Gollin, Lagakos, and Waugh, 2014). In an accounting sense, these gaps are important for understanding why developing countries have such low aggregate income, since a large fraction of workers in developing countries live in rural areas and work in agriculture (Caselli, 2005; Restuccia, Yang, and Zhu, 2008; Vollrath, 2009; McMillan and Rodrik, 2011).

One interpretation of these gaps is that they reflect a *spatial equilibrium*, in which the higher real consumption levels of urban areas are offset by lower non-monetary amenities. A spatial equilibrium – in which utility levels are equalized across locations – is one of the simplest and most appealing concepts that economists have used to analyze how economic agents locate through space (Rosen, 1979; Roback, 1982). The rationale is simple: if any region offered a better bundle of consumption and amenities than the rest, then agents would move into the better region until any arbitrage opportunities were gone. The concept of a spatial equilibrium has proven useful in learning about a wide set of economic phenomena, including the wage and size distribution of U.S. cities (Baum-Snow and Pavan, 2012; Desmet and Rossi-Hansberg, 2013), the dynamics of U.S. manufacturing and service activity (Desmet and Rossi-Hansberg, 2014), and the welfare effects of infrastructure investments (Allen and Arkolakis, 2014; Donaldson and Hornbeck, 2016) and regional development policy (e.g., Kline and Moretti, 2014a). It is no wonder then that Glaeser and Gottlieb (2009, p.3) call a spatial equilibrium the "central theoretical tool" of urban economics.[1]

Almost all papers that assume a spatial equilibrium rely on models to infer the amenity values of different locations, rather than measuring the amenities directly (Glaeser and Gottlieb, 2009). This reflects the intrinsic difficulty of measuring amenities. But in the context of developing countries, this approach is particularly unsatisfying. Given that most developing countries have much higher real average consumption in urban areas, a spatial equilibrium implies that there are some offsetting amenities that are far superior in rural areas. If this is true, researchers and policy makers should know which amenities of urban areas are of so much worse quality. If urban amenities are not worse, then a spatial equilibrium may not be the right interpretation of the significantly higher urban consumption levels in developing countries – with important

---

[1]Numerous other papers assume a spatial equilibrium as part of their analysis. See Kline and Moretti (2014b) and Redding and Rossi-Hansberg (Forthcoming) for highly readable surveys of the recent quantitative literature in urban economics and studies of place-based policies more generally.

implications for policy. For instance, a model with large disamenities of urban life would call for different policy responses than a model in which there are significant frictions associated with rural-urban migration. Distinguishing between these two alternatives requires direct measures of amenities across space.

In this paper, we report on spatial patterns of living standards that we find from constructing a new database of consistently measured amenities covering 20 developing countries in sub-Saharan Africa. We take a broad definition of amenities and include aspects of living standards that are not solely private consumption and that are linked to location. We focus on real measures based on health outcomes, public goods access, crime exposure and pollution exposure.[2] These measures are commonly cited as amenities in the literature on urban and regional economics, they are unlikely to be well captured in existing consumption metrics, and they can be measured consistently within and across countries. To construct our database, we combine nationally representative household survey data, satellite-derived measures of pollution, and population-density data coming from the Gridded Population of the World Version 4 (GPW). The main surveys we employ are the Demographic and Health Surveys (DHS) and Afrobarometer Surveys. The surveys also report the GPS coordinates or location names of each geographic region sampled, which we link to data on population density from the GPW. We draw on global estimates of the spatial distribution of air pollution concentrations for fine particulate matter (PM2.5) and nitrogen dioxide (NO2). An advantage of our approach is that we do not have to rely on reported categories of what is "urban" and "rural." These binary categories often reflect administrative boundaries and cut-offs that are not consistent across space and time. Our linked data allow us to examine in much greater detail how these amenity measures vary across population density within each country, using internationally comparable measures of density.

We find that in almost all countries, and for almost all measures, amenities are non-decreasing in population density. Measures of children's health, which reflect private consumption as well as public health investments and pollution exposure, are increasing in density almost everywhere. Every public good we consider is at least as prevalent in denser areas if not more prevalent. Property crime and violent crime are, unfortunately, high throughout density space in most countries, and marginally higher in urban areas, on average, though with statistically and insignificant differences in most countries. The same is true for the percent of adults reporting fear of crime in their homes and feeling unsafe in their neighborhoods. In terms of magnitudes, we draw on previous studies of willingness to pay to avoid crime to estimate the

---

[2]By "real," we mean that we measure quantities rather than values. In the following section, we discuss how higher housing prices in urban areas are consistent with a spatial equilibrium only if housing quality or consumption levels (or both) are lower in urban areas – hence our analysis of housing quality.

monetary value of the higher crime rates in denser areas. We conclude that in monetary terms, the value of the crime-rate differences between rural and urban areas is dwarfed by the much larger differences in average income.

Perhaps surprisingly, we find that outdoor air pollution is largely unrelated to density in sub-Saharan Africa. Fine particulate matter is most prevalent in areas near the Sahara, due to dust and sand in the air. Even excluding dust and sea salt, PM2.5 concentrations are quite similar in urban areas and rural areas, on average, as are NO2 concentrations, to focus on two widely used measures of air quality. Overall magnitudes are also quite low, particularly relative to other developing countries such as India and China, which we view as a sad commentary on the lack of manufacturing activity in Africa. Arguably more relevant for many Africans is *indoor* air pollution, which comes about from cooking indoors with solid fuels such as coal or wood. In each of our countries, we find that the fraction of households using solid fuels to cook indoors falls sharply with population density. Thus, people in the rural areas of our countries face worse air quality, on average, than their urban counterparts face. Although we have no comparable measures on water pollution, child health indicators – which can be quite sensitive to water-quality issues – are systematically better in densely populated areas.

Our findings provide little support for the existence of a spatial equilibrium in Africa, at least in a simple static sense. With few exceptions, measures of health, public goods, crime and pollution are either similar across density space or strictly better, on average, in denser areas. Previous studies show that real consumption is higher in denser areas; in particular, our data confirm that consumer durables ownership and housing quality are higher, on average, in denser areas. Of course, our results do not imply that a spatial equilibrium does not exist. We did not – and could not – consider all possible amenities.[3] One should view our findings so far as the most careful look to date at the "prime suspects" of amenities mentioned in the literature, rather than as proof of non-existence of a spatial equilibrium, which, even in principle, we could not produce.

To complement our findings for these amenities, we take two further steps. First, we look at data on internal migration rates within our countries. In a spatial equilibrium, net migration should not flow systematically toward denser areas or toward rural areas. What we find is that migration flows are, in fact, strongly oriented toward denser areas in all but one of the 20 countries we consider – consistent with a long-standing literature on rural-urban migration

---

[3]In fact, by definition, there is an "amenity" which is "worse" in denser areas: population density. It is possible that people could simply dislike living in denser areas. Ciccone and Hall (1996), for example, argue that this is a simple and "realistic" description of the data. Still, there are reasons to believe people prefer to be in denser areas. In the developing world, Fafchamps and Shilpi (2009), for example, provide evidence that rural workers dislike the isolation that comes with rural life.

(see, e.g., Young, 2014; Glaeser and Henderson, 2017, and the references therein). Second, we look at some additional "secondary" measures of amenities, moving down the list of suspects; for these, we have fewer countries (and less data). Here, we find one amenity that does seem to decrease with population density: trust in neighbors, though the magnitudes are not overwhelmingly higher in denser areas. For the other secondary metrics that we consider – anxiety, reported lack of food, and reported lack of medicine – we again find worse outcomes in rural areas.

We conclude that the simplest explanation for our findings is that living standards in developing countries are higher, on average in denser areas, and that individuals respond by voting with their feet and migrating to denser areas. In other words, our facts are consistent with a more dynamic notion of spatial equilibrium, in which individuals are not indifferent across space at any point in time but migrate to better locations to close spatial gaps in living standards (Topel, 1986).[4] The fact that we find few permanent amenities of cities that are systemically worse than those in rural areas does not imply that moving costs are small. In fact, many related studies point to large moving costs, such as from loss of social networks (Morten, 2013; Munshi and Rosenzweig, 2016), lack of skill transferability (Bazzi, Gaduh, Rothenberg, and Wong, 2016), poor road infrastructure (Morten and Oliveira, 2017), risk of bad migration outcomes (Harris and Todaro, 1970; Bryan, Chowdhury, and Mobarak, 2014), or policies that restrict migrants' access to public goods (see, e.g., Feler and Henderson, 2011; Brueckner and Lall, 2015). Relatedly, historical evidence points to very persistent transitional dynamics associated with large shocks, such as after the division of Germany post-WWII, when workers reallocated away from border regions for decades (Redding and Sturm, 2008; Ahlfeldt, Redding, Sturm, and Wolf, 2015).

One refinement of the spatial equilibrium theory relies on compositional effects linked to the sorting of heterogeneous agents (Lagakos and Waugh, 2013; Herrendorf and Schoellman, 2014; Young, 2014; Hamory Hicks, Kleemans, Li, and Miguel, 2017). In particular, there could be a spatial equilibrium within narrowly defined education groups, where utility is constant within each education group, and higher for more educated groups. To evaluate this possibility, we look at a subset of our amenity measures, as our data allow, for households within narrowly defined educational groups, including those with no education. We find that amenities are again similar or increasing with density within each education category. This challenges the notion that a spatial equilibrium holds within education groups.

In terms of measurement, our work is close to that of Young (2014), and builds on his innovative use of the DHS surveys. Unlike in his paper, our work covers measures of property and

---

[4]See Eckert and Peters (2017) for an elegant recent formulation of such a model.

violent crime, fear of crime, trust in neighbors, indoor and outdoor air pollution, and public goods like roads, police stations, health clinics and schools. Our paper is also close to that of Chauvin, Glaeser, Ma, and Tobio (2016), who test the spatial equilibrium hypothesis in various forms in Brazil, China and India. They conclude that a spatial equilibrium assumption can be used, "if it is used warily," in Brazil and China, but they reject a spatial equilibrium outright in India. Our data, which all come from countries less developed than India, are consistent with their conclusions for India, suggesting that a spatial equilibrium is not the right description of the world's least developed countries. Finally, our paper builds on recent work measuring components of living standards at a fine level of geographic detail, such as Henderson, Storeygard, and Weil (2012) and Henderson, Storeygard, and Deichman (2017), which use satellite data on night lights to construct proxies for income at a fine level of geographic detail. One way that our work differs is that it tries to isolate the effect of non-monetary amenities, which have not been studied systematically in the developing world. In emphasizing the measurement of amenities, our work parallels that of Albouy (2012), and echoes his conclusion that, once one considers richer spatial data with better measurement, denser areas do not appear to be undesirable places to live.

This paper is structured as follows. Section 2 presents a simple model with a spatial equilibrium and discusses the role of housing prices. Section 3 outlines our data and how we link our outcome measures with population density. Section 4 presents our main findings for amenities by population density. Section 5 looks at amenity measures for households by education group and for rural-urban migrants. Section 6 explores migration flows and our secondary amenities metrics. Section 7 concludes.

2. Simple Spatial Equilibrium Model

In this section, we present a simple spatial model so as to motivate our empirical work. We also use the model to clarify, in particular, how housing prices fit into our analysis, and to explore how including worker heterogeneity may alter our analysis.

Environment: The economy is populated by identical households that each have a utility function $U(c,h,a)$, where c is non-housing consumption, h is housing consumption (or "housing quality"), and a is amenities. The utility function satisfies $\frac{\partial U}{\partial c} \geq 0$, $\frac{\partial U}{\partial h} \geq 0$ and $\frac{\partial U}{\partial a} \geq 0$ for all levels of c, h and a; i.e., utility is everywhere non-decreasing in each input.

The economy is partitioned into J regions. Households are freely mobile and must locate in exactly one region. Regions have three exogenous characteristics: wages, w_j; housing prices, p_j; and amenities a_j. Any household locating in region j earns wage w_j, may consume as

much housing services as it wants at price p_j per unit, and enjoys amenities a_j. Households anywhere may consume non-housing consumption at a normalized price of one. The budget constraint of a household locating in region j is $w_j = c + p_j h$. We make no attempt to clear markets but, rather, focus on the households' choice of where to locate.

Spatial Equilibrium: A spatial equilibrium implies that utility levels are equated across regions. Let this common utility level be denoted $\bar{u}$. The standard approach in the literature is to use the households' (and/or firms') optimality conditions, plus the common utility value $\bar{u}$, to impute amenities, a_j, given data on prices w_j and p_j. In contrast, our approach is to focus on quantities, and to assess whether the model's predictions for amenities are borne out in the data. To this end, we denote the common quantities of non-housing consumption and housing in j by c_j and h_j. Then, note the following basic, almost definitional, properties of the static spatial equilibrium:

1. Consider two regions u and r. If $c_u > c_r$, than $h_u < h_r$ or $a_u < a_r$ or both.

2. Consider two regions u and r. Households do not prefer to migrate from u to r, or vice versa.

Property 1 says simply that if non-housing consumption is higher in one region than in a second region, then either housing consumption or amenities must be higher in the second. No region can have higher quantities of every input into the utility function, or that would violate the assumption that utility is equated across regions. Property 2 says that when offered the choice, households do not systematically prefer another region to the one they are in. This is the heart of the logic of a spatial equilibrium: if households were systematically better off in another region than in their current region, they would migrate to improve their well being. In what follows, we draw on new data on quantities of consumption, amenities and net migration flows to test these two basic predictions of a static spatial equilibrium.

The Role of Housing Prices: How do housing prices fit into this? Suppose that the urban region had higher wages but also higher housing prices. That is, $w_u > w_r$, but $p_u > p_r$. Couldn't this mean there is a spatial equilibrium, even if amenities are the same, – i.e., $a_u = a_r$? If there were a spatial equilibrium, then one of two basic patterns could be found in the quantity data. First, households could use the higher wages in the urban area to obtain higher consumption, but with lower housing consumption. Second, households could have lower consumption but more housing. If both were higher in the urban area then there would not be a spatial equilibrium. The point is that the ubiquity of higher housing prices in high-wage regions is not sufficient to conclude there is a spatial equilibrium. A spatial equilibrium implies that the high-wage, high-price region must have either lower consumption or lower housing, where

these are measured in real terms rather than in value terms. If, instead, as we later find, both housing quality and consumption are higher in the high-wage, high-price denser areas, it must be true that some other amenity is worse in the denser area, or there would not be a spatial equilibrium.

Extension to Heterogeneous Households: The assumption that all households are identical, while quite common, is rather stark. Indeed, many more recent spatial equilibrium models assume that households are heterogeneous in productivity or tastes and sort accordingly. Recent examples include Ahlfeldt, Redding, Sturm, and Wolf (2015), Caliendo, Dvorkin, and Parro (2015) and Faber and Gaubert (2016). Relative to homogeneous-agent spatial equilibrium models, households are not indifferent, and, in general, each household strictly prefers one region. The key is that population shares depend on relative wages, housing prices and amenities, with higher population shares in regions with higher amenities, all else equal.[5]

How would the assumption of heterogeneous households affect our conclusions above? To answer this question, consider the following simple extension of our model based (heavily) on the elegant formulation of Kline and Moretti (2014b). Suppose, for simplicity, that there are two regions: u and r. Households are indexed by i, and the utility of household i is $U(c_j, h_j, a_j) + e_{i,j}$ for region j, where c_j, h_j and a_j are the common consumption, housing and amenities of region j, and $e_{i,j}$ is an idiosyncratic taste term. The $e_{i,j}$ draws are independent and identically distributed across households with mean zero in each region. For tractability, assume further that $e_{i,j}$ are drawn from a Gumbel (extreme value type I) distribution with shape parameter σ.

In this model, the workers that prefer a region u to a second region r are those that have $U(c_u, h_u, a_u) - U(c_r, h_r, a_r) > e_{i,r} - e_{i,u}$, and vice versa. As Kline and Moretti (2014b) show, the Gumbel assumption implies that the fraction of households locating in region u is given by

$$N_u = \Lambda\left(\frac{U(c_u, h_u, a_u) - U(c_r, h_r, a_r)}{\sigma}\right), \tag{1}$$

where $\Lambda(\cdot) \equiv \frac{\exp(.)}{1+\exp(.)}$ is a logistic(0,1) distribution.

In most developing countries, and in each of the African countries studied in the current paper, the majority of the population lives in rural areas. Is there a way to square this observation, plus the higher consumption levels of urban areas, with a spatial equilibrium? In this model, the only way this is possible is if either housing quality or amenities are higher in the *rural* area. To see this, suppose $N_u < 1/2$ as we see empirically. Then it follows from (1) that

---

[5]Of course, in general equilibrium, housing prices may depend on amenities, though since this is not crucial for our argument, we abstract from it here.

$U(c_u, h_u, a_u) < U(c_r, h_r, a_r)$ for any value of the parameter σ. If $c_u > c_r$, it must then be the case that either $h_u < h_r$ or $a_u < a_r$ (or both).

We conclude that adding heterogeneity in household preferences to our model doesn't change our basic conclusion. For there to be a spatial equilibrium with most households located in the rural area, as is empirically the case in most developing countries, amenities still have to be higher in rural areas. We now turn to measuring whether this is the case.

3. Measuring Amenities Across Space

This section describes our data and our method for matching the data on amenities and other real outcome measures with population density.

3.1. Overview of Data Sources

Until recently, measuring amenities and other living standards metrics across space in the developing world was not feasible. Exploiting progress in surveying and mapping technology, we construct a new dataset that spatially links household surveys on health, public goods, crime, mistrust and other reported living standards, and satellite-derived measures of pollution with gridded population density. To do so, we use data from the Demographic and Health Surveys (DHS), the Living Standards Measurement Surveys (LSMS), Afrobarometer and remotely sensed pollution data. The micro surveys are high-quality nationally representative surveys that cover large numbers of households (typically more than 5,000 for DHS and LSMS and 1,200–2,400 for Afrobarometer) in developing countries. The surveys are designed to use consistent methodologies and definitions across countries. The DHS and LSMS data focus on variables related to population, health, and nutrition, while Afrobarometer focuses on attitudes towards democracy and governance, including the availability of public goods and experiences of crimes. Our data on outdoor air pollution come from satellite-derived estimates of fine particulate matter — in particular PM2.5 concentrations from van Donkelaar, Martin, Brauer, and Boys (2015) and NO2 concentrations from Geddes, Martin, Boys, and van Donkelaar (2016). We outline the main choices related to using these data here and in the following section, when we present our results, while referring the reader to the detailed data appendices A - C for further details.

We select countries that satisfy four criteria: (i) the survey was conducted no earlier than 2005; (ii) spatial identifiers of the respondents or clusters were collected and are available; (iii) the country was larger than 50,000 square kilometers; and (iv) the country was classified as a low-income country by the World Bank in 2012 (meaning GNI per capita (Atlas method,

current US$) below $4,126 in 2012.) For the DHS data, we are left with a sample of 276,051 households across 20 African countries, as listed in Tables A.1 and A.5, covering countries with a combined population of about 770 million people. From these data, we get our main metrics of child health (stunting, wasting, anemia, and minimum acceptable diet), indoor air pollution (cook indoors, use of solid fuel for cooking) and measures of durables ownership and housing quality that we use to supplement our main analysis.

3.2. Population Density Measures

To measure population density, we use data from the Gridded Population of the World Version 4 (GPWv4), which provides population density estimates at a resolution of 30 arc-seconds, corresponding to about 1km at the equator (Center for International Earth Science Information Network, 2015). The gridded population data employ a minimal amount of modeling by equally distributing non-spatial population data from censuses among spatial datasets of administrative units (Doxsey-Whitfield, MacManus, Adamo, Pistolesi, Squires, Borkovska, and Baptista, 2015).

One attractive feature of GPWv4 for the purpose of this analysis is that the distribution of population data is transparent and performed without using further auxiliary data. This comes at a cost of a lower resolution than that available in some alternative data sources. For example, one higher resolution dataset is WorldPop, which uses a range of input data and has a resolution of 100m (Linard, Gilbert, Snow, Noor, and Tatem, 2012). For our analysis, one important consideration is that these other input data might introduce circularity in measurement. For example, if nighttime lights data from satellites are used to assign populations to locations, in an effort to allocate population at a finer geographical resolution, then it would hamper our efforts to estimate the relationship between population density and electrification: by construction, higher densities would be associated with higher rates of electrification. We rule out this circularity by using population density data that are not modeled using further input data. The maximum dispersion assumption of GPWv4 within spatial administrative units therefore biases us towards finding no relationship between population density and outcome variables.

The resolution of the census data underlying the GPWv4 varies across countries due to availability of data. Some countries provide their data at the level of the enumeration area, while others share data only at the second administrative level. We restrict our analysis to countries for which the underlying census data have sufficiently high spatial detail, which corresponds approximately to those for which we have data on more than 40 regions per country.

3.3. Spatially Linking Amenity Measures and Density

We next combine the different sources of data step by step. To link the individual data from the DHS and Afrobarometer with population density, we would ideally have the GPS location of each household. The DHS readily collects GPS coordinates for survey clusters, but in order to preserve the anonymity of survey respondents, these have been displaced — i.e., reassigned a GPS location that falls within a specified distance of the actual location. Urban DHS clusters are randomly displaced by 0-2km, and rural clusters are randomly displaced by 0-5km, with one percent of clusters randomly selected to be displaced by up to 10km (Perez-Heydrich, Warren, Burgert, and Emch, 2013). We take into account the random offset of DHS cluster locations when linking DHS GPS data with continuous raster data by taking 5 km buffers around both urban and rural DHS clusters, as suggested by Perez-Heydrich, Warren, Burgert, and Emch (2013). Appendix A provides more detail on this procedure. An important consideration is how representative our samples are across different levels of population density. We discuss the sampling protocol of the surveys in the Appendix and show in Figure A.6 that when we have geo-coded census data, the survey data cover a wide range of densities. All of our results are robust to using the survey weights provided.

Unfortunately, the Afrobarometer did not collect the GPS location for respondents, but the location name was recorded. We develop an algorithm that performs a series of exact and fuzzy matches of location names relying on data from a global gazetteer that contains the latitude and longitude of a location. Depending on the survey round, this involves between 13 and 21 steps in which the village name, district name and region name are sequentially matched against the ascii names of locations, as well as up to four alternative names listed in the gazetteer. To catch misspellings, we perform fuzzy matches based on similar text patterns, using a similarity score of 0.7 and a vectorial decomposition algorithm (3-gram) (Raffo, 2015). Appendix B provides further detail on the matching procedure. Using this algorithm, we are able to geo-locate 92–95 percent of village names in each round.[6] For each location we can then extract the population density value.

Both the pollution data and the population density data are gridded data, making it straightforward to link them. The estimated PM2.5 and NO2 concentrations are available at a resolution of 0.1 decimal degrees (about 10km at the equator), compared to the 30 arc-second resolution of the population data. We construct a fishnet grid of the same resolution of the pollution data (the coarser spatial resolution) and, for each pixel, compute the average pollution measure and the average population density from the GPWv4. PM2.5 is measured in $\mu g/m^3$, while NO2 is

---

[6]Nunn and Wantchekon (2011) manually geo-locate the respondents of the 2005 Afrobarometer round. When we compare their locations with ours, we find that the median distance is 12.5km.

measured in ppb (parts per billion). Appendix C contains further details on how we link the pollution data with the population density data. Section 4.5 discusses the joint distributions of these variables in detail for the Nigerian case, as an illustrative case.

With these different pieces of information in hand, we can test whether the data are consistent with a simple static spatial equilibrium model when considering this key set of real outcome measures.

4. Empirical Findings

In this section, we present our empirical findings. We begin with measures of consumption of durable goods and housing quality across geographic space, to motivate our methodology and as a frame of reference. As in previous studies, our data confirm that durable ownership and housing quality increase substantially with population density. We then turn to our main outcome measures, in particular the amenity metrics covering health, public goods, crime and pollution. We find that health outcomes, which may reflect public health investments or pollution exposure, are generally better in urban areas. Public goods access is either similar across density or better in denser areas. Crime is similar in most countries and, at best, marginally worse in denser areas in some countries. Some pollution metrics are unrelated to density (fine particulate matter levels), while others are worse in rural areas (use of solid fuels burned indoors).

4.1. Consumption of Durables and Housing

To illustrate our methodology, we consider how telephone ownership varies through space. Telephones are a prime example of a consumer durable good in Africa that is important in the aggregate and clearly associated with higher incomes. Figure 1 shows a kernel-weighted local polynomial regression of whether a household has a telephone at different levels of the log of population density for Ethiopia, Nigeria, Senegal and Tanzania. Each regression line shows a 95-percent confidence interval, and the graphs exclude the top and bottom five percentiles of the distribution. Several facts are worth noting. First, there is a large dispersion in telephone ownership from the least populated to the most populated areas, with a support from zero to one. Second, across the whole range of densities, ownership is increasing almost monotonically and continuously. The pattern is not driven by our choice of countries: we find similar pictures for the remaining 16 countries in our data.

To present data in a compact fashion for our whole set of countries, we divide locations within each country into quartiles of population density, and we compute average durable ownership

Figure 1: Telephone Ownership Gradients
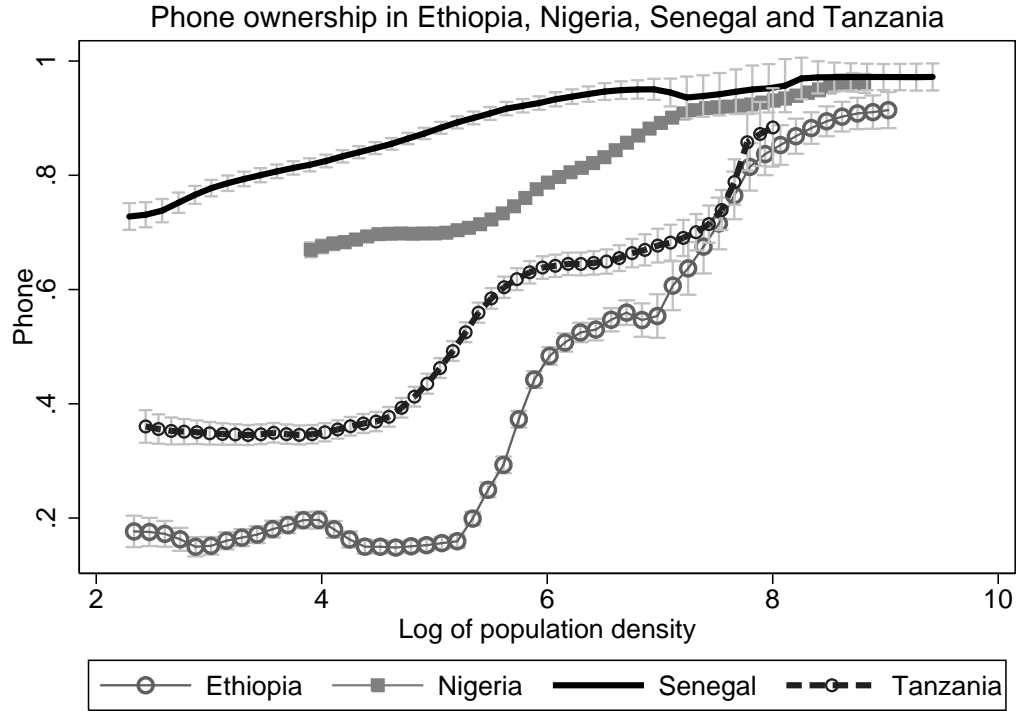
rates by quartile. An advantage of looking at density quartiles is that we avoid the problems of defining "urban" and "rural" locations, which may reflect arbitrary administrative categories and are not standardized across countries.[7]

Table 1 presents the average durable ownership rates by density quartile. The number of countries with statistically significant differences from Q1 are given below each average (both negative and positive differences, so that 0-20 means zero negative differences and 20 positive differences across our 20 countries). Telephone ownership rates are strongly increasing with density, with the least dense quartile having a 41 percent ownership rate, compared to 83 percent in the densest quartile. These differences are statistically significant in all of our 20 countries. Television ownership is similarly higher in denser areas, with just 11 percent having televisions in the least dense quartile and 59 percent in the densest, and, again, all 20 having significant differences. Automobile ownership rates are low everywhere, but still higher in the densest quartile on average and significant in all but one country. Motorcycle ownership rates are similar across quartiles, with five countries having significantly lower ownership rates in denser areas, six having higher rates, and the rest insignificant. Households are likely

---

[7]We use sample weights when computing quartiles and averages across quartiles; when looking at averages across quartiles within countries, we define these within countries; when we aggregate across countries, we define quartiles over the whole sample of countries.

Table 1: Durables Ownership by Density Quartiles

	Population Density Quartile			
	Q1	Q2	Q3	Q4
Telephone	.41	.49	.60	.83
		0-9	0-15	0-20
Television	.11	.15	.29	.59
		1-4	0-16	0-20
Automobile	.01	.02	.03	.10
		0-4	0-8	0-19
Motorcycle	.13	.13	.13	.14
		2-4	4-4	5-6

Note: This table reports the average fraction of respondents who own durables across quartiles across our set of 20 countries. The two numbers below the averages x-y represent the number of countries with a difference between the least dense quartile (Q1) that is statistically significant at the one-percent level and either negative (x) or positive (y).

substituting motorcycles for automobiles to some extent in denser areas, though, overall, they have higher ownership rates of motorized transportation goods. Overall, these results confirm that our metrics are consistent with the existing finding in the literature that real consumption levels are relatively highest in urban areas.

We turn next to measures of housing consumption. We focus on eight measures of housing quality: the percentages of households having: (1) electricity, (2) tap water, (3) a constructed floor, (4) a flush toilet, (5) a finished roof, and (6) finished walls; plus (7) the average number of minutes each day spent collecting water, and (8) the number of sleeping rooms per person over the age of 5. The latter could be seen as a measure of housing quantity. Together, these measures provide a fairly comprehensive view of housing quality for households in the developing world.

Table 2 reports the average housing quality metrics by density quartiles. The table shows that, for each metric, housing quality rises with density on average (see Appendix Figure A.2 for the corresponding plots). Differences between the second and first quartiles are statistically significant for the majority of countries. Differences between the first and fourth (densest) quartiles are quite large in magnitude and statistically significant for all but a few countries. In the least dense quartile, people spend, on average, 29.4 minutes per day collecting water, while those in the densest quartiles spend, on average, 12.2 minutes. In 17 countries, the difference is

Table 2: Housing Quality by Density Quartile

	Population Density Quartile			
	Q1	Q2	Q3	Q4
Electricity	.12	.16	.32	.64
		0-4	0-14	0-20
Tap water	.15	.19	.39	.67
		0-2	0-11	0-20
Constructed floor	.33	.38	.52	.82
		1-5	0-12	0-19
Flush toilet	.02	.05	.12	.31
		0-3	0-10	0-19
Water collection (min)	29.4	24.3	20.7	12.2
		4-0	6-0	17-0
Finished roof	.41	.5	.67	.88
		2-8	1-14	0-19
Finished walls	.35	.42	.53	.82
		0-6	1-14	0-20
Sleeping rooms per person (aged> 5)	.61	.62	.61	.61
		0-2	1-5	4-6

Note: This table reports housing quality across quartiles across our set of 20 countries. The two numbers below the averages x-y represent the number of countries with a difference between the least dense quartile (Q1) that is statistically significant at the one-percent level and either negative (x) or positive (y).

statistically significant from zero. Households in the densest quartiles are substantially more likely to have finished walls, a finished roof or a constructed floor. We conclude that by these six measures, housing quality is unambiguously higher in denser areas in these developing countries.[8] Our one measure of housing quantity – the number of sleeping rooms per person – does not suggest any strong differences across density quartiles: households in densely populated locations do not have substantially less housing space, by this measure, than those living in rural areas.

---

[8]Some measures of housing quality may already be incorporated into measured rural-urban income differences. For example, surveys reporting average consumption expenditures at the household level may deflate housing expenditure by regional quality-adjusted price indices. In practice, it is an open question how well quality adjustments are done in these surveys.
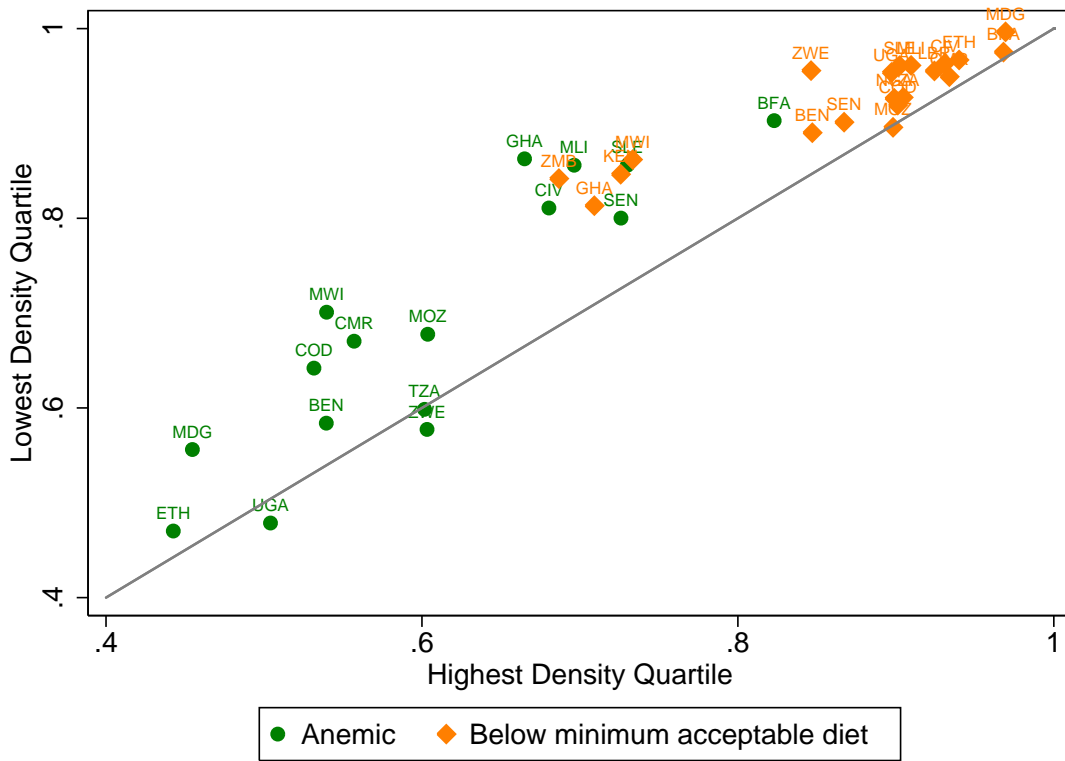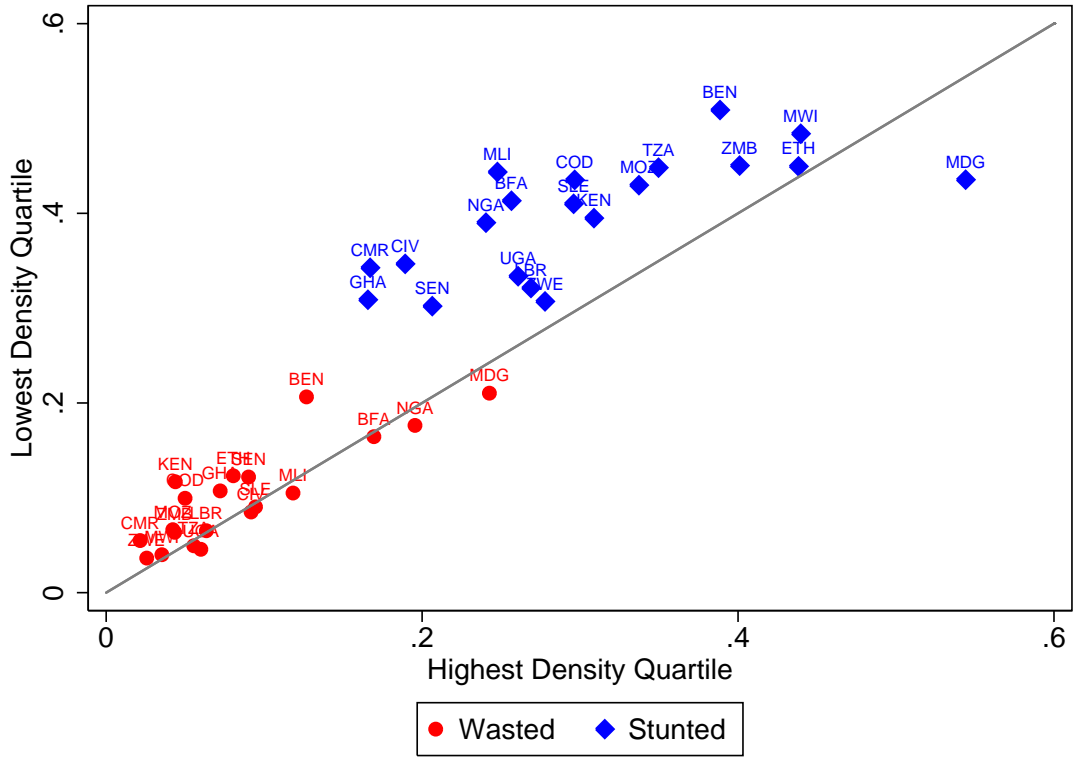
4.2. Child Health

We next look at measures of child health. Child health is informative about both the consumption of the household and amenities such as access to health facilities or medicine. One potential pitfall of using child health measures is that household members living in urban areas might be more informed about children's health problems, thus affecting their propensity to report a health problem. Mindful of this, we selected only outcomes that are objectively measured; in other words, they are not dependent on reporting by caretakers and, thus, do not reflect inaccurate information by respondents.

We look at four main objective measures of child health: stunting, wasting, anemia and consuming an acceptable minimum diet. Stunting is defined as having low height for age, and is perhaps the most commonly used indicator of poor child health calculated using the DHS (see, e.g., Cummins, 2013). Wasting is defined as having low weight for height. The minimum acceptable diet is is defined by minimum meal frequency and dietary diversity (World Health Organization, 2015). For all metrics we calculate the fraction of children that have the condition in question and then aggregate by population density; see Appendix A for more detail on how these variables are defined.

Figure 2 shows the fraction of children with stunting, wasting, anemia and consumption below the acceptable minimum diet. For each metric, there is one point per country, capturing the average value in the lowest-density (y-axis) and the highest-density (x-axis) quartiles. The upper panel covers the percent stunted (darker circles) and the percent wasted (lighter diamonds). The lower panel covers the percent anemic (darker circles) and the percent below an acceptable minimum diet (lighter diamonds).

Two main features stand out in Figure 2. First, for all four metrics, most countries lie near the 45 degree line. This implies that rates of child malnourishment are largely similar in the least dense and densest areas in most countries. Second, in all but a handful of countries, rural areas have higher rates of child malnourishment. Rates of stunting and wasting are similar or higher in rural areas in all countries but Madagascar. Anemia and dietary inadequacy are worse in rural areas for all countries except Zimbabwe and Tanzania, which have marginally worse values in the most densely populated areas. We find that differences in stunting rates are significantly lower in urban areas in 11 out of 20 countries, while around half of the countries have significant differences in the other metrics (see Appendix Table A.2). It is hard for us to conclude whether these findings are largely about better food consumption by urban children or better access to health care in urban areas. In either case, it is safe to conclude that, in most

Figure 2: Child Health

countries, child health is substantially better in more densely populated areas.[9]

There is an important comparison to be made with the historical evidence from currently rich countries such as the United States and those in Western Europe. Interestingly, these rich countries generally have better average health outcomes in cities than in rural areas, though this wasn't previously the case (Costa and Kahn, 2006). In the early 20th century, U.S. cities had higher death rates than rural areas. Since then, these patterns have reversed, with higher life expectancies among urban residents, and public health investments playing a central role in raising the quality of city life (Cutler and Miller, 2005; Costa and Kahn, 2006; Kesztenbaum and Rosenthal, 2016). Thus, while the developing countries in the present study are poorer than the United States was in 1900, they do not share the same pattern of cities being associated with worse health environments that once characterized the United States.[10]

4.3. Public Goods

This section documents the prevalence of public goods across space, such as the presence of a sewerage system, cell phone service, schools and health facilities. So far, we have not distinguished between "consumption" that is privately provided and publicly provided. Many of the goods that we consider as "consumption" or "amenities" incorporate some element of private provision and some element of public provision. For instance, our child health indicators are influenced both by private expenditure (e.g., on food) and public investments (e.g., in health service provision). The same is true, to some degree, even of durable goods ownership, where the private expenditure requires accompanying public investments or public institutional arrangements — e.g., for mobile phones or motorcycles, which require mobile phone coverage and road construction, respectively.

Some public goods (e.g., legal systems) are difficult to assign to geographic locations. However, certain other categories of public goods can be assigned to locations in density space, and we now briefly consider these. If these goods were more abundantly provided in sparsely populated rural areas than in densely populated areas, they might offset the urban-rural differences in consumption goods.

Table 3 shows a number of public goods across density quartiles. Overall, none of these mea-

---

[9]We also looked for geo-referenced data on infectious diseases as well, but found that they were largely unavailable for our set of countries. One exception are malaria test samples taken as part of the Demographic and Health Surveys, Malaria Indicator Surveys and AIDS Indicator Surveys. Using all data available for our countries from these surveys, we find that malaria incidence is decreasing with population density consistently in most of the countries. This is in line with Tatem, Gething, Smith, and Hay (2013), who document a negative relationship between malaria and urbanization on a global scale.

[10]Jedwab and Vollrath (2017) argue that much of the recent growth of cities in the developing world is due to their relatively low mortality rates, unlike cities in Western Europe in the past.

sures is worse in denser areas, on average. Denser areas are more likely to have electricity grids, piped water systems, sewerage systems and cell phone coverage. A number of services are also more likely to be present in denser areas: post offices, health facilities, police stations and market stalls. Finally, transportation infrastructure is better – at least when measured by whether a road is paved. These findings are consistent with the theory that, although there may be some congestion issues in urban areas that offset the benefits to individuals, the per capita cost of providing many public goods in remote and sparsely populated locations has led governments (rationally) to supply lower quantities of public goods in these areas.

It is also worth pointing out that the patterns across density are very similar to those emerging from the household-level DHS data, both in terms of the monotonically increasing relationship as well as the differences across, for example, the bottom and top quartiles. An important limitation of the data is that we can observe only the presence of a public good, but not its attributes, such as the availability of doctors and nurses or teachers, their qualifications, or how crowded a health facility or school is; we also do not observe whether a sewer is open or closed, or how reliable the electricity off the grid is. We still think that these measures contain valuable information on the distribution of amenities across space.

4.4. Crime

We next turn to measures of crime, which is arguably one of the most important non-monetary amenities that varies across space. The challenge we face is that data on crime at small levels of geographic detail are difficult to obtain for African countries. Official administrative records are either not stored centrally or not available to researchers.

The best and most comparable data on crime come from the Afrobarometer surveys, which were collected in 2005, 2009 and 2011 for a large subset of our countries (see Table A.6 in the Appendix for more on the surveys we employ). The surveys are designed to use consistent methodologies and definitions across countries. The questionnaire focuses on attitudes towards democracy and governance and includes questions on crime, safety and trust. An advantage of survey data over administrative data on crime is that the latter are likely to be biased towards areas with police presence or better record-keeping capacity. In contrast, our surveys are administrered in the same way within and across countries, as well as across the three years of our survey. Hence, our data are unlikely to be biased toward any particular geographic area.

We consider four main metrics: property crime, violent crime, feeling safe in one's neighborhood, and fear of crime in one's home. To measure property and violent crime, we use the survey questions: *"Over the past year, how often (if ever) have you or anyone in your family had*

Table 3: Public Goods by Density Quartile

	Population Density Quartile			
	Q1	Q2	Q3	Q4
Electricity grid that most houses could access	0.39	0.42	0.48	0.72
		1-0	0-4	0-10
Piped water system that most houses could access	0.36	0.35	0.42	0.67
		0-0	0-2	0-11
Sewage system that most houses could access	0.14	0.13	0.18	0.37
		0-0	0-0	0-7
School in EA or within easy walking distance	0.91	0.90	0.90	0.90
		0-0	0-0	0-0
Health Clinic in EA or within easy walking distance	0.59	0.58	0.62	0.73
		1-0	0-0	2-4
Police Station in EA or within easy walking distance	0.29	0.30	0.33	0.47
		0-0	0-0	1-4
Market Stalls in EA or within easy walking distance	0.62	0.64	0.64	0.76
		0-1	0-0	0-4
Road on way to interview was at the start paved	0.27	0.30	0.35	0.54
		1-0	1-2	0-6

Note: This table reports the average fraction of enumeration areas having access to a public good. Electricity grid, piped water system and sewerage system are equal to one if they can be accessed by most houses in an enumeration area. Post-office, school, health clinic, police station and market stalls are equal to one if they are present in the enumeration area or within easy walking distance. Road to interview is equal to one if the road taken on the way to the interview was at the start paved/tarred/concrete. Numbers below the averages in each row are the number of countries with a difference from the least dense quartile (Q1) that is statistically significant at the one-percent level. In this table we use within-country quartiles rather than across-country quartiles to ensure that all countries have enumeration areas across quartiles.

something stolen from your house?" and *"Over the past year, how often (if ever) have you or anyone in your family been physically attacked?"* For each region, we compute the fraction of respondents reporting at least one theft (property crime) or attack (violent crime).

To measure feeling safe and fear of crime, the questions are: *"Over the past year, how often, if ever, have you or anyone in your family felt unsafe walking in your neighborhood?"* and *"Over the past year, how often have you or anyone in your family feared crime in your own home?"* The possible answers to these questions on experienced crime and perceived safety are: *"never,"*

Figure 3: Crime
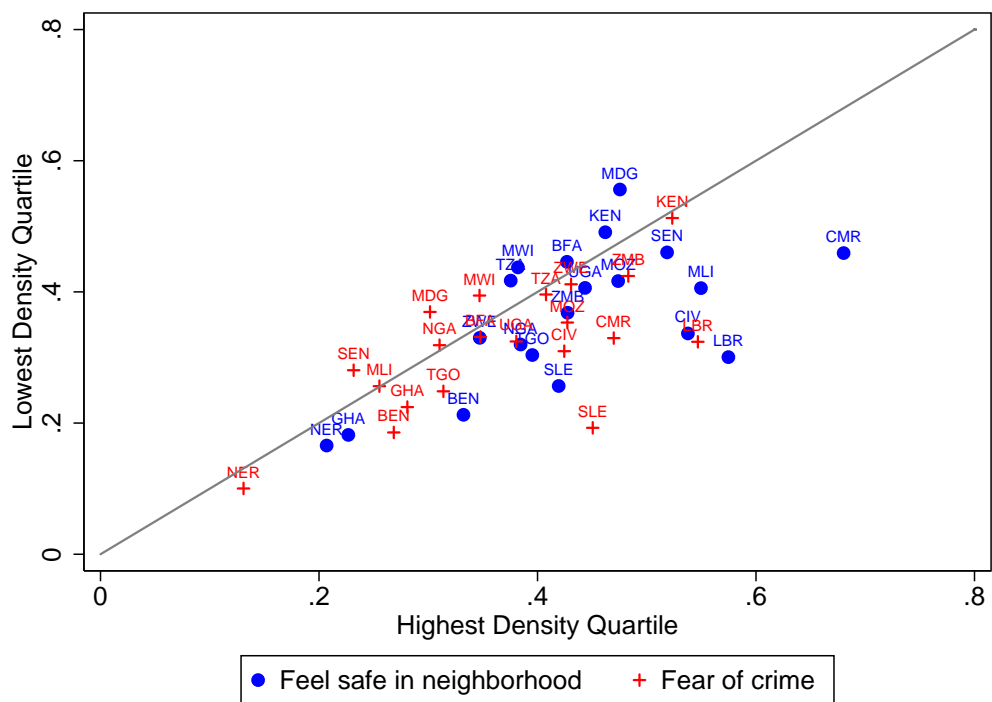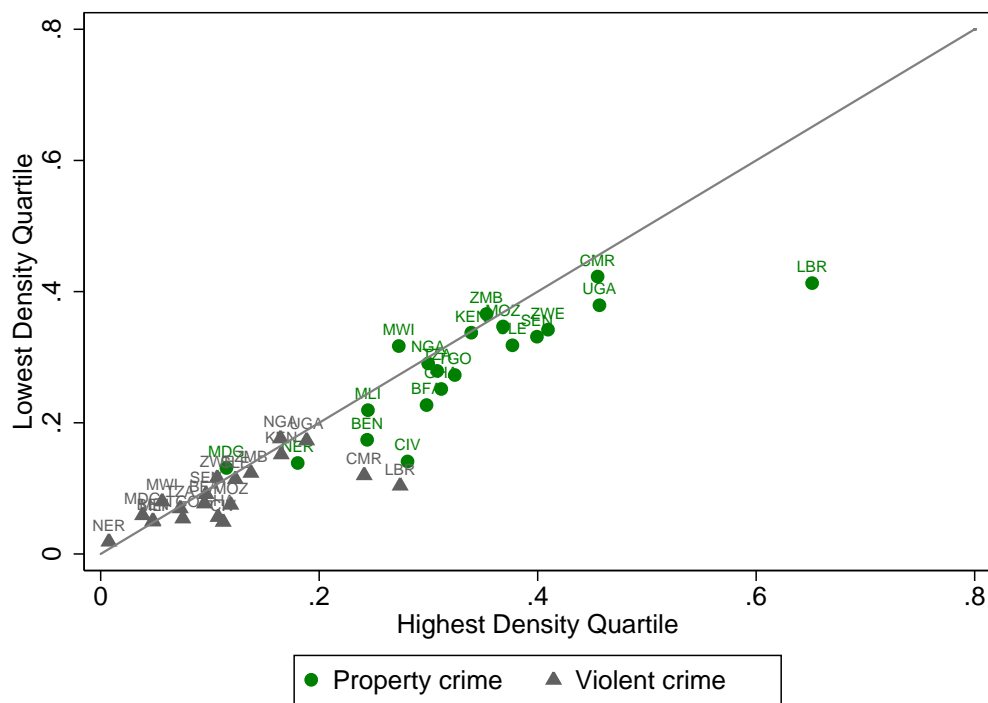
"just once or twice," "several times," "many times," and *"always."* We define a dummy variable

as equal to one if a respondent's reply is anything more than "never."

Overall, we find that crime is quite common in Africa. About one third of respondents report a theft from their house in the previous year. The highest rates of theft are in Liberia (49 percent), Uganda (42 percent) and Senegal (39 percent), and the lowest rates are in Madagascar (13 percent), Niger (18 percent) and Mali (21 percent). The heterogeneity in physical attacks follows a similar pattern for most countries, and the pairwise correlation coefficient at the country level between theft and attack is 0.7 and highly significant. Exceptions include Senegal, where theft is high but attacks are reported infrequently. Across the whole sample, more than one third of respondents report that they felt unsafe in their neighborhood at least once in the past year, and that they feared crime in their own home.

Figure 3 shows differences in experienced crime and fear of crime across space. We show both of these categories of variables, as fear of crime might matter at least as much as experiences of crime for location choices. Both figures illustrate that most countries are located close to the 45-degree line. Property crime appears to be slightly higher in denser areas, but the differences for most countries are fairly small when comparing them with observed differences in living standards, for example. One limitation of the theft variable is that it does not consider livestock theft, a type of crime common in rural areas. It is, therefore, likely that the difference is even smaller when taking into account livestock theft. The results are similar for fear of crime and perceived feeling of safety in the neighborhood, where most countries cluster around the 45-degree line.

Table 4 presents the average crime rates by density quartile across all the countries. Around 29 percent of households in the least dense quartile experience property crimes, compared to 33 percent in the densest quartile. Violent crime affects ten percent of households in the least dense quartile compared to 12 percent in the densest quartile. Fear of crime and feeling unsafe are similarly increasing in population density on average, with similarly modest differences by density. Only a handful of countries have statistically significant differences in crime rates through density, with the majority insignificant.[11]

Still, our evidence so far suggests that crime is the leading contender for the amenity that

---

[11]Studies by Fafchamps and Moser (2003) and Demombynes and Ozler (2002) from Madagascar and South Africa point to somewhat higher crime rates in less dense areas. Crime information is also available in some LSMS surveys, though with questions that are hard to compare across countries. Appendix Figure A.9 shows the fractions of LSMS households having experienced a crime in the last twelve months, in five countries for which geo-references are available: Ethiopia (2013), Malawi (2010), Tanzania (2009-2010), Nigeria (2012) and Uganda (2010-2011). While patterns of crime across space are different in each of these five countries, one common theme is a lack of evidence that rates of crime are systematically increasing in density, or a lack of evidence that the densest areas have the highest rates of theft; visually, a constant rate of crime seems as though it would just about fit within the confidence intervals.

Table 4: Crime and Fear of Crime by Density Quartile

	Population Density Quartile			
	Q1	Q2	Q3	Q4
Property crime	.29	.31	.31	.33
		0-3	0-1	0-4
Violent crime	.1	.09	.1	.12
		0-1	0-0	0-4
Fear of crime	.32	.33	.34	.36
		0-1	0-2	0-4
Feel unsafe	.37	.39	.38	.45
		0-2	0-1	0-3

Note: This table reports the average fraction of respondents reporting property crime, violent crime, fear of crime in one's home, and feeling unsafe in one's neighborhood. The two numbers below the averages x-y represent the number of countries with a difference between the least dense quartile (Q1) that is statistically significant at the one-percent level and either negative (x) or positive (y).

gets worse with density. Could it be that crime rates are enough to offset the higher income and consumption levels of more-urban areas? Several previous studies in the literature have estimated the value of living in an areas with less crime, proxied by willingness to pay. What is the value of having 33-percent chance of theft in the densest areas relative to just 29 percent in the least dense, as we find on average? Or having a 12 percent chance of violent crime in the densest areas compared to just ten percent in the least dense areas?

Relative to the large differences in average income across space, the estimated valuations of crime implied by estimates in the literature are quite modest. For example, using crime and property-value data, Bishop and Murphy (2011) estimate a dynamic model and infer that San Francisco residents are willing to pay $472 per year to avoid a ten-percent increase in violent crime. On an average income per head of $57,276, this amounts to 0.8 percent of average yearly income. Using direct survey questions, Cohen, Rust, Steen, and Tidd (2001) estimate that, in 2000, U.S. residents were willing to pay $120 to reduce the chance of armed robbery by ten percent. This amounts to 0.4 percent percent of average income ($120 / $34,432). Similarly, Ludwig and Cook (2001) estimate that U.S. households in 1998 were willing to pay $240 per year to reduce the chance of gunshot injury by 30 percent, which amounts to 0.5

percent of average household income ($240 / $51,939). Taken together, these studies suggest that the modest differences in crime rates with density may not be large enough to offset the much higher average incomes in urban areas.

4.5. Outdoor Air Pollution

Pollution is a widely studied amenity that varies through space. Sources of outdoor pollution include vehicles, electricity generation, industry, waste and biomass burning, and re-suspended road dust from unpaved roads. Banzhaf and Walsh (2008) find pollution to be an important determinant of locational choice in the United States, and exposure to pollutants significantly affects health, human capital and productivity (Adhvaryu, Kala, and Nyshadham, 2014; Currie and Walker, 2011; Currie, Hanushek, Kahn, Neidell, and Rivkin, 2009; Graff Zivin and Neidell, 2012).[12]

In this section, we use satellite-derived estimates of two measures of outdoor air pollution — fine particulate matter (PM2.5) and nitrogen dioxide (NO2) — and document how they vary with population density in our set of countries and several reference countries. We measure both pollution measures in micrograms per cubic meter ($\mu g/m^3$). For frame of reference, the World Health Organization recommends mean annual exposures of 10 $\mu g/m^3$ or less for PM2.5 and 40 $\mu g/m^3$ or less for NO2, at the same time highlighting that there are no levels of pollution exposure that have been proven not to negatively affect health (Geddes, Martin, Boys, and van Donkelaar, 2016; WHO, 2006). Indeed, Pope and Dockery (2006) conclude from a meta analysis of dozens of studies that the relationship between particulate-matter exposure and life expectancy is approximately linear.

Figure 4 illustrates this procedure and shows the distributions of pollutants and population density across space in Nigeria. The top left graph shows the distribution of population density; the top right graph shows the NO2 distribution; and the two bottom graphs show PM2.5, where the graph on the right removing sea salt and dust. Warmer colors denote higher values, and the bins are formed by dividing the data into deciles. Population density in the North is highest around Kano; in the center around Abuja; in the South West close to Lagos and Ibadan; and in the South East between Benin City, Port Hartcourt and Enugu.

Moving to the pollution measures, several observations are worth highlighting: first, at least visually, population density does not appear to be strongly correlated with either of the pollution measures. Nitrogen dioxide levels are very low, with a maximum of 0.7 ppb (1.316 $\mu g/m^3$), far below the WHO recommended thresholds of 40 $\mu g/m^3$. Values are higher over

---

[12]For a comprehensive survey of the literature on pollution and individual welfare see Graff Zivin and Neidell (2013); on environmental amenities and city growth see Kahn and Walsh (2015).

Figure 4: Outdoor Air Pollution (PM2.5) in Nigeria
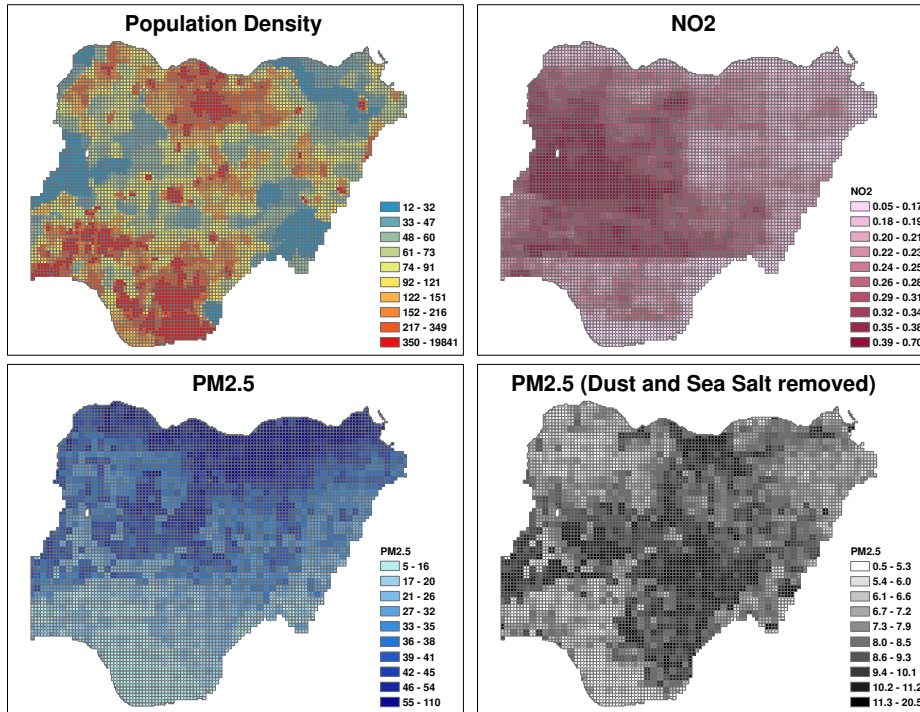
Notes: The top left graph shows the distribution of population density, the top right graph shows the NO2 distribution, and the two bottom graphs show PM2.5, where the graph on the right removes sea salt and dust. Warmer (darker) colors denote higher values, and the bins are formed by dividing the data into deciles.

Lagos, Ibadan, Abuja, Kanduna and Kano, but not over cities in the South East in the Delta, and there are high levels in the center towards the West of the country, where few people live. PM2.5 levels appear to be driven mainly by dust from the Sahara when inspecting the bottom left graph. Removing sea salt and dust produces quite a different distribution, with higher levels in the center and over some cities.

The Nigerian example illustrates further that looking separately at these two indicators for pollution is instructive.[13] The pairwise correlation between PM2.5 and NO2 is -0.0085 with a p-value of 0.4633. Across our whole set of African countries, the correlation of these two measures ranges from 0.65 (Cameroon) to -0.47 (Senegal).

Figure 5 plots (as black circles) the average PM2.5 concentrations for the highest and lowest population density quartiles across all our countries. Those with the highest overall concentrations of PM2.5 tend to have high concentrations both in the least and most dense areas. These countries are essentially the ones that border or contain parts of the Sahara desert, such as

---

[13]This is in line with what Geddes, Martin, Boys, and van Donkelaar (2016) find when they inspect population weighted average PM2.5 and NO2 levels and trends.

Figure 5: Outdoor Air Pollution
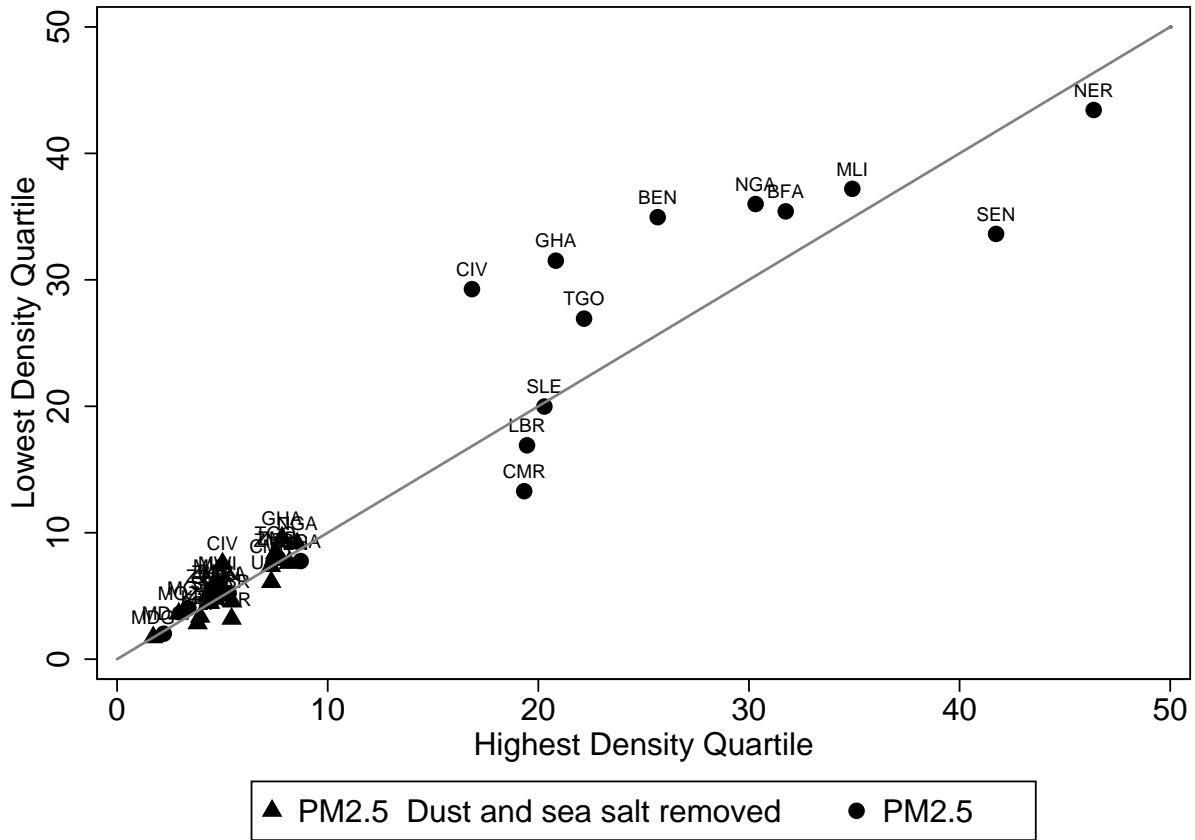
Niger, Senegal, Mali and Nigeria. Therefore, we look plot PM2.5 concentrations with dust and sea salt removed (as black triangles).[14] There is no evidence suggesting that anthropogenic sources of PM2.5 are more or less harmful for health than are natural sources, but it is possible that individuals perceive anthropogenic sources as more hazardous to their health. These lead to far lower levels of PM2.5, again with very modest apparent differences by population density. Figure A.3 plots NO2 by population density. While there is some variation across countries, there is, again, little apparent difference by density. When we look at average pollution levels by density quartile, we find that NO2 concentrations have largely statistically insignificant differences (see Appendix Table A.3). PM2.5 on average is lower in denser areas, though the differences are statistically significant only in around half the countries, and the average differences are again modest. We conclude that outdoor air pollution concentrations in Africa are at best loosely linked to population density.

---

[14]About half of the population-weighted ten-year mean of PM2.5 concentration in Eastern Sub-Saharan Africa is estimated to be due to dust and sea salt; in Western Sub-Saharan Africa the proportion amounts to about three quarters (van Donkelaar, Martin, Brauer, and Boys, 2015).

What do these gradients look like in other parts of the world, known to suffer from high levels of urban pollution, such as China? The global coverage of the PM2.5 map and the population density map allows us to look at additional countries. Figure A.4 shows the relationship between pollution and population density in China, India, and the United States. The left-hand side axis denotes levels for India and China, while the right-hand axis indicates the level in the United States. All three countries show clearly positive density gradients. In China, PM2.5 levels for the top population density decile amount to 66 $\mu g/m^3$, more than six times the WHO recommended threshold; the lowest population density decile has a level of 13 $\mu g/m^3$. In India, the top decile has a level of 41 $\mu g/m^3$, still four times the WHO recommended threshold, compared to 6 $\mu g/m^3$ in the lowest decile. The bottom figure shows NO2 distribution for these other countries. The levels are much lower, but there are gradients again for China, India and the United States. These positive gradients are very much what we might expect of a world in which cities have high concentrations of industrial activity and automobile traffic. Although urban areas in Africa are growing rapidly, there is little industrial activity (Gollin, Jedwab, and Vollrath, 2016), and, consequently, industrial air pollution is relatively low.

To be clear, we are not making a claim that outdoor pollution does not matter in African cities. Our satellite-derived pollution estimates do not capture pollution exposure in some dimensions: they are annual series and, therefore, average over temporally high values. At a 10km resolution, they are spatially rather coarse, ignoring local effects such as proximity to roads that have been demonstrated to matter significantly. For example, Kinney, Gichuru, Volavka-Close, Ngo, Ndiba, Law, Gachanja, Gaita, Chillrud, and Sclar (2011) find average PM2.5 concentrations at four traffic sites between 7.30am and 6.30pm in Nairobi to amount to between 58.1 and 98.1 $\mu g/m^3$; the maximum multi-annual average PM2.5 concentration for Kenya in our sample is 13.9 $\mu g/m^3$, and this pixel is at Lake Turkana, the world's largest desert lake (Avery, 2012). Finally, satellite-derived measures reflect the column of pollution as observed from space, rather than the concentration experienced on the ground. Nevertheless, as the graphs from India, China, and the United States illustrate, the two datasets still capture meaningful variation in concentration levels across space. What emerges, rather, is that in Africa, cities are not large enough and concentration of industries is not significant enough to create large clouds of pollution around cities; however, background non-anthropogenic pollution is high. This combination produces different pollution gradients from those observed in more industrialized parts of the world.

Figure 6: Indoor Air Pollution
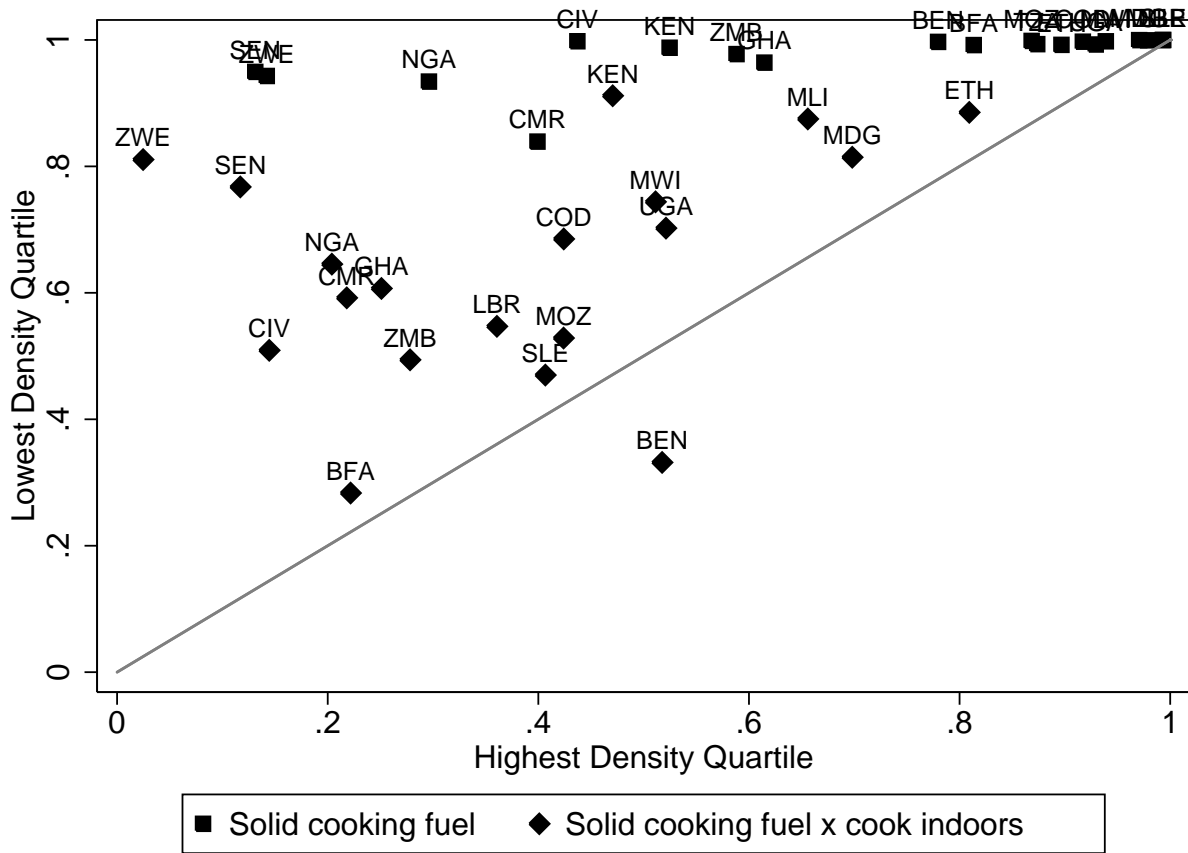
4.6. Indoor Air Pollution

Arguably a more serious health risk in developing countries is indoor air pollution, largely related to the use of unvented fires for cooking. As a proxy for indoor air quality, we examine the main material used for cooking as reported in the DHS. The World Health Organization estimates that over four million people suffer from premature deaths due to illnesses attributable to cooking with solid fuels, such as wood and charcoal (WHO, 2014). Indoor air pollution is also estimated to contribute significantly to outdoor air-pollution-related deaths.

Figure 6 shows the proportion of households using solid fuels for cooking across population density (as black squares). Across all of our countries, virtually everyone in the least dense areas uses a solid cooking fuel. In the urban areas there is more variation, with the densest areas ranging from 20 percent solid fuel in Zimbabwe up to around half in countries such as Nigeria, Cote d'Ivoire and Ghana, and the majority over 80 percent. Overall, the pattern is clearly one of much higher use of solid fuels in rural areas than in urban areas.

One potential advantage of rural areas is that there might be more space to accommodate outdoor cooking, thereby somewhat mitigating the negative effect of using solid fuels. Therefore we also show the interaction effect: among the households using solid fuels, we look at the proportion cooking indoors (as black diamonds). The fraction of people cooking primarily indoors with solids drops everywhere, and it is still higher in rural areas in almost every country (Benin being the sole exception). The last row of Table A.3 shows that the fraction cooking inside with solid fuels is six percentage points lower in the third quartile than in the least dense quartile, and 25 percentage points lower in the densest quartile. All but one country have a significant difference. In summary, indoor air pollution is worse almost everywhere in rural areas than in African cities.

We note in passing that our pollution measures do not include adequate data on trash and refuse, which might be more abundant and more visible in urban areas than rural areas. We also do not capture measures of water quality directly – although a number of our health measures (especially child health outcomes) are quite sensitive to certain types of water pollution and sanitation problems. We cannot eliminate the possibility that there are some disamenities related to pollution and waste in urban areas. But we also note that rural areas have their own issues of water and air pollution, and the data for our African countries are not consistent with the idea that pollution is systematically worse in urban areas or locations of high population density.

5. Heterogeneity by Education Level and Migration Status

In this section we consider how our results vary by education level, and how rural-urban migrant households differ from other urban households.

5.1. Heterogeneity by Education Level

One explanation for the patterns that we observe is that workers select into different regions and occupations according to comparative advantage. This type of sorting has been emphasized recently in the macroeconomics literature by, e.g., Lagakos and Waugh (2013), Young (2014), Bryan and Morten (2015) and others, and by many in the urban economics literature, including Glaeser and Mare (2001) and Combes, Duranton, and Gobillon (2008), as a way of explaining regional average income differences. One variant of the spatial equilibrium hypothesis building on the sorting literature is that there is a spatial equilibrium within each permanent-income (education) category, but not across categories. This theory could potentially reconcile the higher average consumption and amenity levels of urban areas by simply having more of the educated households living in urban areas. To address this issue further,

we compare how measured living standards vary by density within specific educational groups. Educational attainment is perhaps the simplest measure of permanent income that is available at the individual level in the data, though of course there are components of ability not well captured by educational attainment.

For simplicity, we divide households into two education groups: those whose head finished primary school, and those whose head did not finish primary school. Figure A.5 plots one metric by educational group: electricity access. The top graph shows the proportion of households that have electricity by the highest education of the household head. The histograms show that the different educational groups are represented at various population density levels. Households with more-educated household heads experience better access to electricity at almost all levels of population density. Still, the urban-rural gradients documented earlier persist even within these education categories.

To study these slopes by education group, we estimate the following linear projection for households i in country c:

$$x_{ic} = \theta_0 + \theta_1 P_{ic} + \theta_2 E_{ic} + \theta_3 (P_{ic} * E_{ic}) + \epsilon_{ic}$$
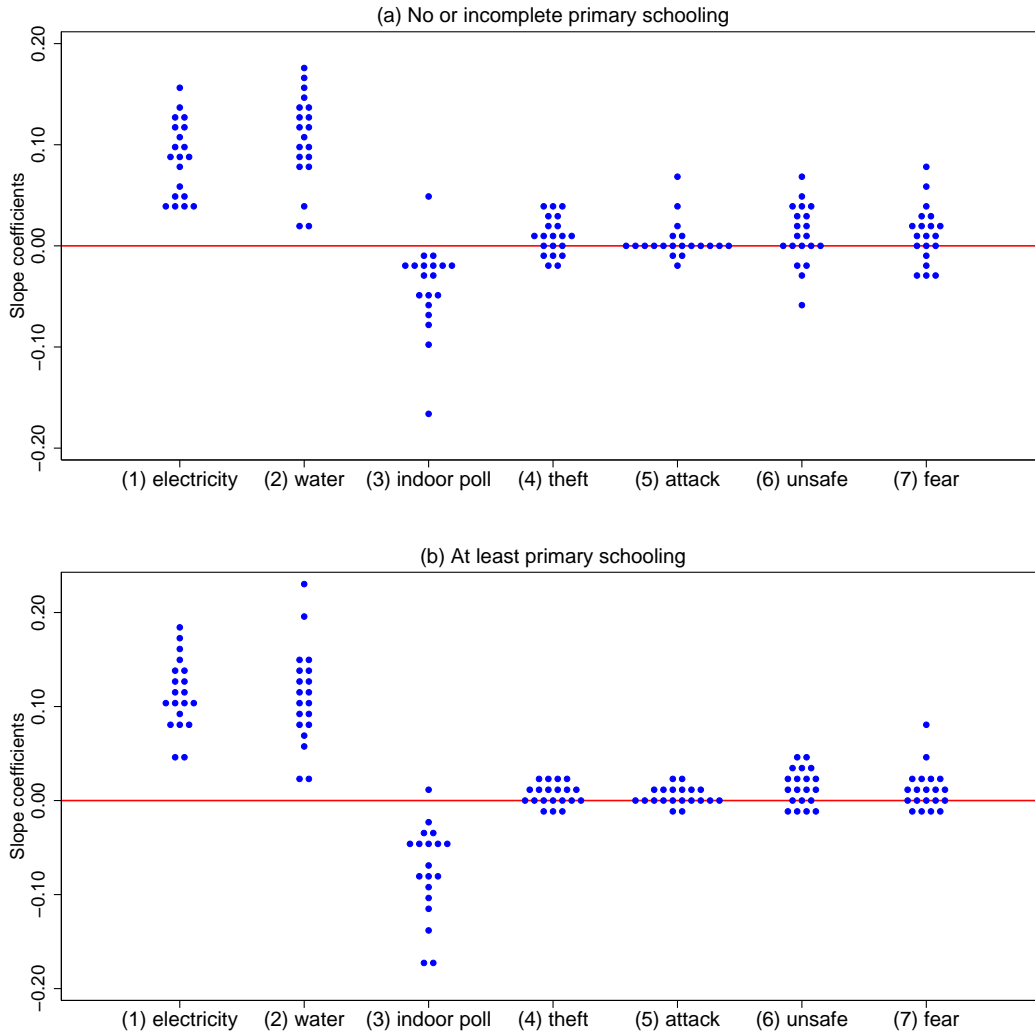
where x_{ic} is a measure of consumption or amenities, P_{ic} is the log of population density and E_{ic} is a dummy variable that is equal to one if the household head has completed primary education or more.

Figure 7, Panel (a), shows the linear gradients for households with heads who have less than complete primary education, and panel (b) shows gradients for households with heads who have complete primary education or more. Each dot represents a slope estimate for one country. The y-axis indicates the size of the coefficient. The figures show that in virtually all countries and for virtually all of our living-standards measures, these urban-rural gradients persist *within* populations at similar educational levels. That is, almost without exception, the relationship between population density and housing consumption is positive for the two main education groups. These findings provide evidence against the variant theory in which a spatial equilibrium holds within educational categories.

5.2. Rural-Urban Migrant Households

Migration decisions are likely to be shaped by the conditions in arrival destinations rather than by averages of a population. It might be the case that conditions in the location to which individuals would consider migrating are much worse than the averages we have documented so far. Our data do not contain any information about whether a cluster is a migration destina-

Figure 7: Density Gradients by Education of Household Head

(a) No or incomplete primary schooling

(b) At least primary schooling

Note: Each dot in the Figure represents the coefficient estimate of a linear projection for household i in country c: $x_{ic} = \theta_0 + \theta_1 P_{ic} + \theta_2 E_{ic} + \theta_3 (P_{ic} * E_{ic}) + \epsilon_{ic}$ where x_{ic} is a measure of consumption, P_{ic} is the log of population density, and E_{ic} is a dummy variable that is equal to one if the household head has completed primary education or more. Panel (a) shows the linear gradients for households with household heads/respondents who have less than complete primary education; and panel (b) shows gradients for households with household heads/respondents who have complete primary education or more. The y-axis indicates the size of the coefficient.

tion, and the random displacement of clusters makes it infeasible to spatially merge the DHS data with secondary data (i.e., "slum" settlements). Still, we can use the individual-level migration data to characterize rural-urban migrant households and compare them to other urban households and to rural households.

Table 5 reports measures of durables ownership, housing quality, public goods and indoor air quality for rural-urban migrant households in the densest quartile. The first two data columns

Table 5: Rural-Urban Migrants

| | Population Density Quartile | | | | |
| | Q1 | Q4 | Q4 | Q4 | Q4 |
	whole sample	whole sample	>=2 adult rural-urban migrants	>=50% rural-urban migrants	>=50% rural-urban migrants & hh head <=30 years
Telephone	0.20	0.68	0.79	0.68	0.63
		0-12	0-12	0-12	0-10
Electricity	0.09	0.57	0.64	0.58	0.56
		0-12	0-11	0-12	0-10
Tap water	0.12	0.63	0.68	0.64	0.65
		0-12	0-10	0-12	0-10
Constructed floor	0.28	0.79	0.87	0.81	0.79
		0-11	0-12	0-11	0-10
Flush toilet	0.02	0.30	0.33	0.29	0.25
		0-12	0-9	0-10	0-7
Number of sleeping rooms	0.59	0.59	0.48	0.59	0.65
		1-3	5-0	3-3	0-4
Cook inside with solid fuel	0.64	0.38	0.34	0.33	0.32
		6-0	6-0	6-0	6-0

Note: This table reports measures of durables ownership, public goods, housing quality and indoor air pollution across the 12 countries for which we have migration data. The fourth and fifth data columns report the averages for regions in the densest quartile of the country with the fewest migrants and most migrants. The two numbers below the averages x-y represent the number of countries with a difference between the least dense quartile (Q1) that is statistically significant at the one-percent level and either negative (x) or positive (y).

reproduce the averages for the least dense quartile (Q1) and densest quartile (Q4) in our whole sample. The next three data columns report averages for several alternative definitions of migrant households. We are somewhat limited by the way the DHS defines the individual-level sample: individual questionnaires are only administered to a subset of household members that

fulfill certain eligibility criteria (mostly being in the right age range); furthermore, men are only surveyed in a fraction of all households. If we were to count migrant adults per household we would falsely assign a lower number of migrants to households in which fewer individuals were asked about their migration history. To overcome this, we limit the sample to households that were selected for both male and female individual-level interviews, which leaves us with about 50 percent of the whole sample.

The first definition of migrant households we consider is households with at least two adults who have migrated from a rural area in the last five years (the third data column). By this definition, migrant households in Q4 have similar characteristics to other households in Q4, with somewhat higher telephone ownership rates and electricity connections, on average, and somewhat fewer sleeping rooms per adult. More importantly, migrant households in Q4 are still far better than the average households in Q1. We next consider households where at least 50 percent of the adults are rural-urban migrants. The averages for this metric again point to comparable characteristics between migrant households and other urban households, and far better characteristics of migrant households than the average rural household.

Finally, we look at rural-urban households where at least 50 percent of household residents are rural-urban migrants, and where the household head is 30 years old or younger. This definition allows us to make sure we are excluding migrants that have gone to live with (say) a wealthy, established older relative. These migrant households appear similar or marginally worse than other urban households by most of our metrics. But they are still far better than the average rural households of Q1. In 10 of 12 countries, the migrants have statistically significantly higher telephone ownership rates, electricity and tap water connections, and constructed floors. Migrant households have marginally more sleeping rooms per adult, and substantially lower rates of cooking with solid fuels indoors.

Overall, at least among these living standards metrics, there is little support for the hypothesis that rural-urban migrants tend to be worse off than households living in rural areas, on average. Instead, rural-urban migrants appear to have far better living conditions than households that remain in rural areas.

6. Net Migration and Secondary Amenities

The descriptive statistics of outcomes across population densities in Section 4 suggest that there is no easily observable measure of consumption or amenity that is decreasing with population density. We turn now to measures of migration. Taken literally, the spatial equilibrium in our model predicts that we should observe no migration across regions. Taken more liberally, the

model suggests that we should see roughly similar movements of workers from less dense to more dense regions, than from more dense to less dense areas. We ask whether this is the case. We then turn to some secondary amenities to ask how those vary across space.

6.1. Net Migration Rates

We now compute, for the subset of countries with appropriate data, the fraction of all surveyed individuals in the DHS that are rural-to-urban migrants and the fraction of all individuals that are urban-to-rural migrants. Ideally, we would know the exact location from which an individual migrated. Unfortunately the DHS data do not contain this information. However, we know if an individual moved from the capital, from a large city or town, or from the countryside. We define an urban-rural migrant as someone who has been residing in the lowest-density quartile for five years or less, and who previously lived in the capital or a large city. Similarly, we define a rural-urban migrant as someone who has been residing in the highest-density quartile for five years or less, and who previously lived in the countryside.[15]

Table 6 displays the fractions of all individuals that are rural-to-urban migrants, urban-to-rural migrants, and their difference. In every country, there are substantially more rural-to-urban migrants than the opposite. The differences are starkest in Kenya, where 7.6 percent are rural-urban migrants, compared to 0.6 percent urban-to-rural migrants, and Malawi, which has 7.2 percent rural-urban migrants and less than 0.5 percent moving in the opposite direction. All other countries but one have positive net migration.[16]

The table shows clearly that for all but one country, rural-urban migration is larger in absolute terms than urban-rural movements, and in all cases, the net flows are significantly positive. This view is consistent with cities being seen as attractive places to live, and workers voting with their feet to move there. It is hard to reconcile this evidence with a spatial equilibrium. To most development economists, of course, this may not be a new or controversial claim; the literature has long emphasized the importance of rural-urban migration as one feature of structural transformation. But policy makers continue to worry about excessive urbanization, and many academic economists use models that explicitly or implicitly assume that population movements are associated with some kind of sorting that is consistent with a steady-state dis-

---

[15]We do not consider individuals who reported that they previously lived in a town due to the difficulty of assigning the person to a high- or low-density area. We keep them in the sample to compute the appropriate population-weighted cut-off for density quartiles.

[16]Our finding is not in contrast to Young (2014), who reports that urban-to-rural migration as a fraction of the rural population is as large as rural-to-urban migration as a fraction of the urban population. He states that "the difference relative to shares of destination arises because of the smaller average urban population share (0.41 versus 0.59 for rural). Overall, net migration is in favor of urban areas with, on average, 0.126 of the aggregate young adult female population moving to urban areas and only 0.07 to rural areas."

Table 6: Rural-Urban and Urban-Rural Migrants as Percent of Adults

	Rural-to-Urban	Urban-to-Rural	Difference
	Percent of Adults		
Dem. Republic of the Congo (2007)	2.39	0.47	1.92***
Ethiopia (2005)	3.08	0.15	2.93***
Ghana (2008)	4.82	1.18	3.64***
Kenya (2008-2009)	7.60	0.58	7.02***
Liberia (2007)	2.46	2.24	0.23
Madagascar (2008-2009)	4.16	0.19	3.97***
Malawi (2010)	7.23	0.45	6.77***
Mali (2006)	4.46	0.66	3.80***
Nigeria (2008)	4.83	0.37	4.46***
Senegal (2005)	2.75	0.92	1.83***
SierraLeone (2008)	4.44	0.36	4.08***
Zambia (2007)	4.00	0.56	3.44***

Note: The first column lists the country and year of survey. The first two data columns report the percent of adults that are rural-to-urban migrants and urban-to-rural migrants, respectively, in the last five years. The third data column reports the simple difference. ***,**,* mean statistically significant at the one-, five- and ten-percent levels. Test statistics are computed taking into account the stratified sampling design.

tribution of population. For these reasons, we find it useful to emphasize that the net flow of people in these economies is clearly single-directional.

6.2. Secondary Amenities Measures

We turn next to a set of issues that we view as secondary amenities, by which we mean that they are somewhat less tangible than, for example, pollution and crime, considered above. In particular, we focus on measures of social networks, social insurance and stress. By bringing together large numbers of people, cities can create stress and lead to feelings of isolation and anonymity. Urban dwellers may find that they have weaker social links to supportive networks and communities than rural people have. This may be important if would-be migrants worry about their access to the insurance provided by their social networks. If individuals have sufficiently strong preferences for the social integration – and perhaps also for smoother

consumption or lower deviation from the consumption of those around them – then these characteristics of urban life could explain some of the stark differences in living standards that we observe.

The Afrobarometer asks a number of questions that allow us to further investigate these possible explanations related to shortages of necessities, stress and trust. To examine consumption variability, we look at whether and how often a household has gone without enough food or medicines in the past year. Representing crucial consumption goods, these variables are possibly good proxies for insurance provided by community members. We define a household as lacking food or medicines if they state that they lacked the good several times, many times or always. Finally, the Afrobarometer asks: *"In the past month, how much of the time: Have you been so worried or anxious that you have felt tired, worn out, or exhausted?"* If respondents reply with *"many times"* or *"always"* we take this as a measure of anxiety.

There are also a number of questions related to trust towards people in general, as well as towards various distinct groups of people, including neighbors and relatives. The questionnaire asks: *"Generally speaking, would you say that most people can be trusted or that you must be very careful in dealing with people?"* We create a dummy variable that is equal to one if the individual responds with *"most people can be trusted"* and zero otherwise. The questions related to specific groups are framed slightly differently: whether respondents trust their neighbors or relatives on a scale from 0 to 3 (not at all, just a little, somewhat, a lot). In line with our framing of lower levels of trust as a downside of city life, we define a dummy variable equal to one if respondents say that they trust their neighbors not at all or just a little.

The top panel of Figure 8 plots the proportion of individuals reporting a lack of medicines and food in the top and bottom quartile while the bottom panel shows trust in general and towards relatives; Table A.4 presents the quartile averages. Overall, the data do not support the idea that people find cities to be places with less certain access to food or medicine. To the extent that this represents the effective functioning of social insurance and networks, the data are not consistent with the idea that people in cities face weaker insurance. Individuals residing in the highest-density quartile are eight percentage points less likely to have reported lacking food, and 12 percentage points less likely to have reported lacking medicine. In the majority of countries in our sample, this difference is significant. One explanation is that rural communities face covariate shocks, causing a failure of mutual insurance. But especially in the case of medicine, we would expect a significant idiosyncratic component. To further explore this question, we also used data from the Financial Inclusion Insights program, which surveys a representative sample of individuals in Uganda, Kenya, Nigeria and Tanzania. The survey asks whether an individual could get extra money from relatives in an emergency. In all four

Figure 8: Social Insurance and Anxiety

36
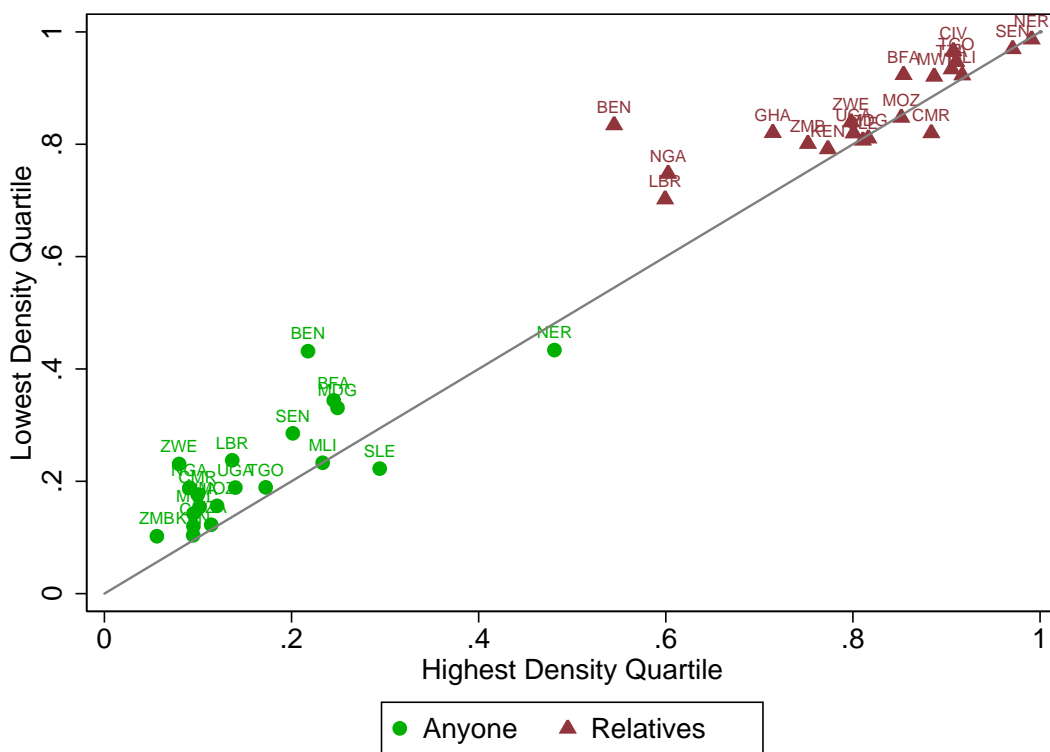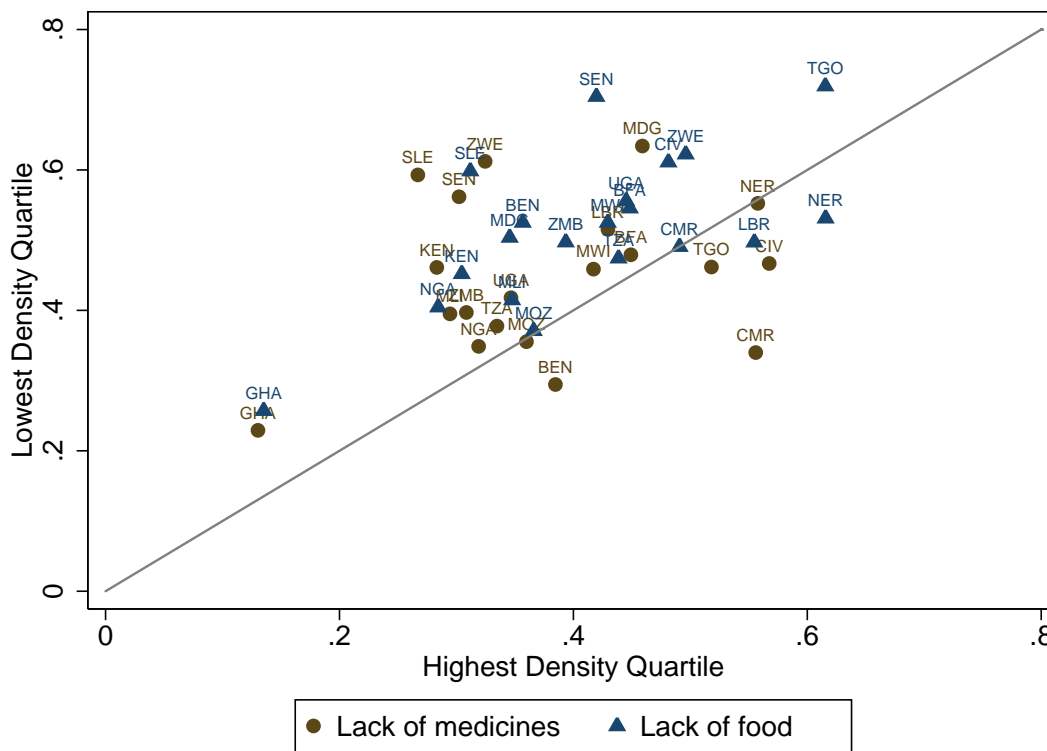
countries, the proportion answering "yes" is higher in the most densely populated quartile (31 percent) compared to the least densely populated quartile (25 percent). A parsimonious explanation is that higher income levels in urban areas lead to strengthening of insurance, though future research is needed on this topic.

As for anxiety, measured levels generally appear to be higher in the countryside than in cities, weakening the argument that people prefer the countryside due to its higher level of tranquillity and reduced stress. However, the differences are significant only in a small number of countries. For trust, however, the story is different. First, all measures of trust – trust towards anyone, relatives, neighbors and one's own ethnic group – are higher in the lowest-population density quartile than in the quartile with the highest population density: cities appear to be low-trust environments. Second, there appear to be very low levels of trust in general among the population: most trusting are individuals in Madagascar, where 30 percent think that people, in general, can be trusted. Other countries have much lower levels of trust; for example, about 90 percent of respondents in Tanzania and Kenya report that "one must be very careful with people." One could imagine that the problem in cities is simply that people live with more-diverse neighbors, including people from other ethnicities, whom they trust less than co-ethnics. But it is striking that trust in co-ethnics also declines in cities relative to rural areas. However, the difference is significant only for a handful of countries.[17]

There are many issues with these kinds of self-reported data, both within the same cultural contexts and even more so across different contexts. Still, the higher level of mistrust offers one possible compensating differential for the many tangible advantages of densely populated areas, but we lack a clear way of quantifying this in monetary terms, to compare with the higher income and other benefits of city life. Nevertheless, we think these results are interesting and potentially point towards avenues for further research.

7. Conclusion

One appealing explanation of the higher real consumption levels of urban areas in the developing world is a "spatial equilibrium," where the higher amenities of rural life are enough to offset the greater monetary rewards of living in cities. In this paper, we go searching for a

---

[17]We are not the first to find that mistrust is higher in denser areas in Africa. Nunn and Wantchekon (2011) use the 2005 Afrobarometer survey, demonstrating that a higher historical exposure to the slave trade appears to be linked to reduced levels of trust in the present day. They control for urban location as classified by the survey, but the coefficient is not reported in the main paper. We geo-locate two further rounds of the same survey and link the data with spatial population density data. Replicating their results, we find that the coefficient on the urban dummy is negative in all their models and highly significant. The patterns found in the two papers are, therefore, similar in that urban locations are associated with higher levels of mistrust.

spatial equilibrium in 20 African countries using new spatial disaggregated evidence on non-monetary amenities. We focus on measures of child health, public goods access, property and violent crime, fear of crime, and outdoor and indoor air pollution, which can be measured consistently across space within and across our countries. We find that almost all metrics in almost all countries are unrelated or improving with population density. The only real exception is crime, though when differences exist they are modest in magnitude and dwarfed in monetary value by the size of the rural-urban wage gaps. When looking at migration rates, every country for which we have data shows that rural-to-urban migrants are far more common than urban-to-rural migrants. That is, net migration is overwhelmingly towards denser areas.

Our findings are hard to reconcile with a spatial equilibrium in the developing world, at least in a static sense. Instead, they point to a world in which individuals in developing countries are moving, on net, to cities, which, on average, offer a better mix of consumption and amenities. This kind of mobility may be limited by a range of frictions, which prevent utility from being instantaneously equalized. In this way, our conclusions are consistent with the hypothesis of Chauvin, Glaeser, Ma, and Tobio (2016) that a spatial equilibrium emerges only when economies are sufficiently developed.

The distinction between a model with spatial equilibrium and a model with frictions to mobility is not simply a theoretical nicety. These two types of models have very different implications for our understanding of rural-urban differences and for policies and government investments. In a world in which urban disamenities account for a spatial equilibrium with big differences in average incomes and productivity between rural and urban locations, policies might focus on increasing well-being still further in urban areas — for example, by increasing the quality of housing and transport, reducing pollution, and otherwise investing in cities. These are surely worthwhile investments, and they would presumably induce further movement from rural areas to urban ones. But these are not investments that would *equalize* utility across locations; if our analysis is correct, such investments would tend to *exacerbate* those disparities in utility.

An alternative approach might focus on reducing frictions to movement – perhaps by addressing the social disruptions and start-up costs associated with mobility. If a model with frictions is better matched to the data, investments of this kind might serve to equalize utility across locations and to allocate labor to its most productive uses. Although our data cannot conclusively reject the existence of a spatial equilibrium, we argue that the preponderance of the evidence suggests that, at least for this set of African countries, the data are more consistent with a model in which developing countries are undergoing dynamic structural change and moving, albeit with frictions, toward more-productive urban areas and non-agricultural activities.

References

ADHVARYU, A., N. KALA, AND A. NYSHADHAM (2014): "Management and Shocks to Worker Productivity: Evidence from Air Pollution Exposure in an Indian Garment Factory," Unpublished Working Paper, University of Michigan.

AHLFELDT, G., S. J. REDDING, D. M. STURM, AND N. WOLF (2015): "The Economics of Density: Evidence from the Berlin Wall," *Econometrica*, 83(6), 2127–2189.

ALBOUY, D. (2012): "Are Big Cities Bad Places to Live? Estimating Quality of Life across Metropolitan Areas," Unpublished Working Paper, University of Michigan.

ALLEN, T., AND C. ARKOLAKIS (2014): "Trade and the Topography of the Spatial Economy," *Quarterly Journal of Economics*, 129(3), 1085–1140.

AVERY, S. (2012): "Lake Turkana and the Lower Omo: Hydrological Impacts of Major Dam and Irrigation Developments," Unpublished Working Paper, African Studies Centre, Oxford.

BANZHAF, H. S., AND R. P. WALSH (2008): "Do People Vote With Their Feet? An Empirical Test of Tiebout's Mechanism," *The American Economic Review*, pp. 843–863.

BAUM-SNOW, N., AND R. PAVAN (2012): "Understanding the City Size Wage Gap," *Review of Economic Studies*, 79(1), 88–127.

BAZZI, S., A. GADUH, A. ROTHENBERG, AND M. WONG (2016): "Skill Transferability, Migration, and Development: Evidence from Population Resettlement in Indonesia," *American Economic Review*, 106(9), 2658–98.

BISHOP, K. C., AND A. D. MURPHY (2011): "Estimating the Willingness to Pay to Avoid Violent Crime: A Dynamic Approach," *American Economic Review Papers and Proceedings*, 101(3), 625–629.

BRUECKNER, J. K., AND S. V. LALL (2015): "Cities in Developing Countries: Fueled by Rural-Urban Migration, Lacking in Tenure Security, and Short of Affordable Housing," in *Handbook of Regional and Urban Economics*, ed. by G. Duranton, J. V. Henderson, and W. Strange, vol. 5B, pp. 1399–1451.

BRYAN, G., S. CHOWDHURY, AND A. M. MOBARAK (2014): "Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh," *Econometrica*, 82(5), 1671–1748.

BRYAN, G., AND M. MORTEN (2015): "Economic Development and the Spatial Allocation of Labor: Evidence From Indonesia," Unpublished Working Paper, Stanford University.

CALIENDO, L., M. DVORKIN, AND F. PARRO (2015): "Trade and Labor Market Dynamics," Unpublished Working Paper, Yale University.

CASELLI, F. (2005): "Accounting for Cross-Country Income Differences," in *Handbook of Economic Growth*, ed. by P. Aghion, and S. Durlauf., 679-741. Elsevier.

CENTER FOR INTERNATIONAL EARTH SCIENCE INFORMATION NETWORK (2015): "Gridded Population of the World, Version 4 (GPWv4), Preliminary Release 2 (2010)," Accessed 23 September 2015; available: http://www.ciesin.columbia.edu/data/gpw-v4, Columbia University.

CHAUVIN, J. P, E. GLAESER, Y. MA, AND K. TOBIO (2016): "What is Different about Urbanization in Rich and Poor Countries? Cities in Brazil, China, India and the United States," *Journal of Urban Economics*.

CICCONE, A., AND R. E. HALL (1996): "Productivity and the Density of Economic Activity," *American Economic Review*, 86(1), 54–70.

COHEN, M. A., R. T. RUST, S. STEEN, AND S. T. TIDD (2001): "Willingness-to-Pay for Crime Control Programs," Unpublished Working Paper, Vanderbilt University.

COMBES, P-P, G. DURANTON, AND L. GOBILLON (2008): "Spatial Wage Disparities: Sorting Matters!," *Journal of Urban Economics*, 63(2), 723–742.

COSTA, D. L., AND M. E. KAHN (2006): "Public Health and Mortality: What Can We Learn from the Past?," in *Poverty, the Distribution of Income, and Public Policy*, ed. by D. C. Alan Auerbach, and J. Quigley. Russell Sage.

CUMMINS, J. R. (2013): "On the Use and Misuse of Child Height-for-Age Z-score in the Demographic and Health Surveys," Unpublished Working Paper, University of California Riverside.

CURRIE, J., E. A. HANUSHEK, M. E. KAHN, M. NEIDELL, AND S. G. RIVKIN (2009): "Does Pollution Increase School Absences?," *Review of Economics and Statistics*, 91(4), 682–694.

CURRIE, J., AND R. WALKER (2011): "Traffic Congestion and Infant Health: Evidence from E-ZPass," *American Economic Journal. Applied Economics*, 3(1), 65.

CUTLER, D., AND G. MILLER (2005): "The Role of Public Health Improvements in Health Advances: the Twentieth-Century United States," *Demography*, 42(1), 1–22.

DEMOMBYNES, G., AND B. OZLER (2002): "Crime and Local Inequality in South Africa," World Bank Policy Research Working Paper 2925.

DESMET, K., AND E. ROSSI-HANSBERG (2013): "Urban Accounting and Welfare," *American Economic Review*, 103(6), 2296–2327.

————— (2014): "Spatial Development," *American Economic Review*, 104(4), 1211–1243.

DHS (2013): "DHS Sampling and Household Listing Manual," Discussion paper, ICF International.

DONALDSON, D., AND R. HORNBECK (2016): "Railroads and American Economic Growth: A "Market Access" Approach," *Quarterly Journal of Economics*, 131(2), 799–858.

DOXSEY-WHITFIELD, E., K. MACMANUS, S. B. ADAMO, L. PISTOLESI, J. SQUIRES, O. BORKOVSKA, AND S. R. BAPTISTA (2015): "Taking Advantage of the Improved Availability of Census Data: A First Look at the Gridded Population of the World, Version 4," *Papers in Applied Geography*, 1(3), 226–234.

ECKERT, F., AND M. PETERS (2017): "Spatial Structural Change," Unpublished Working Paper, Yale University.

FABER, B., AND C. GAUBERT (2016): "Tourism and Economic Development: Evidence from Mexico's Coastline," Unpublished Working Paper, UC Berkeley.

FAFCHAMPS, M., AND C. MOSER (2003): "Crime, Isolation and Law Enforcement," *Journal of African Economies*, 12(4), 625–671.

FAFCHAMPS, M., AND F. SHILPI (2009): "Isolation and Subjective Welfare: Evidence from South Asia," *Economic Development and Cultural Change*, 57(4), 641–683.

FELER, L., AND J. V. HENDERSON (2011): "Exclusionary Policies in Urban Development: Underservicing Migrant Households in Brazilian cities," *Journal of Urban Economics*, 69(3), 253–272.

FERRÉ, C., F. H. G. FERREIRA, AND P. LANJOUW (2012): "Is There a Metropolitan Bias? The Relationship Between Poverty and City Size in a Selection of Developing Countries," *World Bank Economic Review*, 26(3), 351–382.

GEDDES, J. A., R. V. MARTIN, B. L. BOYS, AND A. VAN DONKELAAR (2016): "Long-Term Trends Worldwide in Ambient NO2 Concentrations Inferred from Satellite Observations," *Environmental Health Perspectives*, 124(3), 281–289.

GLAESER, E., AND J. V. HENDERSON (2017): "Urban Economics for the Developing World: An Introduction," *Journal of Urban Economics*, 98, 1–5.

GLAESER, E. L., AND J. D. GOTTLIEB (2009): "The Wealth of Cities: Agglomeration Economies and Spatial Equilibrium in the United States," *Journal of Economic Literature*, 47(4), 983–1028.

GLAESER, E. L., AND D. C. MARE (2001): "Cities and Skills," *Journal of Labor Economics*, 19(2), 316–342.

GOLLIN, D., R. JEDWAB, AND D. VOLLRATH (2016): "Urbanization with and without Industrialization," *Journal of Economic Growth*, 21, 35–70.

GOLLIN, D., D. LAGAKOS, AND M. E. WAUGH (2014): "The Agricultural Productivity Gap," *Quarterly Journal of Economics*, 129(2), 939–993.

GRAFF ZIVIN, J., AND M. NEIDELL (2012): "The Impact of Pollution on Worker Productivity," *American Economic Review*, 102(7), 3652–73.

——— (2013): "Environment, Health, and Human Capital," *Journal of Economic Literature*, 51(3), 689–730.

HAMORY HICKS, J., M. KLEEMANS, N. Y. LI, AND E. MIGUEL (2017): "Reevaluating Agricultural Productivity Gaps with Longitudinal Microdata," Unpublished Working Paper, U.C. Berkeley.

HARRIS, J. R., AND M. P. TODARO (1970): "Migration, Unemployment and Development: A Two-Sector Analysis," *American Economic Review*, 60(1).

HENDERSON, J. V., A. STOREYGARD, AND U. DEICHMAN (2017): "Has Climate Change Driven Urbanization in Africa?," *Journal of Development Economics*, 124, 60–82.

HENDERSON, V., A. STOREYGARD, AND D. WEIL (2012): "Measuring Economic Growth from Outer Space," *American Economic Review*, 102(2), 994–1028.

HERRENDORF, B., AND T. SCHOELLMAN (2014): "Wages, Human Capital, and Structural Transformation," Unpublished Working Paper, Arizona State University.

JEDWAB, R., AND D. VOLLRATH (2017): "The Urban Mortality Transition and Poor Country Urbanization," Unpublished Working Paper, George Washington University.

KAHN, M. E., AND R. WALSH (2015): "Cities and the Environment," in *Handbook of Regional and Urban Economics*, ed. by G. J. Duranton, V. Henderson, and W. C. Strange, vol. 5 of *Handbook of Regional and Urban Economics*, chap. 7, pp. 405 – 465. Elsevier.

KESZTENBAUM, L., AND J.-L. ROSENTHAL (2016): "The Democratization of Longevity: How the Poor Became Old in Paris, 1880-1913," in *New Approaches to Death in Cities during the Health Transition*, ed. by D. R. Fariñas, and M. Oris. Springer, Cham.

KINNEY, P. L., M. G. GICHURU, N. VOLAVKA-CLOSE, N. NGO, P. K. NDIBA, A. LAW, A. GACHANJA, S. M. GAITA, S. N. CHILLRUD, AND E. SCLAR (2011): "Traffic Impacts on PM 2.5 Air Quality in Nairobi, Kenya," *Environmental Science and Policy*, 14(4), 369–378.

KLINE, P, AND E. MORETTI (2014a): "Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority," *Quarterly Journal of Economics*, 129(1), 275–331.

———— (2014b): "People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs," *Annual Review of Economics*, 6, 629–662.

LAGAKOS, D., AND M. E. WAUGH (2013): "Selection, Agriculture, and Cross-Country Productivity Differences," *American Economic Review*, 103(2), 948–980.

LINARD, C., M. GILBERT, R. W. SNOW, A. M. NOOR, AND A. J. TATEM (2012): "Population Distribution, Settlement Patterns and Accessibility Across Africa in 2010," *PLoS One*, 7(2), 1–8.

LUDWIG, J., AND P. J. COOK (2001): "The Benefits of Reducing Gun Violence: Evidence from Contingent-Valuation Survey Data," *Journal of Risk and Uncertainty*, 22(3), 207–226.

MCMILLAN, M. S., AND D. RODRIK (2011): "Globalization, Structural Change and Productivity Growth," NBER Working Paper No. 17143.

MORTEN, M. (2013): "Temporary Migration and Endogenous Risk Sharing in Village India," Unpublished Working Paper, Stanford University.

MORTEN, M., AND J. OLIVEIRA (2017): "The Effects of Roads on Trade and Migration: Evidence from a Planned Capital City," Unpublished Working Paper, U.C. Berkeley.

MUNSHI, K., AND M. ROSENZWEIG (2016): "Networks and Misallocation: Insurance, Migration, and the Rural-Urban Wage Gap," *American Economic Review*, 106(1), 46–98.

NUNN, N., AND L. WANTCHEKON (2011): "The Slave Trade and the Origins of Mistrust in Africa," *American Economic Review*, 101(7), 3221–52.

PEREZ-HEYDRICH, C., J. WARREN, C. BURGERT, AND M. EMCH (2013): "Guidelines on the use of DHS GPS data," Discussion paper, Demographic and Health Surveys.

POPE, C. A. I., AND D. W. DOCKERY (2006): "Health Effects of Fine Particulate Air Pollution: Lines that Connect," *Journal of Air and Waste Management Association*, 56, 709–742.

RAFFO, J. (2015): "MATCHIT: Stata Module to Match Two Datasets Based on Similar Text Patterns," Statistical Software Components, Boston College Department of Economics.

REDDING, S., AND E. ROSSI-HANSBERG (Forthcoming): "Spatial Development," *Annual Review of Economics*.

REDDING, S. J., AND D. M. STURM (2008): "The Costs of Remoteness: Evidence from German Division and Reunification," *American Economic Review*, 98(5), 1766–1797.

RESTUCCIA, D., D. T. YANG, AND X. ZHU (2008): "Agriculture and Aggregate Productivity: A Quantitative Cross-Country Analysis," *Journal of Monetary Economics*, 55, 234–250.

ROBACK, J. (1982): "Wages, Rents, and the Quality of Life," *Journal of Political Economy*, 90(6), 1257–1278.

ROSEN, S. (1979): "Wage-Based Indexes of Urban Quality of Life," in *Current Issues in Urban Economics*, ed. by P. Mieszkowski, and M. Straszheim. Johns Hopkins University Press.

TATEM, A. J., P. W. GETHING, D. L. SMITH, AND S. I. HAY (2013): "Urbanization and the Global Malaria Recession," *Malaria Journal*, 12(1), 133.

TOPEL, R. H. (1986): "Local Labor Markets," *Journal of Political Economy*, 94(3), S111–S143.

VAN DONKELAAR, A., R. MARTIN, M. BRAUER, AND B. BOYS (2015): "Use of Satellite Observations for Long-Term Exposure Assessment of Global Concentrations of Fine Particulate Matter.," *Environmental Health Perspectives*, 123(2), 135–144.

VOLLRATH, D. (2009): "How Important are Dual Economy Effects for Aggregate Productivity?," *Journal of Development Economics*, 88(2), 325–334.

VRIJHEID, M., D. MARTINEZ, S. MANZANARES, P. DADVAND, A. SCHEMBARI, J. RANKIN, AND M. NIEUWENHUIJSEN (2011): "Ambient Air Pollution and Risk of Congenital Anomalies: A Systematic Review and Meta-Analysis," *Environmental Health Perspectives*, 119(5), 598–606.

WHO (2006): *Air Quality Guidelines: Global Update 2005: Particulate Matter, Ozone, Nitrogen Dioxide, and Sulfur Dioxide*. World Health Organization.

——— (2014): *Indoor air quality guidelines: household fuel combustion*. World Health Organization.

WORLD HEALTH ORGANIZATION (2015): "Nutrition Landscape Information System (NLiS): Children 6-23 months of age who receive a minimum acceptable diet," Accessed 1 september 2015; available: here, World Health Organization.

YOUNG, A. (2014): "Inequality, the Urban-Rural Gap and Migration," *Quarterly Journal of Economics*, 129(2), 939–993.

Appendix (For Online Publication)

Figure A.1: Consumption and Amenity Gradients in a Spatial Equilibrium
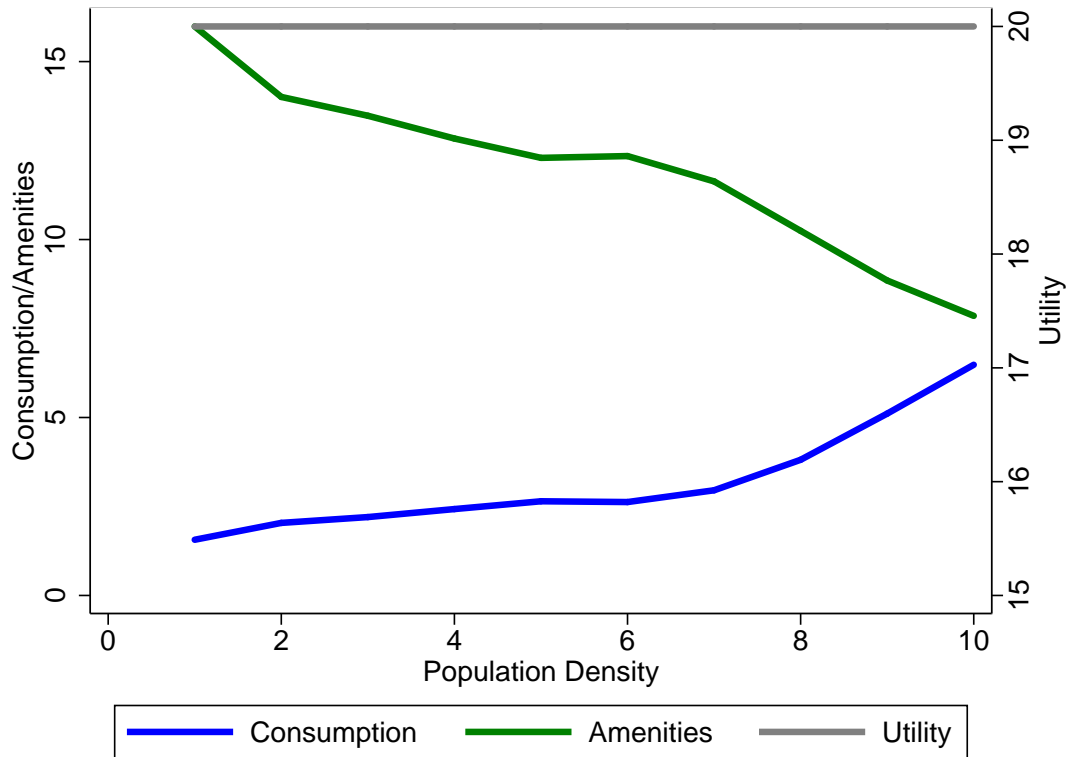
Notes: The figure illustrates the relationship between consumption, amenities, and utility as predicted in a standard spatial equilibrium, described in Section 2, using $\bar{U} = 20$, $\alpha = 0.5$, and $h = 1$. Consumption c is proxied with an asset and housing quality index, taking an average across 20 African countries.

Table A.1: Set of Countries Studied

Country	Households in Sample	Country Population
Benin	17,332	10,050,702
Burkina Faso	13,617	16,460,141
Cameroon	14,189	21,699,631
Dem. Republic of Congo	16,344	65,705,093
Ethiopia	16,037	91,728,849
Ghana	11,574	25,366,462
Ivory Coast	9,394	19,839,750
Kenya	9,033	43,178,141
Liberia	9,333	4,190,435
Madagascar	17,578	22,293,914
Malawi	24,210	15,906,483
Mali	10,105	14,853,572
Mozambique	13,899	25,203,395
Nigeria	38,170	168,800,000
Senegal	7,780	13,726,021
Sierra Leone	12,629	5,978,727
Tanzania	9,282	47,783,107
Uganda	8,939	36,345,860
Zambia	7,164	14,075,099
Zimbabwe	9,442	13,724,317
Total	276,051	769,082,846

Figure A.2: Housing Quality
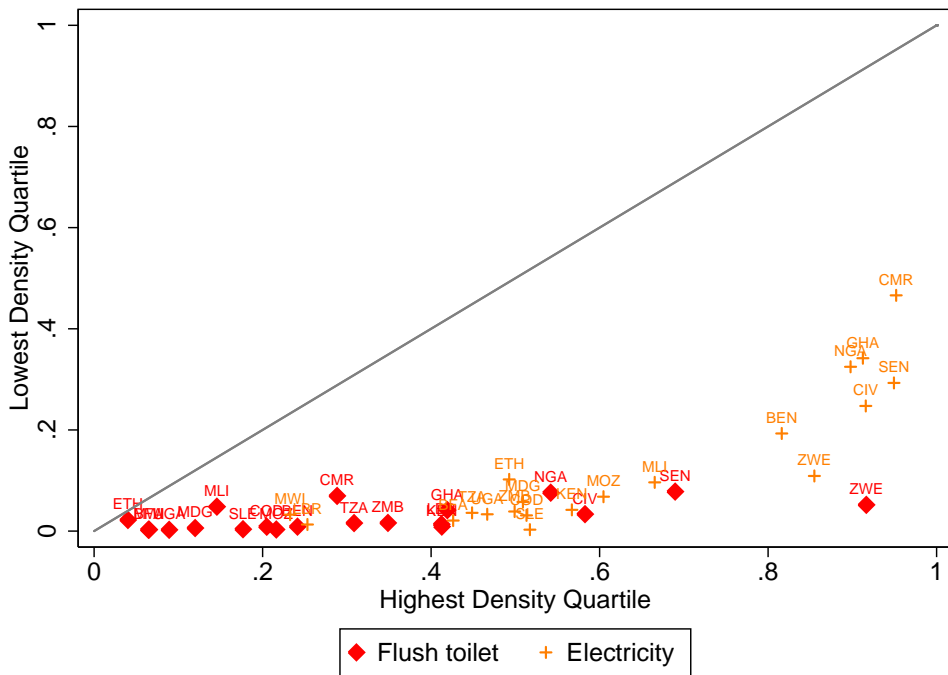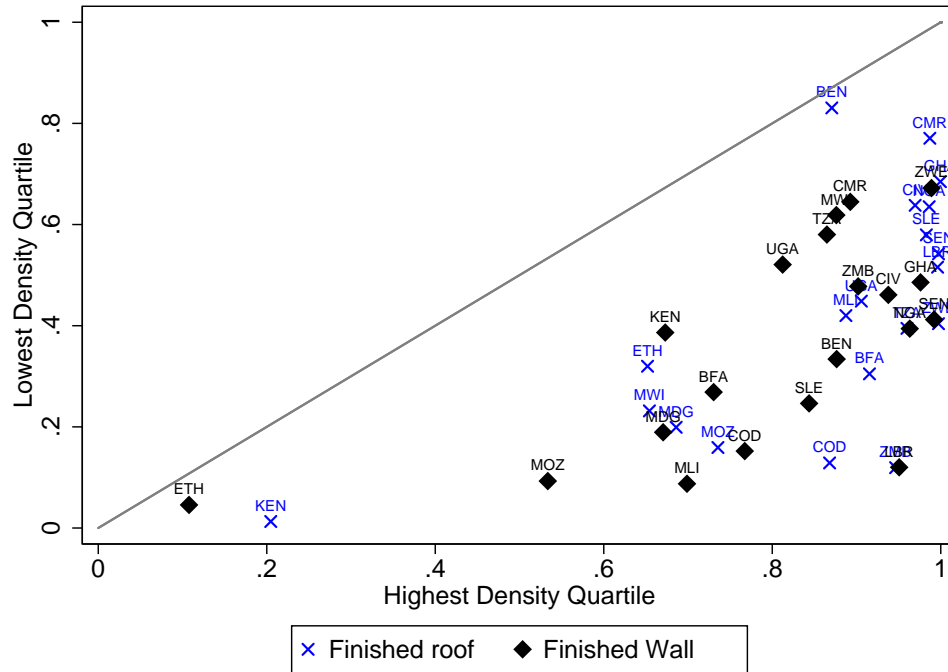
Table A.2: Child Health Characteristics by Density

| | Population Density Quartile | | | |
	Q1	Q2	Q3	Q4
Anemic	.69	.68	.64	.6
		0-0[†]	3-0[†]	7-0[†]
Stunted (low height for age)	.4	.4	.38	.29
		0-0	3-1	11-0
Wasted (low height for weight)	.1	.1	.09	.09
		1-0	3-1	6-1
Below minimum diet	.93	.91	.9	.87
		1-0	2-0	5-0

Note: This table reports child health across quartiles across our set of 20 countries. The two numbers below the averages x-y represent the number of countries with a difference between the least dense quartile (Q1) that is statistically significant at the one-percent level and either negative (x) or positive (y).

[†] We do not have data on anemia for Kenya, Liberia, Nigeria and Zambia, so the total number of countries for this variable is 16.

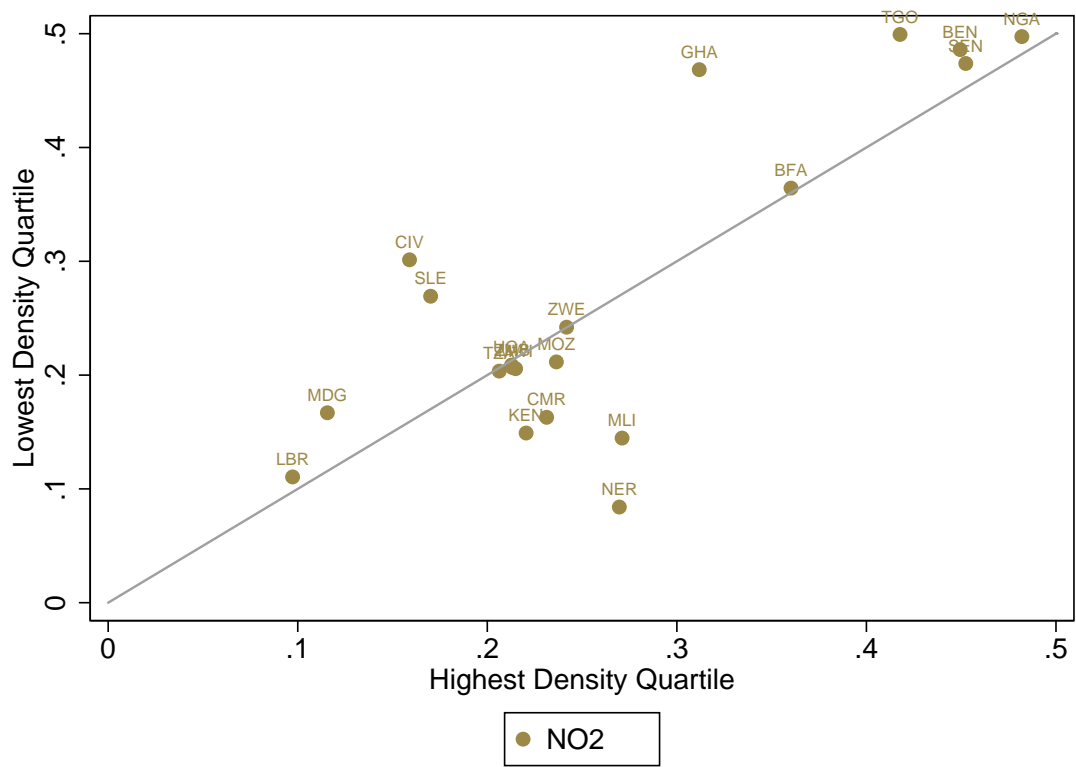Figure A.3: Outdoor Air Pollution – NO2

Table A.3: Pollution Differences by Density Quartile

	Population Density Quartile			
	Q1	Q2	Q3	Q4
Nitrogen Dioxide	.27	.28	.26	.27
		6-7	8-6	9-6
PM2.5	20.1	20.42	19.1	18.56
		4-6	9-5	12-6
Cook inside with solid fuel	.63	.64	.58	.38
		0-4†	5-3†	15-1†

Note: This table reports pollution exposure quartiles across our set of 20 countries. The two numbers below the averages x-y represent the number of countries with a difference between the least dense quartile (Q1) that is statistically significant at the one-percent level and either negative (x) or positive (y). Here, we use within-country quartiles rather than across-country quartiles to ensure that all countries have observations across quartiles; results are similar if we use across-country quartiles and exclude countries that do not have observations across all quartiles. † The Tanzania DHS does not contain information on the type of cooking fuel and indoor cooking, so the total number of countries for this variable is 19.

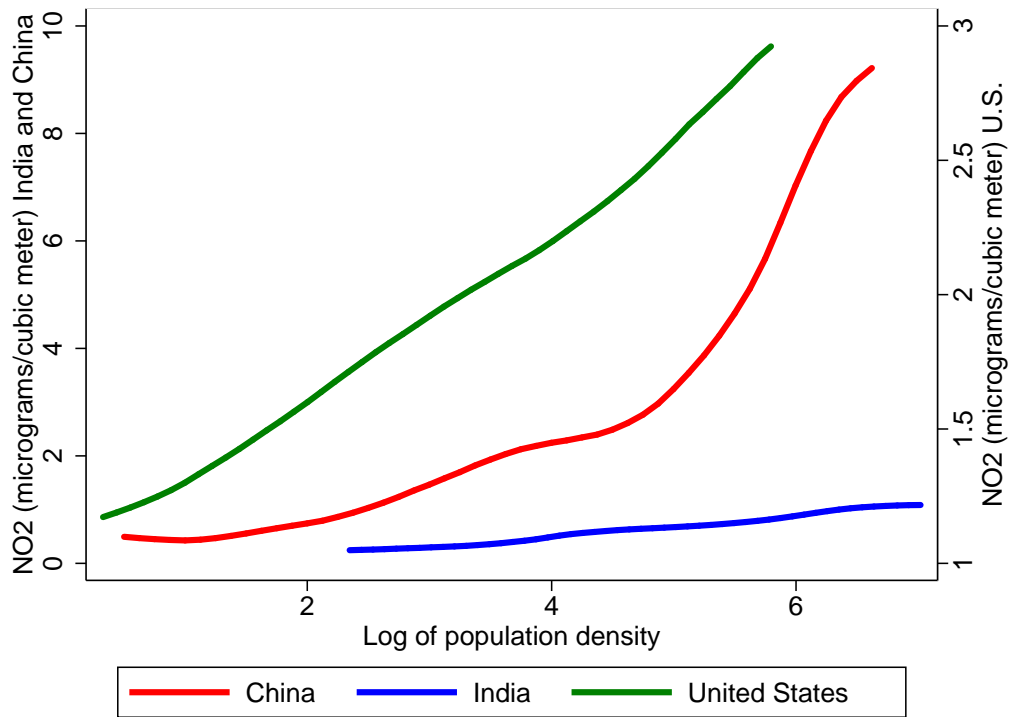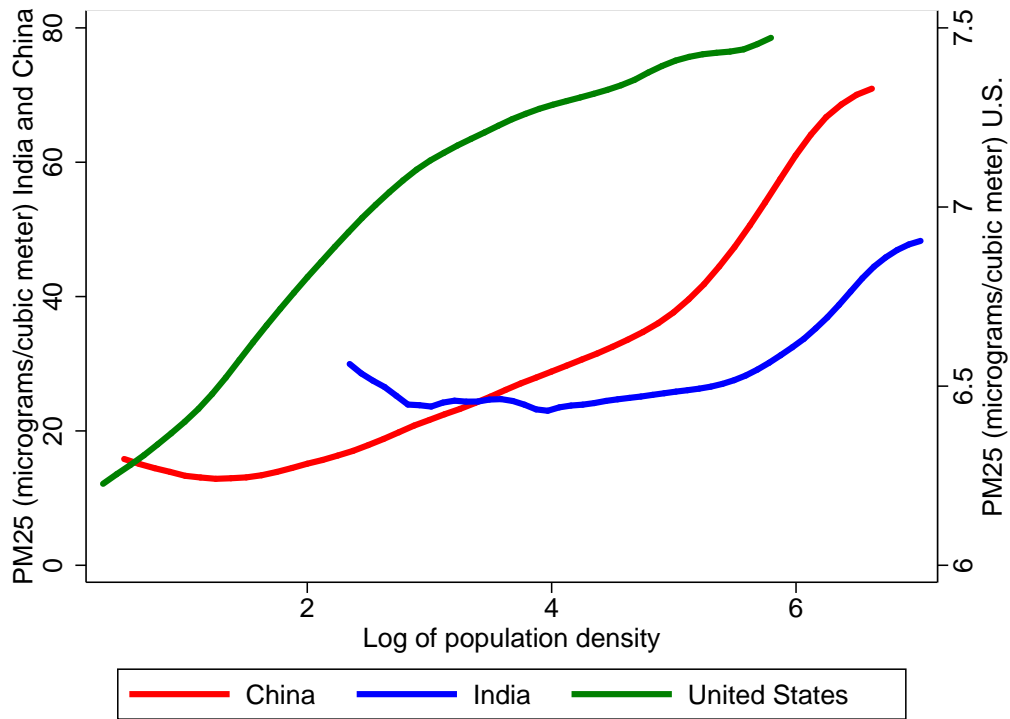Figure A.4: Outdoor Air Pollution in China, India and the United States

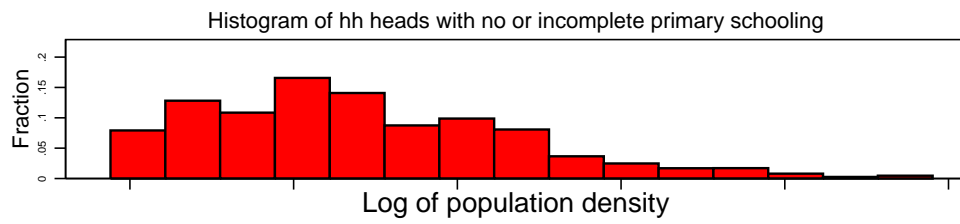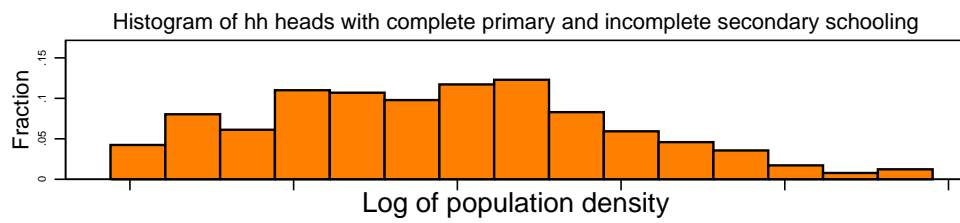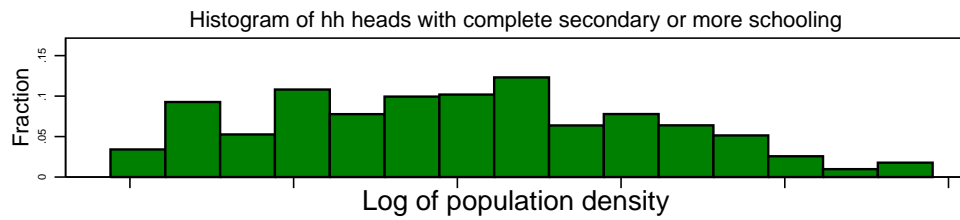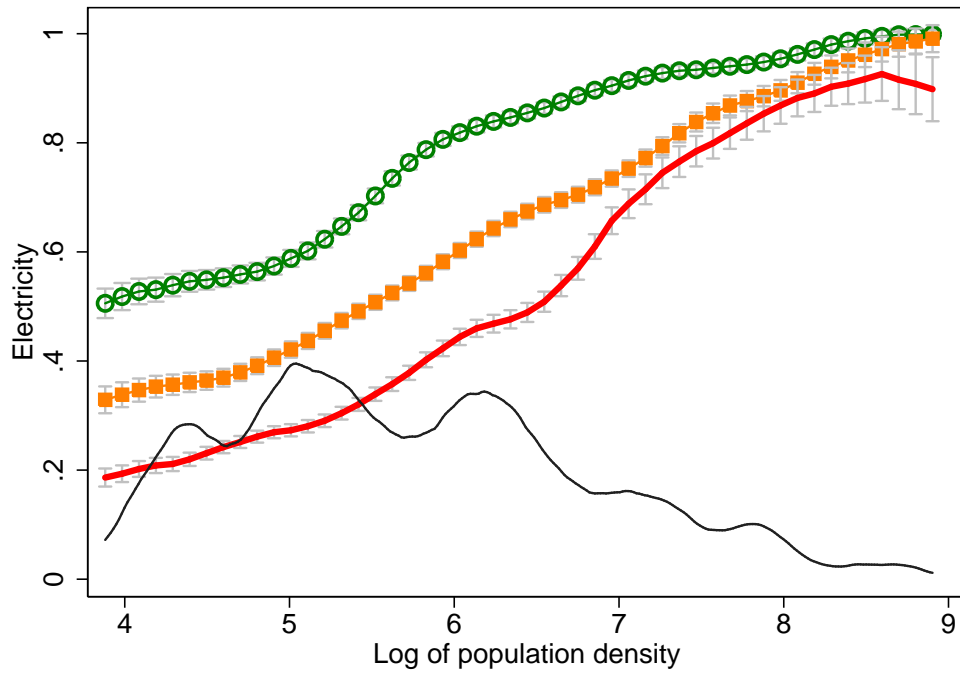Figure A.5: Electricity by Highest Education of Household Head in Nigeria

Histogram of hh heads with complete secondary or more schooling

Histogram of hh heads with complete primary and incomplete secondary schooling

Histogram of hh heads with no or incomplete primary schooling

Table A.4: Secondary Amenity Differences by Density Quartile

	Population Density Quartile			
	Q1	Q2	Q3	Q4
Lack of food	.44	.47	.43	.36
		1-1	2-1	11-1
Lack of medicine	.52	.52	.47	.4
		0-1	1-0	11-0
Anxiety	.34	.34	.33	.29
		0-1[†]	2-1[†]	3-0[†]
Trust anyone	.22	.22	.22	.14
		0-1	2-1	7-0
Trust relatives	.87	.85	.83	.81
		3-0	3-0	5-0
Trust neighbor	.67	.65	.6	.52
		2-0	5-0	9-0
Trust own ethnic group	.67	.65	.59	.55
		0-0[†]	1-1[†]	5-0[†]

Note: The first three columns report the average differences from the second, third and fourth density quartiles relative to the first (least dense) quartile across our set of 20 countries. The fourth column reports the average standard deviation across regions across our set of 20 countries. Numbers below the average differences are the number of countries with a difference that is statistically significant at the one-percent level.

[†] Data on anxiety and trust towards one's own ethnic group is not available for Burkina Faso, Cameroon, Cote d'Ivoire, Liberia, Niger, Sierra Leone and Togo; Zimbabwe also does not record information on trust towards one's own ethnic group. The total number of countries for anxiety and trust towards one's own ethnic group is therefore 13 and 12, respectively.

# A. Demographic and Health Surveys (DHS) Data

1.1. DHS sample

While the DHS aim to make survey instruments and samples comparable across countries, the exact sampling differs according to the particular survey.[18] The target population of most DHS surveys are women aged 15-49 and children under the age of five living in residential households with the most common sampling following a two-stage cluster sampling procedure (DHS, 2013). If a recent census is available, the sampling frame of the census is used to define primary sampling units which are usually enumeration areas. Alternative sample frames include lists of electoral zones, estimated structures per pixel derived from high-resolution satellite imagery or lists of administrative units. Clusters will then be stratified depending on the number of domains that are desired for the particular survey, where a typical stratification is first at the geographical level and then at rural/urban clusters. In the first stage, from each of the strata a random sample of enumeration areas is selected inversely proportional to size. Unless a reliable listing of households exists, households will be listed for each of the selected primary sampling units. In the second stage, households are selected with equal probability.

If the sampling frame is not specifically selected to match the population along the lines of population density, it is likely that the distribution of the survey sample according to population density might not match that of the entire population. In practice, the cases we have examined show very little effort to oversample or undersample with respect to population density. For Tanzania, we can compare the population density distribution of the Afrobarometer and DHS clusters with those of the overall population from the census data, where we weight the population density of enumeration areas by the population. As is evident from Figure A.6, both the Afrobarometer survey and the DHS appear to capture a sample that covers a wide range of population densities.

---

[18]For further information, see: http://www.dhsprogram.com.

Figure A.6: Distribution of population, DHS and Afrobarometer respondents in Tanzania
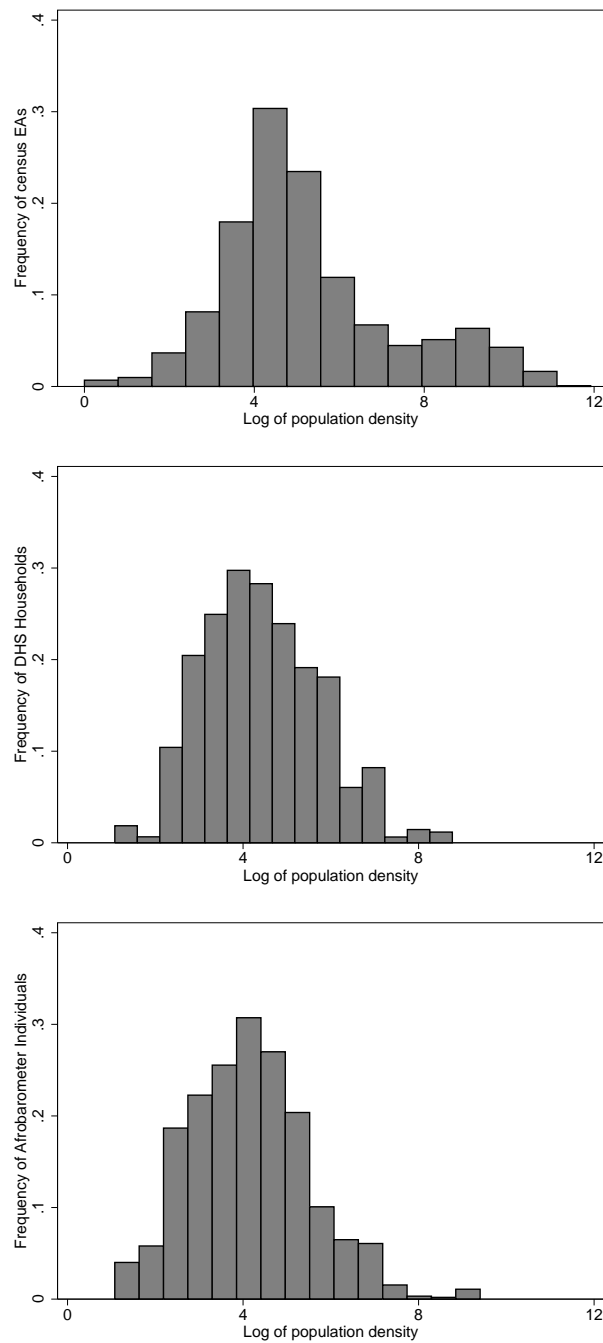
Notes: The top figure shows the distribution of the population using the 2002 enumeration area census data and the total population in each enumeration area as sample weights. The middle graph shows the distribution of population densities from the DHS data. The bottom graph shows distribution of clusters from Afrobarometer data. For expositional simplicity the top graph excludes 112 enumeration areas that have a log of population density above 12.

Table A.5: Surveys

	Household	Malaria	Migration	Crime (Afrobarometer)	Crime (LSMS)
Benin	Benin 2011-12 Standard DHS	Benin 2011-12 Standard DHS		x	
BurkinaFaso	Burkina Faso 2010 Standard DHS	Burkina Faso 2010 Standard DHS			
Cameroon	Cameroon 2011 Standard DHS	Cameroon 2011 Standard DHS			
DRC	Congo Democratic Republic 2013-14 Standard DHS	Congo Democratic Republic 2013-14 Standard DHS	Congo Democratic Republic 2007 Standard DHS		
Ethiopia	Ethiopia 2011 Standard DHS	Ethiopia 2011 Standard DHS	Ethiopia 2005 Standard DHS		2013/14 Ethiopian Socioeconomic Survey
Ghana	Ghana 2008 Standard DHS		Ghana 2008 Standard DHS	x	
IvoryCoast	Cote d'Ivoire 2011-12 Standard DHS	Cote d'Ivoire 2011-12 (14) Standard DHS			
Kenya	Kenya 2008-09 Standard DHS		Kenya 2008-09 Standard DHS	x	
Liberia	Liberia 2013 Standard DHS	Liberia 2011 MIS	Liberia 2007 Standard DHS		
Madagascar	Madagascar 2008-09 Standard DHS	Madagascar 2013 MIS DHS-VI	Madagascar 2008-09 Standard DHS	x	
Malawi	Malawi 2010 Standard DHS	Malawi 2012 MIS	Malawi 2010 Standard DHS	x	LSMS 2004/05
Mali	Mali 2012-13 Standard DHS	Mali 2012-13 Standard DHS	Mali 2006 Standard DHS	x	
Mozambique	Mozambique 2011 Standard DHS	Mozambique 2011 Standard DHS		x	
Nigeria	Nigeria 2013 Standard DHS	Nigeria 2010 MIS	Nigeria 2008 Standard DHS	x	NGHS, Panel Wave 2, 2012-2013; Post-harvest household questionnaire
Senegal	Senegal 2010-11 Standard DHS	Senegal 2010-11 Standard DHS	Senegal 2005 Standard DHS	x	
SierraLeone	Sierra Leone 2013 Standard DHS		Sierra Leone 2008 Standard DHS		
Tanzania	Tanzania 2010 Standard DHS	Tanzania 2011-12 Standard AIS		x	Tanzania NPS 2008
Uganda	Uganda 2011 Standard DHS	Uganda 2009 MIS		x	Uganda NPS 2009/10
Zambia	Zambia 2007 Standard DHS		Zambia 2007 Standard DHS	x	
Zimbabwe	Zimbabwe 2010-11 Standard DHS				

1.2. Linking DHS data with population density data

To link the DHS data with population density, we draw a buffer of 5km around each cluster and extract the average population density around each cluster. We perform these calculations in WGS1984, since the different areas of the pixel sizes when moving away from the equator has been taken into account when constructing the population density grid, which is defined as the population count divided by the area. Many urban DHS clusters are in proximity closer than 5 km so that buffer polygons around clusters are overlapping. Therefore, we compute our zonal statistics using the Spatial Analyst Supplemental Tools in ArcGIS, a supplemental toolbox that allows computing zonal statistics for overlapping polygons. All computations were performed in ArcGIS 10.4.

1.3. Spatial linking of DHS

Figure A.7: DHS clusters in Dar Es Salaam
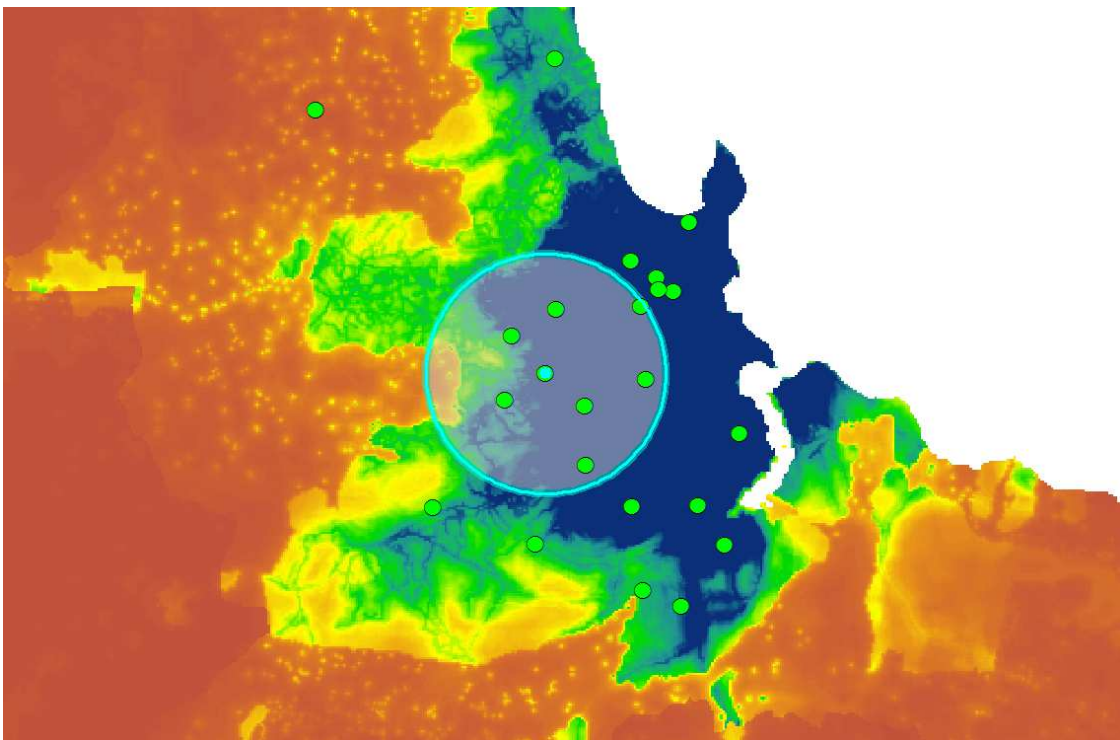
Notes: Figure shows a 5km circle around a DHS cluster in Dar Es Salaam; the gridded data come from WorldPop.

1.4. DHS Variables

The main variables to measure durables are whether the household has a television (hv208=1), mobile phone, or landline (hv221=1 or hv243a=1). For housing consumption, we examine

whether the household has electricity (hv206=1), running water (hv201<21), a constructed floor (hv213>13), a flush toilet (hv205<20), finished walls (hv214>=30), a finished roof (hv215>=30) and the number of sleeping rooms per adult over five (hv216/(hv009-hv014)). For the highest level of education completed by the household head we use variable hv109. Indoor cooking is determined using variables hv226 and hv241. Variable hv204 captures the time required to fetch water.

The World Health Organization (2015) defines a minimum acceptable diet as follows:

"The composite indicator of a minimum acceptable diet is calculated from: (i) the proportion of breastfed children aged 6-23 months who had at least the minimum dietary diversity and the minimum meal frequency during the previous day and (ii) the proportion of non-breastfed children aged 6-23 months who received at least two milk feedings and had at least the minimum dietary diversity not including milk feeds and the minimum meal frequency during the previous day.

Dietary diversity is present when the diet contained four or more of the following food groups:

- grains, roots and tubers;

- legumes and nuts;

- dairy products (milk, yogurt, cheese);

- flesh foods (meat, fish, poultry, liver or other organs);

- eggs;

- vitamin A-rich fruits and vegetables; and

- other fruits and vegetables.

The minimum daily meal frequency is defined as:

- twice for breastfed infants aged 6-8 months;

- three times for breastfed children aged 9-23 months;

- four times for non-breastfed children aged 6-23 months."

To compute a child's minimum dietary diversity from the DHS, we use data on breastfeeding (m4), age of the child (hw1), the number of times the child ate (m39), the type of food groups the child consumed (v414a-v414j), the consumption of powdered/tinned/fresh milk and infant

formula (469e and 469f). If the survey is from Phase 5, there is no information on milk feedings for children who are no longer breastfed; in this case, we calculate the minimum acceptable diet only for children who are breastfed.

Migration status is determined using the years lived in the current location (v104) and the type of the previous residence which is classified into capital, large city; city; town; countryside; and abroad (v105). Following Young (2014), we exclude individuals who lived abroad, and check that all variables are coded consistently across countries; for example, abroad is sometimes coded as 4 and sometimes as 5.

Our data for malaria incidence comes from a combination of DHS, MIS and AIS. In a subset of countries, blood samples were collected for children aged 6-59 months as part of the DHS data collection in households that were selected for the men's questionnaire (every one out of eight households). If the data are not available as part of the DHS, we use data from the most recent geo-referenced AIS survey or Malaria Indicators Survey (MIS) during which blood samples from children aged 0-59 or 6-59 months were taken depending on the specific survey. Table A.5 shows the exact survey used for each country. Malaria tests were administered via rapid diagnostic testing and blood smear microscopy.[19] We construct a dummy variable that is equal to one if the malaria rapid test for a child was positive (hml25).

B. Afrobarometer

2.1. Geo-locating Afrobarometer respondents

Afrobarometer surveys collect data on attitudes towards democracy and governance, as well as a range of other quality-of-life measures.[20] The Afrobarometer surveys do not record coordinates of respondents, but record the village, district and region names. The 2011 round provides four different administrative names. We use a matching algorithm that matches village names and other provided administrative names to locations as listed in gazetteers; specifically, we follow Nunn and Wantchekon (2011) and use the geonames gazeteer available on www.geonames.com. This website provides a list of locations, each assigned an ID along with several names: the geographical name of the point in utf8 and plain ascii characters; alternative names and the associated latitude and longitude coordinate. There is also auxiliary information such as the modification date of each entry, administrative codes, elevation, and feature classes. If a name is associated with several entries, we keep the most recent entry.

Our matching algorithm uses a mixture of exact matches and fuzzy matches in multiple stages

---

[19]Note that in Madagascar, areas without malaria have been excluded from the survey.
[20]For further information, see http://www.afrobarometer.org.

(depending on the survey round, between 13 and 21). Whenever a location name is identified, we assign it the latitude and longitude and remove it from the dataset that is fed into the next stage. In essence, matching is achieved in the following way: first, we perform a series of exact matches based on the village name from Afrobarometer with the asciiname listed in the gazetteer; if there are no exact matches with the village name and the asciiname, we search through the next four alternative names listed in the gazeteer for the specific location. In this first stage we find almost forty percent of locations. We then use the most precise administrative classification. For example, if the data set has information on the village name, district and region, this would be the district. We perform the exact same series of matches on the district name, using again the ascii name and four possible alternative names listed in the gazetteer. In rounds three and four of the survey in which we have only district and region names in addition to the village names, this step finds 49–52 percent of the locations.

Third, we match on the region name which finds another four to six percent of the sample. Finally, to catch any remaining misspellings, we perform a fuzzy match based on similar text patterns between the village name and the ascii name using a command developed by Raffo (2015). We use a similarity score of above 0.70 and a vectorial decomposition algorithm (3-gram). This finds another one to three percent of locations. In total, we are able to match between 92 and 95 percent of individuals in each round.

In addition to random checks of the identified locations, we use the 2005 data to check the consistency between our algorithm and the location data of Nunn and Wantchekon (2011). For the subset of locations for which they provide geo-locations, we find that the median distance between their location and our locations is 12.5 km. Further, considering that the population density data vary largely at the district and region levels, we expect the difference to be even smaller when looking at the resulting population densities. Indeed, the correlation coefficient between the population density from their and our data is 0.65 with a p-value of 0.000.

We use variables related to feeling unsafe walking in one's neighborhood (q9a), fear of crime in one's own home (q9b); theft (q9b); physical attack (q9c); trust in general (q83); trust towards relatives (q84a); trust toward neighbors (q84b); trust toward one's own ethnic group (q84c); frequency of lack of food (q8a) and medicine (q8c); and anxiety (q96b).

C. Pollution data

Our pollution data for PM2.5 and NO2 concentrations come from van Donkelaar, Martin, Brauer, and Boys (2015) and Geddes, Martin, Boys, and van Donkelaar (2016), respectively. As the date for the Gridded Population of the World v4 (GPWv4) data is approximately 2010,

we take the pollution measures that are closest in time: the tri-annual mean (2009-2011) for both PM2.5 series; for NO2, we have the exact year 2010. The estimated PM2.5 and NO2 concentrations are available at a resolution of 0.1 decimal degrees (about 10km at the equator). We construct a fishnet of the same resolution and, for each pixel, compute the average pollution measure and the average population density from the GPWv4. PM2.5 is measured in $\mu g/m^3$, while NO2 is measured in ppb (parts per billion). Following Vrijheid, Martinez, Manzanares, Dadvand, Schembari, Rankin, and Nieuwenhuijsen (2011), we use a conversion of 1ppb= 1.88 f $\mu g/m^3$, which assumes ambient pressure of 1 atmosphere and a temperature of 25 degrees celsius.

Table A.6: Afrobarometer Sample

	Individuals	Round 3	Round 4	Round 5
Benin	3,543	x	x	x
Burkina Faso	2,255		x	x
Cameroon	1,072			x
Cote D'Ivoire	1,192			x
Ghana	4,089	x	x	x
Kenya	4,659	x	x	x
Liberia	2,282		x	x
Madagascar	3,881	x	x	x
Malawi	4,784	x	x	x
Mali	3,663	x	x	x
Mozambique	4,744	x	x	x
Niger	1,199			x
Nigeria	6,961	x	x	x
Senegal	3,596	x	x	x
Sierra Leone	1,190			x
Tanzania	4,791	x	x	x
Togo	1,056			x
Uganda	7,191	x	x	x
Zambia	3,590	x	x	x
Zimbabwe	4,344	x	x	x

Note: Column (2) shows the number of individuals in our sample for each of the countries; columns (3)–(5) indicate when a country was added to the Afrobarometer sample. Round 3 took place in 2005, Round 4 in 2008, and Round 5 in 2011.

Figure A.8: Pollution and population density in Nigeria
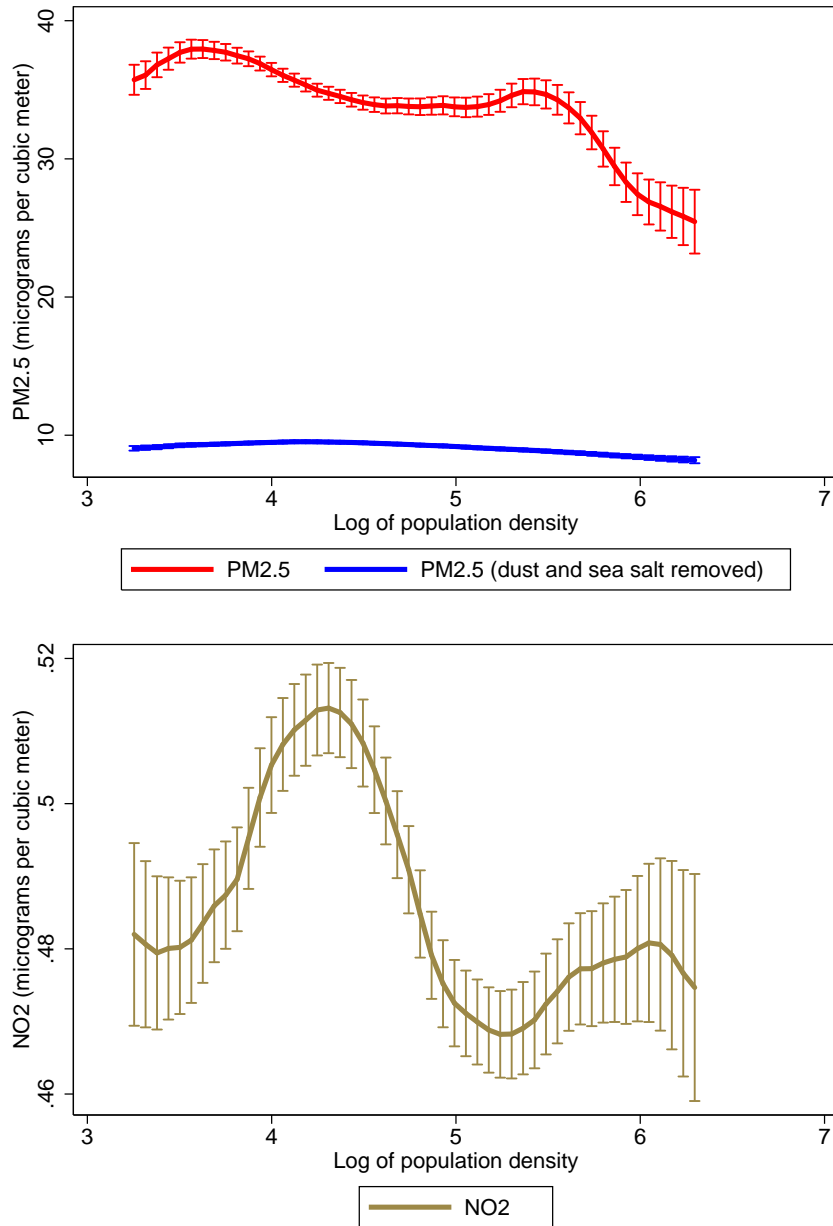
Notes: The figure shows a kernel-weighted local polynomial regression of the level of pollution on the log of population density in Nigeria using data from the entire country, and plotting 95 percent confidence intervals. The top panel shows the results for PM2.5, and the bottom panel shows NO2 levels across population density space. Taking the log of population density removes uninhabited pixels. We remove the top and bottom five percentiles of the population density distribution.

Figure A.9: Crime - LSMS

Ethiopia
Theft, robbery, other violence

Malawi
Anything stolen, personal attack

Nigeria
Theft/kidnapping/hijacking/robbery/assault

Tanzania
Anything stolen/personal assault

Uganda
Theft/violence/conflict