



Do urban wage premia reflect lower amenities? Evidence from Africa[☆]

Douglas Gollin^a, Martina Kirchberger^{b,*}, David Lagakos^c

^a University of Oxford and CEPR, United Kingdom

^b Trinity College Dublin, Ireland

^c Boston University and NBER, United States

ABSTRACT

In most developing countries, wages are substantially higher in cities than in rural areas. One possible explanation is that the higher wage levels of urban areas are offset by lower non-monetary amenities. This paper draws on new high-resolution evidence to document how non-monetary amenities vary between urban and rural areas within 20 Sub-Saharan African countries. We focus on measures of public goods, crime and pollution. We find that in almost all countries, and for almost all measures, the quality of amenities is at least as high in cities as in rural areas. This finding casts doubt on the hypothesis that urban wage premia in the developing world represent a compensation for lower amenities.

1. Introduction

Until around a century ago, the cities of Europe and North America were death traps (Costa and Kahn, 2006). Life expectancy for a newborn in Paris in the 19th century was around five years lower than for a child born in the rest of France (Kesztenbaum and Rosenthal, 2016). Londoners died of infectious diseases, like cholera and typhoid, at a much higher rate than did their rural counterparts (Riley, 2001). Cities were packed with people and consequently full of human and animal waste; the advent of industrial production injected toxic chemicals and particulate pollution into the mix. Perhaps unsurprisingly, urban wages at the time were high, presumably providing some compensating differentials for the high living costs, disamenities, and health risks of urban life. Indeed, Williamson (1982) documents a clear positive correlation between real wage levels and infant mortality rates across 19th century British towns.

To what extent do the same dynamics characterize urban life in today's developing countries? A body of recent evidence documents that consumption, income, and real wage levels are substantially higher in

cities than in rural areas (Ferré et al., 2012; Young, 2013; Gollin et al., 2014; Herrendorf and Schoellman, 2018). Do these wage premia serve to compensate urban residents for utility losses from unhealthy, unsafe, and unpleasant living conditions? Many observers look at the difficult conditions in developing world slums today and conclude that rural-urban migrants today face a trade-off similar to the one which confronted people historically in today's rich countries (e.g. Davis, 2006; Boo, 2014).

In this paper, we ask whether rural-urban income gaps in the developing world today should similarly be understood as compensation for disamenities of urban life. We draw on new and detailed spatial evidence on three categories of non-monetary amenities that are frequently invoked as prime candidates for rural-urban differences in real consumption in developing countries. These categories are: public goods provision, crime, and air pollution. All three of these could plausibly generate urban disamenities that might serve to offset the income and consumption advantages of urban life. For instance, vast settlements within urban areas in the developing world may lack access to amenities like electricity and piped water; and in fact the absence of these amenities is often

[☆] For helpful comments we thank the editor, Ed Glaeser, two anonymous referees, David Albouy, Deborah Balk, Joe Cummins, Banu Demir-Pakel, Stefan Dercon, Dave Donaldson, Cheryl Doss, Ben Faber, Marcel Fafchamps, Roland Hodler, Darby Jack, Kelsey Jack, Matt Kahn, Pat Kinney, Pramila Krishnan, Horacio Larreguy, Ethan Ligon, Matthew Neidell, Natalia Ramondo, Chris Small, Daniel Sturm, Luke Valin, David Weil, and seminar participants at Barcelona, Belfast, Bristol, Columbia, Copenhagen, Edinburgh, the Ethiopian Development Research Institute, the IMF, LSE, Notre Dame, Nottingham, Sussex, Trinity College Dublin, UC Riverside, as well as conference participants at the Edinburgh Structural Change Conference, the NEUDC, CEPR and CSAE conferences, the NBER Summer Institute, the 5th Urbanization and Poverty Reduction Research Conference, and the Workshop on Geodata and Economics. For excellent research assistance we thank Jessie Wang, Eoin Dignam and Abed Mutemi, as well as the Bill and Melinda Gates Foundation for supporting Mutemi's work on this project. We would like to thank Jeff Geddes and Aaron van Donkelaar for sharing the pollution data and their advice. Kirchberger is grateful for many helpful conversations and suggestions while hosted by the Center for International Earth Science Information Network at Columbia University. This paper was partly supported through the Global Research Program on Spatial Development of Cities, funded by the Multi Donor Trust Fund on Sustainable Urbanization of the World Bank and supported by the UK Department for International Development. All potential errors are our own.

* Corresponding author.

E-mail addresses: douglas.gollin@qeh.ox.ac.uk (D. Gollin), martina.kirchberger@tcd.ie (M. Kirchberger), lagakos@bu.edu (D. Lagakos).

used to define urban “slums” (Marx et al., 2013). Urban areas are frequently thought to have higher rates of crime, as is the case in the United States (see Glaeser and Sacerdote, 1999). Air pollution could relate to the presence of more automobiles or more concentrated industrial activity. Indeed, in the United States and Western Europe, urban areas have long suffered from air pollution; even in pre-industrial times, rural areas were seen as places where people could escape the odors and “miasmas” of urban air (see e.g. Cutler and Miller, 2005). In this sense, better air quality has long been viewed as an important positive amenity of rural life. If these disamenities are more prevalent in urban areas than in rural areas, then we should arguably interpret rural-urban income gaps as compensating differentials; otherwise, we may need to invoke different explanations for the ubiquity of gaps between rural and urban incomes and living standards.

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 micro surveys on health, public goods, crime, and other reported living standards, and satellite-derived measures of pollution with gridded population density. An advantage of our approach is that we do not rely on administrative definitions of “urban” and “rural.” These binary categories often reflect arbitrary boundaries and cut-offs that are inconsistent across space and time. Our data allow us to examine the distribution of amenities and disamenities across locations that differ in population density. We can see how amenities vary for people living in locations with differing population densities, along a continuum. We focus our attention on 20 Sub-Saharan African countries that together represent about three-fourths of the region’s population.

Our empirical analysis begins by examining a large number of outcomes that are closely related to private consumption or expenditure. These include measures of housing quantity and quality, durable goods ownership, and some other measures influenced by income and wealth. Consistent with previous research, we find strong support for the proposition that households in urban areas are better off, by these measures, than households in rural areas. We add to the previous literature in showing that almost all these measures improve more or less monotonically with population density across the communities within each country. This implies that the effect does not come simply from those living in the largest cities or the wealthiest communities.¹

Next, we turn to amenities that are either publicly provided or jointly produced from public and private inputs. Our main empirical finding is that, within the countries that we study, almost all amenities that we consider improve at least weakly with population density. Publicly provided utilities like electricity, piped water, and sewage systems are far more likely to be found in areas with higher population density. The same is true for paved roads. Schools are about equally distributed through density space, but the available evidence suggests that the quality of schooling is higher, if anything, in urban areas. Police stations are more prevalent in urban areas, although the quality of police services is challenging to measure. Health facilities are more common in urban areas, and child health outcomes, which at least in part reflect the quality of healthcare services, are better in urban areas.

Property crime and violent crime are, unfortunately, high throughout density space in most countries. These categories of crime appear to be marginally higher in urban areas, but the differences are statistically insignificant 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. Using plausible values of willingness to pay to avoid crime, the differences between rural and urban areas are dwarfed by the much larger gaps in average income.

Air pollution – perhaps surprisingly – is similar or worse in rural areas in the majority of our countries. The reason is that rural areas in sub-Saharan Africa – or at least in our 20 countries – tend to be exposed to dust blown from the continent’s deserts. Although this is a naturally occurring phenomenon, the particulate matter originating in this fashion appears to have the same health consequences as anthropogenic particulates pollution. Even after excluding desert dust and sea salt, however, concentrations of fine particulate matter are not higher on average in urban areas in our data. Nitrogen dioxide concentrations – largely the result of industrial pollution – are no higher in urban areas than in rural areas, serving as a reminder, in case one were needed, of how under-developed Africa’s manufacturing sector is.

Thus, at least for these three categories of amenities, it seems unlikely that rural-urban income gaps simply compensate for disamenities of urban life. Still, we obviously cannot address all possible amenities, and there could be some other undesirable non-monetary feature of urban life that compensates for its higher real wages. To address this possibility, at least in part, we consider two additional measures of well-being: subjective welfare measures and rural-urban migration rates. We show, for a smaller set of countries, that reported happiness and life satisfaction are generally higher on average in urban areas than in rural areas. In addition, people are voting with their feet: in almost every country we examine, net migration flows show strong movements of people from rural areas to towns and cities. Migrants to cities appear to be drawn by the greater amenities as well as the higher nominal and real wages.

Put differently, our findings suggest that urban disamenities do not obviously explain the urban-rural consumption gaps in Africa. Instead, a more plausible explanation may be that spatial wage gaps reflect frictions – perhaps stemming from high migration costs and lifecycle effects, as in the dynamic models of Eckert and Peters (2018) and Morten and Oliveira (2018), and consistent with the ideas of Topel (1986). Future work should focus on understanding the types of frictions that impede movement of workers to areas with higher average productivity.

This paper is structured as follows. Section 2 presents a simple framework to interpret recent evidence on real private consumption gaps and motivate our look at non-monetary amenities. Section 3 describes our data. Section 4 offers evidence on variables reflecting real consumption and wealth. Section 5 presents our results for public goods by population density. Section 6 documents how crime varies across space across our countries, and Section 7 does the same for air pollution. Section 8 explores two indirect measures of whether amenities compensate for urban wage premia, in particular self-reported well-being in urban and rural areas and net migration flows to urban areas. Section 9 returns to the historical evidence described above and asks why Africa is different today. Section 10 concludes.

2. Motivating framework and evidence

Consider a simple and general model of endogenous location choice. People derive utility from private consumption and public amenities, where private consumption depends on local wages and prices, and public amenities are location dependent. A model of this kind implies that in equilibrium, if private consumption is higher in one region than in a second region, then 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. This is the central 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. A well-established literature has characterized the conditions for a spatial equilibrium to hold; the basic insight that high prices and disamenities can offset nominal wage

¹ We acknowledge from the outset that our data lack sufficient granularity for us to identify people living within the densest urban slums, however; as we will discuss below, privacy issues prevent us from pinpointing the exact location of households in the survey data. Thus, we cannot rule out the possibility that we might find, in slums, a breakdown in the monotonicity of the relationship between living standards and population density. However, we can be confident that this effect is not so large as to offset the more general pattern. Furthermore, we show below in Section 4 that the amenity levels experienced in cities by internal migrants are also better on average than those in rural areas.

differences dates back to Rosen (1979) and Roback (1982). More recently, Glaeser and Gottlieb (2009) provide a comprehensive review of the recent literature in urban economics, much of which draws heavily on the concept of spatial equilibrium.

A model of location choice with frictionless mobility predicts that locations with higher real wages (and hence higher private consumption) will tend to have lower amenities (or, equivalently, more disamenities). To test this proposition, one would need high-quality location-level data on real wages (or consumption) and amenities. Such data are often available for rich countries. For example, Glaeser and Gottlieb (2009) show that data for Metropolitan Statistical Areas in the United States are broadly consistent with a spatial equilibrium. In today's developing countries, however, comparisons of real wages across locations are conceptually problematic. Similar problems apply to comparisons of proxies for real wages, such as real expenditure or real consumption.

The difficulty is not absence of data, *per se*. Real wages are of course not directly observed, even in rich countries. However, many countries collect data across locations on both nominal wages and prices. For instance, censuses and labor force surveys often collect data on nominal incomes or wages. Many smaller household surveys also collect detailed data on wages, self-employment income, and expenditures on a wide range of goods. These data give a clear and generally consistent picture of nominal wages and income. Unsurprisingly, within countries, urban wages are typically far higher than rural wages. (See, for example, Herrendorf and Schoellman (2018) for evidence from 13 countries of all income levels; see also Hnatkovska and Lahiri (2016) who analyze rural-urban wage gaps for China and India.) However, deflating or adjusting these nominal wage differences so that they correspond to real wage differences has proven much more difficult.

Deflating appropriately requires spatially disaggregated price data. Many countries collect price data at multiple locations – though typically at major markets, rather than in villages and rural areas. Thomas (1980) shows that one can infer regional price differences from regional expenditure data if the quality of each good is assumed constant across regions. This assumption presumably holds for a set of tradeable goods (e.g., packaged manufactured goods); but it is particularly problematic for goods and services that are locally produced and consumed. Some of these play an important role in the consumption basket; e.g., housing or transportation. Other items, such as clothing, are likely to differ significantly in quality and type between rural and urban areas, because of differing social norms and income levels. More fundamentally, in a developing country context, baskets of goods may differ substantially across locations, and many goods are simply non-comparable between rural and urban areas.² Since income levels differ between rural and urban locations, the differing baskets of goods reflect income effects as well as location-specific relative prices, as pointed out by Munshi and Rosenzweig (2016).

In spite of these challenges, several studies have combined data on nominal incomes and location-specific prices to compare real wages in rural and urban areas. In particular, Ravallion et al. (2007) construct rural and urban price indices for a large set of developing countries. They find that prices in urban areas are typically around 30 percent higher than prices in rural areas, offsetting a substantial fraction of the difference in nominal wages. Similarly, Ferré et al. (2012) construct location-specific poverty lines based on price data and consumption baskets. They use these poverty lines to assess the incidence of poverty in cities of different sizes; they conclude that poverty is more widespread and

² Furthermore, rural households frequently produce many of the goods that they consume, including food, housing, and many household goods. For example, a rural household may use gourds instead of plastic buckets, or chewing sticks instead of toothbrushes. Rural households are also more likely than urban households to home produce services, such as cooking, hair dressing, or home repair. For home-produced goods and services, it is unclear what prices we should apply to construct a price index that allows for cross-location comparisons, as discussed in Deaton and Dupriez (2011).

acute in towns and smaller cities than in larger urban settings. Munshi and Rosenzweig (2016) use location-specific prices and nominal wages to compute rural-urban wage gaps for Indian workers with less than primary education – a group relatively homogeneous in terms of skills – and find that urban wages are 27 to 35 percent higher than rural wages, in real terms, depending on the price deflator used.³

Given the methodological challenges of these comparisons, however, we have opted for a different approach in comparing living standards across space within countries. We do not work directly with data on income or prices; instead, we draw heavily on data describing real outcomes. To a significant degree, we follow Young (2013), who constructs a complex index of real consumption for a number of developing countries. We draw on the same raw data but analyze it in a simpler format, looking at each component separately. Our analysis also differs from Young's in that we move beyond a binary treatment of urban and rural to consider the whole continuum of population density. In addition, we add data on non-monetary amenities, or public consumption. We focus on amenities that we can measure clearly and consistently, and for which we have abundant data of high quality.

Our measures typically embody both elements of private consumption and 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). Similarly, a household's access to electricity depends both on the availability of electricity in the vicinity and on the household's investment in private goods, such as paying to be connected to the grid. Even for many durable goods, the household's ownership is contingent on underlying public investments or public institutional arrangements. For instance, mobile phone ownership is of little benefit to a household without a working cell signal; ownership of a motorcycle offers little benefit to a household living in a village that is inaccessible by road.

By looking closely at the data, we hope to assess the relative merits of competing explanations for differences in real wages between urban and rural areas. By focusing on observable dimensions of private consumption, we hope to avoid some of the problems of comparisons that depend on accurate measures of nominal wages, nominal prices, and consumption baskets. In our emphasis on observable attributes of amenities, we seek to capture the effects of location-specific characteristics, including those that are determined by public policy choices. And ultimately, in looking at migration rates and life satisfaction measures, we are able to look at location choices directly.

3. Data

Our core results use data from the Demographic and Health Surveys (DHS), Afrobarometer, and remotely sensed pollution data. The micro surveys are high-quality nationally representative surveys that cover large numbers of households (typically more than 5000 for each DHS survey and 1,200–2,400 for each Afrobarometer survey) in developing countries. The surveys are designed to use consistent methodologies and definitions across countries. The DHS 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 et al. (2015) and NO2 concentrations from Geddes et al. (2016). Appendix A provides an overview of the different surveys used. 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 Appendices B to D for further details.

³ Munshi and Rosenzweig (2016) also make a strong argument that the wage gaps they observe are not entirely driven by sorting and selection – an important issue that we will return to later in this paper.

Our analysis focuses on 20 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 Table B.1, covering countries with a combined population of about 770 million people.⁴

A valuable feature of the DHS data sets is that they include geocoded locations for survey clusters, subject to an important caveat. The DHS preserves the anonymity of survey respondents by displacing their locations to make them approximate. Each cluster is 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 et al., 2013). We take into account the random offset of DHS cluster locations when linking DHS GPS data with continuous raster data on population density by taking 5 km buffers around both urban and rural DHS clusters, as suggested by Perez-Heydrich et al. (2013). Appendix B 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 B.1 that when we have geocoded 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 (see Appendix C). 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). Using this algorithm, we are able to geo-locate 92 to 95 percent of village names in each round. For each location we can then extract the population density value.⁵

The use of geocoded data allows us to provide a richer picture of spatial disparities than that provided in previous studies. Previous literature has generally focused on comparing wages and amenities across locations that are defined, by administrative classifications, as either urban or rural. However, this binary classification conceals important differences. Small cities may be lumped with political or economic capitals; peri-urban areas may be treated as though they are rural and remote. Different countries define urban areas differently and inconsistently, making cross-country comparisons and syntheses problematic.

We go beyond the previous literature by making use of the more detailed data on population densities for locations where our households live. We characterize locations within countries by gridcell-level population densities, which we can then use for relative comparisons (i.e., of more and less densely populated locations within countries) or for absolute comparisons (i.e., specified levels of density). We rely on 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 data sets of administrative units (Doxsey-Whitfield et al., 2015).

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 40 or more regions per country.

Both the pollution data and the population density data are gridded data, making it straightforward to link them. The estimated PM_{2.5} 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. Appendix D contains further details on how we link the pollution data with the population density data.

We use these data to assign people and places within each country into density quartiles. We define the lowest-density quartile of the population as the 25 percent of people who live in the most sparsely populated locations. Correspondingly, when we talk about any geographic features (natural or human) characterizing the lowest quartile, these are the averages or totals across all of the gridcells inhabited by the lowest-density people. Symmetrically, the highest density quartile of the population is the 25 percent of people who live in the most densely populated gridcells. When examining outcomes across quartiles within countries, we define quartiles within countries; when we aggregate across countries, we define quartiles over the whole sample of countries. Appendix Figures F.1–F.4 show the densities for different quartiles when we use within-country and across-country quartiles for our DHS countries. How we define quartiles does not matter substantively for the relationship between density and our outcomes. We always use the DHS sample weights when computing quartiles and averages across quartiles.

4. Private consumption

We begin by comparing a set of goods that relate to private wealth and consumption. We note that urban-rural disparities in wages and total expenditure have been documented extensively in the academic literature, so the novelty of the results that we present here is primarily that we show patterns across density space and to show the evidence for individual real measures that offer a clear comparison of living conditions for urban residents in relation to more rural people. Our results here are close to those of Young (2013) but offer detail on individual components of the measures that he uses. For each of our measures, we present data in a compact fashion for our whole set of countries by comparing rates of ownership or consumption across populations within each country. We divide the populations within each country into quartiles based on the population density in the locations where people live. As noted above, an advantage of looking at density quartiles is that we avoid the problems of defining “urban” and “rural” locations. We can also look at deciles or other density bins. We show our core results as average differences across quartiles and complement this by presenting figures displaying one data point per country to explore heterogeneity across countries. Appendix G shows durable goods and housing outcomes across the entire continuum of density space.

4.1. Durable goods

Table 1 presents the average durable ownership rates by density quartile. The number of countries with statistically significant differences from the lowest-density quartile (which we denote Q1) are given below each average. (Our labels include both negative and positive differences, so that ‘0–20’ denotes 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

⁴ The 20 DHS countries are slightly different from the 20 Afrobarometer countries, as shown in Appendix A.

⁵ 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.

Table 1
Durables ownership by density quartiles.

	Population Density Quartile			
	Q1	Q2	Q3	Q4
Telephone	0.41	0.49 0–9	0.60 0–15	0.83 0–20
Television	0.11	0.15 1–4	0.29 0–16	0.59 0–20
Automobile	0.01	0.02 0–4	0.03 0–8	0.10 0–19
Motorcycle	0.13	0.13 2–4	0.13 4–4	0.14 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).

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 displaying no significant gradient. Households are likely 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.

4.2. Housing

We turn next to measures of housing consumption. We focus on seven measures of housing characteristics: the percentages of households having: (1) a constructed floor, (2) a flush toilet, (3) a finished roof, and (4) finished walls; (5) the average number of minutes each day spent collecting water, (6) the number of sleeping rooms per person over the age of 5, and (7) whether a household cooks indoor with solid types of cooking fuels. The number of rooms variable could be seen as a measure of housing quantity. Solid cooking fuels, such as charcoal and wood, when used indoors with inadequate ventilation, create very serious problems of air quality and particulate exposure for households.⁶ 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 Figures G.5-G.11 for the corresponding density 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 statistically significant from

Table 2
Housing quality by density quartile.

	Population Density Quartile			
	Q1	Q2	Q3	Q4
Constructed floor	0.33	0.38 1–5	0.52 0–12	0.82 0–19
Flush toilet	0.02	0.05 0–3	0.12 0–10	0.31 0–19
Water collection (min)	29.4	24.3 4–0	20.7 6–0	12.2 17–0
Finished roof	0.41	0.5 2–8	0.67 1–14	0.88 0–19
Finished walls	0.35	0.42 0–6	0.53 1–14	0.82 0–20
Sleeping rooms per person	0.61	0.62 0–2	0.61 1–5	0.61 4–6
Cook inside with solid fuel	0.63	0.64 0–4 [†]	0.58 5–3 [†]	0.38 15–1 [†]

Note: This table reports housing quality across quartiles across our set of 20 countries. Sleeping rooms per person includes only those aged 6 and above. 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).[†]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.

zero. Households in the densest quartiles are substantially more likely to have finished walls, a finished roof or a constructed floor.

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. The last row of Table 2 shows that the fraction cooking inside with solid fuels is five percentage points lower in the third quartile than in the least dense quartile, and 25 percentage points lower in the densest quartile, and the difference between the most and least dense quartile is statistically significant in 15 countries. In summary, indoor air pollution is worse almost everywhere in rural areas than in African cities.

We conclude that by these seven measures, housing quality is unambiguously higher in denser areas in these developing countries.⁷ 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.

4.3. Child health metrics

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 and the overall disease environment. 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 five main objective measures of child health: stunting, wasting, malaria status, 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

⁶ 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 (World Health Organization, 2014). Indoor air pollution is also estimated to contribute significantly to outdoor air-pollution-related deaths by exacerbating health conditions linked to pollution.

⁷ 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.




Fig. 1. Child health *Note:* This figure plots the percent of children in the least dense and most dense areas that are anemic, consuming below an acceptable minimum diet, wasted, stunted, and testing positive for Malaria. Each dot represents one country. See the text for the definition of each variable.

low weight for height. The minimum acceptable diet is defined by minimum meal frequency and dietary diversity (World Health Organization, 2015). Malaria is determined from blood samples. For all metrics we calculate the fraction of children that have the condition in question and then aggregate by population density; see Appendix B for more detail on how these variables are defined.

Fig. 1 shows the fraction of children with stunting, wasting, anemia, consumption below the acceptable minimum diet and malaria rates. 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), the percent wasted (lighter diamonds), and the percent testing positive for malaria (black squares). The lower panel covers the percent anemic (darker circles) and the percent consuming below an acceptable minimum diet (lighter diamonds).

Two main features stand out in Fig. 1. First, for stunting, wasting, anemia, and consumption below the acceptable minimum, 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. Table 3 shows that stunting rates are significantly lower in urban areas in 11 out of 20 countries, while around quarter to one 12:45 half of the countries have significant differences in the other metrics.

It is hard for us to conclude whether these findings are largely driven by 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 countries, child health is substantially better in more densely populated areas. One measure for infectious diseases is the incidence of malaria using data from test samples taken as part of the Demographic and Health Surveys, Malaria Indicator Surveys and AIDS Indicator Surveys. The figure shows that malaria incidence is lower in more densely populated areas. This is in line with the findings of Tatem et al. (2013), who document a negative relationship between malaria and urbanization on a global scale.

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

Table 3
Child health characteristics by density .

	Population Density Quartile			
	Q1	Q2	Q3	Q4
Anemic [†]	0.69	0.68	0.64	0.6
		0-0	3-0	7-0
Stunted (low height for age)	0.4	0.4	0.38	0.29
		0-0	3-1	11-0
Wasted (low weight for height)	0.1	0.1	0.09	0.09
		1-0	3-1	6-1
Below minimum diet	0.93	0.91	0.9	0.87
		1-0	2-0	5-0
Malaria [†]	0.35	0.32	0.29	0.13
		1-0	4-0	9-0

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). [†] Anemia is available for 16 countries and malaria for 13.

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 historical pattern of cities being associated with worse health environments.⁸

4.4. Heterogeneity, sorting and selection

One possible explanation for the urban-rural disparities that we observe in the data on private consumption is that there may be sorting and selection into locations. For example, if high-ability people move disproportionately to urban areas (or more densely populated rural areas), then the patterns that we observe could arise simply as a consequence of selection and sorting. This issue has been addressed extensively elsewhere, and we can only do limited analysis with our data. However, in this section we consider how our results vary by education level, and how rural-urban migrant households differ from other urban households. In this analysis we also include electricity as one measure of public amenities that is relevant for sorting and selection.

4.4.1. 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 (2013), Bryan and Morten (2019) and others, and by many in the urban economics literature, including Glaeser and Maré (2001) and Combes et al. (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.

⁸ Jedwab and Vollrath (2019) 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.

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 E.2 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.

Fig. 2, 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 quality is positive for the two main education groups.

4.4.2. 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 destination, 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 4 reports measures of durables ownership, housing quality, and indoor air quality for rural-urban migrant households in the densest quartile. We have also added to the table two categories of public goods availability: tap water and electricity connection. Although these do not strictly measure private consumption, they are relevant for addressing the question of heterogeneity in housing quality for migrant households.

The first two data columns 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




Fig. 2. Density gradients by education of household head *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 who have less than complete primary education; and Panel (b) shows gradients for households with household heads who have complete primary education or more. The y-axis indicates the size of the coefficient.

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 to other urban households by most of our metrics, or marginally worse off by some criteria. But they are still faring better than the average rural households of Q1. In 10 of 12 countries, the migrants have statistically significantly higher rates of telephone ownership, 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 better living conditions than households that remain in rural areas.

4.5. Summary of private consumption results

To conclude this section, we note simply that our data on private consumption and wealth measures are consistent across countries in showing that essentially all measures appear to show households in urban areas as better off than those in rural areas. More precisely, we see these measures improving with population density. These measures are

Table 4
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 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).

all based on real and observable measures, rather than deflated nominal wages. Although the data are surely imperfect, they are not subject to the same measurement concerns that pose challenges to standard measures of real wages, derived from nominal wages or household income and expenditure data, with tenuous adjustments for location-specific prices. At the same time, the spatial disparities that we observe are entirely consistent with those calculated using other methods. We cannot reject fully the possibility that selection and sorting account for some fraction of the disparities that we observe; indeed, it would be surprising if these forces did not play a role in explaining the data. However, the spatial disparities that we observe seem to hold when we compare similarly educated populations. We can also reject the possibility that migrants to cities experience conditions far worse than the averages.

5. Public goods and amenities

We turn next to the prevalence of publicly provided amenities across space, such as the presence of a sewerage system, schools and health facilities. We begin with a range of specific publicly provided goods and services, such as physical infrastructure. These may enter utility directly, but more typically, they require complementary private goods. For instance, the provision of an electricity connection does not directly offer utility to people; utility comes from complementarities with private purchases of electricity-using goods (lights, televisions, etc.). The same is true of piped water provision; the sheer existence of sewerage in a community does not necessarily make a household better off, but when households are able to pay for indoor plumbing, the utility gains may be high. An additional concern is that the presence of a publicly provided good or amenity does not necessarily speak to the quality of the good. Access to a road confers benefits that vary strongly with the condition of the road; access to a health clinic is of little value if the staff are not present.

5.1. Public goods provision by population density in Africa

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. As in the previous sec-

Table 5
Public goods by density quartile.

	Population Density Quartile			
	Q1	Q2	Q3	Q4
Electricity grid	0.39	0.42	0.48	0.72
		1–0	0–4	0–10
Piped water system	0.36	0.35	0.42	0.67
		0–0	0–2	0–11
Sewage system	0.14	0.13	0.18	0.37
		0–0	0–0	0–7
School	0.91	0.90	0.90	0.90
		0–0	0–0	0–0
Health Clinic	0.59	0.58	0.62	0.73
		1–0	0–0	2–4
Police Station	0.29	0.30	0.33	0.47
		0–0	0–0	1–4
Market Stalls	0.62	0.64	0.64	0.76
		0–1	0–0	0–4
Paved Road	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. Paved road is equal to one if the road taken on the way to the interview was at the start paved/tarred/concrete. 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). In this table we use within-country quartiles rather than across-country quartiles to ensure that all countries have enumeration areas across quartiles.

tion, we present results for different quartiles of the population density distribution.

Table 5 shows a number of public goods across density quartiles. Overall, none of these measures is worse in denser areas, on average. Denser areas are more likely to have electricity grids, piped water systems, and sewerage systems. 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.

The relatively high provision of public goods and infrastructure in urban areas may also reflect a legacy dating to colonial times. As argued in Michalopoulos and Papaioannou (2013), many colonial powers did not seek to exercise direct control (or were unable to do so) outside of capital cities, and local and provincial leaders may have lacked the resources or capability to build formal infrastructure. Certainly the pattern of low public investment in rural areas of sub-Saharan Africa has long been recognized, at least anecdotally. As far back as the late 1970s, Yap (1977) compiled fragmentary evidence from a handful of countries to suggest that, relative to rural people, "... the average urban dweller seems to have more access to piped water, sewage connections, and electricity; better medical care; and better educational opportunities." At the time, a literature in development economics and political science attributed these advantages to "urban bias" (see, for instance, Lipton (1977) and Bates (1981)). In this view, the relative political power of urban people – often more educated, more active, and closer to the seat of government than rural people – implied that governments were more likely to devote resources to supplying urban locations with amenities, all else equal. Along with the lower cost per household of providing many services in urban areas, urban bias provided a persuasive explanation for the spatial gaps in amenities.

5.2. Quality of public goods and services

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 covered. Yet there does exist some evidence on the quality of public goods at points across space, and this evidence points in the direction of quality that is at least as high in urban areas, if not higher.

For instance, in an Afrobarometer survey from Round 7 (2016/18), respondents were asked whether they had an electricity connection to their dwelling – and if so, how frequently electricity was actually available from this connection. The survey covered more than thirty countries in Africa and over 45,000 respondents. Of rural respondents (who made up 55 percent of the sample), 62 percent lacked an electrical connection, compared to 17 percent of urban respondents, reinforcing the disparity in access. But quality differences exacerbate the disparity: even among those who reported having an electrical connection, 31 percent of rural respondents reported that electricity was available at most half the time, compared to 21 percent for urban respondents. The same survey reported that 54 percent of rural respondents, compared with 46 percent of urban respondents, said that "the current government is... maintaining roads and bridges" either "very badly" or "fairly badly," suggesting that differential provision of roads is compounded by differential quality of maintenance.

Perhaps the clearest data on quality differences pertain to schooling. In a set of 14 African countries, Lee et al. (2005) examine student test scores by region and find robust evidence of lower test scores in rural areas. This finding is certainly consistent with lower quality schools in rural areas, though it could also reflect lower quality parental inputs. Chaudhury et al. (2006) provide even more direct evidence about schooling quality in the developing world in their survey of teacher absence. It is hard to think of many ways to reduce schooling quality more directly than to remove the teacher from the classroom, and this is of course just what happens when the teacher is absent. When looking by region, Chaudhury et al. (2006) find that teacher absence is more prevalent in rural areas in Indonesia, India and Peru. In Ecuador, Bangladesh

and Uganda they find a statistically insignificant difference in teacher absence between rural and urban areas.⁹

Thus, in terms of the provision of public goods and services, urban areas seem systematically to do better both in terms of the availability of infrastructure and public services, and very likely also in terms of the quality of those public goods, conditional on their provision. Certainly there is nothing here to suggest disamenities of urban life that would offset the advantages in nominal or real wages.

6. Crime

We turn from public goods provision to one of the potential disamenities that is most frequently invoked as a characteristic of urban areas: crime. Data on crime are problematic in all countries, and perhaps these data issues are particularly germane in developing countries. Official administrative records are either not stored centrally or not available to researchers. Given that, 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 Appendix Table C.1). One potential 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 administered 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. Survey data may also elicit responses from people who have not bothered to report crimes to police; Afrobarometer data (Round 7, 2016/2018) show that 88 percent of rural respondents and 85 percent of urban respondents say that they have not requested any help from police in the past twelve months, even though around 30 percent of respondents say that they have been victims of theft at least once during that period.¹⁰

6.1. Patterns of crime by population density


We consider four main metrics: property crime, violent crime, feeling safe in one's neighborhood, and fear of crime in one's home. Crime, perhaps more than our other measures, might be sensitive to phrasing. For clarity, we therefore list the exact questions. 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 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," "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

⁹ Chaudhury et al. (2006) have data on health-care worker absence in rural and urban areas for just one country: Peru. There, health workers in urban areas 4.4 percentage points less likely to be absent than rural counterparts.

¹⁰ Possibly this reflects expectations of police follow-through on reported crime. More than one-fourth of the total sample thought that it was "not at all likely" or "not very likely" that police would "take action."

**Fig. 3.** Crime by density quartile

Note: This figure plots the percent of respondents in the least dense and most dense areas experiencing property crime, experiencing violent crime, feeling safe in their neighborhood, and reporting fear of crime. Each dot represents one country. See the text for the definition of each variable.

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.

Fig. 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 (and particularly in remote 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 6 presents the average crime rates by density quartile across all the countries. Around 29 percent of households in the least

Table 6

Crime and fear of crime by density quartile.

	Population Density Quartile			
	Q1	Q2	Q3	Q4
Property crime	0.29	0.31	0.31	0.33
Violent crime	0.1	0–3	0–1	0–4
		0–1	0–0	0–4
Fear of crime	0.32	0–1	0–2	0–4
		0–2	0–1	0–3
Feel unsafe	0.37	0–2	0–1	0–3
		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).

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

6.2. Willingness to pay to avoid crime

Still, our evidence so far suggests that crime is the leading contender for the amenity that 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 area with less crime, proxied by willingness to pay. What is the value of having a 33-percent chance of theft in the densest areas relative to just 29 percent in the least dense areas, 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 ones?

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 et al. (2004) 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). Cook and Ludwig (2000) arrive at a similar estimate in percentage terms of willingness to avoid gun violence.

Still, a concern with these calculations is that they are all based on willingness-to-pay estimates from advanced economies. To shed some light on how much of a difference this may make, suppose that the crime differences were indeed enough to compensate rural households for the substantially lower average wages they earn. What would the implied willingness to pay to avoid crime have to be in these African countries? Would it be sensible?

Take violent crime, for which individuals are relatively most willing to pay to avoid (compared to theft and property crime). Recall that across our countries, violent crime affects 10 percent of those in the most rural areas and 12 percent in the densest areas. For this to explain a wage gap of roughly a factor two between rural and urban areas, it would have to be the case that African households were willing to give up 50 percent of their wages to avoid each additional 1 percentage point increase in violent crime. This is two orders of magnitude larger than the estimates we found in the U.S. literature, which seems quite im-

plausible. Taken together, the evidence in this section suggests that the modest differences in crime rates with density are unlikely to be large enough to offset the much higher average incomes in urban areas.

7. 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 location choice in the United States, and exposure to pollutants has been shown to significantly affect health, human capital and productivity (Currie et al., 2009; Currie and Walker, 2011; Graff Zivin and Neidell, 2012; Adhvaryu et al., 2020).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 ($\mu\text{g}/\text{m}^3$). For frame of reference, the World Health Organization recommends mean annual exposures of 10 $\mu\text{g}/\text{m}^3$ or less for PM2.5 and 40 $\mu\text{g}/\text{m}^3$ 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 et al., 2016; World Health Organization, 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.

7.1. Air pollution by population density

Fig. 4 shows the distributions of PM2.5 and population density across space in Nigeria.

The left graph shows the distribution of population density; the right graph shows the PM2.5 distribution. Darker 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 Harcourt and Enugu.

Moving to the pollution measure, several observations are worth highlighting. First, at least visually, population density does not appear to be strongly correlated with PM2.5 concentrations. PM2.5 levels appear to be driven mainly by dust from the Sahara when inspecting the graph. Removing sea salt and dust produces quite a different distribution, with higher levels in the center and over some cities, but still with little obvious correlation with population density. It is instructive to look separately at these two indicators for pollution as shown in Figure D.1.¹³ The pairwise correlation between PM2.5 and NO2 is -0.1158 with a p-value of 0.000. Across our whole set of African countries, the correlation of these two measures ranges from 0.64 (Cameroon) to -0.49 (Senegal).

Fig. 5 plots (as 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 in both sparsely populated areas and densely populated areas. Most of these countries (e.g., Niger, Senegal, Mali and Nigeria) border the Sahel and are thus exposed to dust blown off the Sahara. To account for this, we plot PM2.5 concentrations with dust and sea salt removed (as triangles). There is no evidence suggesting that anthropogenic sources of PM2.5 are more or less harmful for health than

¹¹ Studies by Fafchamps and Moser (2003) and Demombynes and Özler (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. We have also examined the fractions of LSMS households having experienced a crime in the last twelve months, in five countries for which we could spatially link respondents with population density data: Ethiopia (2013), Malawi (2004–05), 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, Figure E.1 shows that a constant rate of crime seems as though it would just about fit within the confidence intervals.

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).

¹³ This is in line with what Geddes et al. (2016) find when they inspect population weighted average PM2.5 and NO2 levels and trends.




Fig. 4. Outdoor air pollution (PM2.5) in Nigeria *Notes:* The left graph shows the distribution of population density and the right graph shows PM2.5. Darker colors denote higher values, and the bins are formed by dividing the data into deciles.




Fig. 5. PM2.5 concentrations by population density *Note:* This figure plots average concentrations of fine particulate matter (PM2.5) in the atmosphere in the least dense and most dense areas. Each dot represents one country.

are natural sources, but it is possible that individuals perceive anthropogenic sources as more hazardous to their health. When dust and sea salt are taken out of the calculations, the data show far lower levels of PM2.5 across all locations, but it is still the case that there are very modest differences by population density.

Turning to NO₂, Fig. 6 shows that 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 in Table 7, we find that NO₂ concentrations do not vary in statistically significant ways across density space. If anything, PM2.5 levels are lower in more densely populated areas.

We conclude that outdoor air pollution concentrations in Africa are at best loosely linked to population density.

We note that this surprising result is highly specific to the countries in our sample. In other parts of the world, the same pollution measures display quite strong patterns of (positive) correlation with population density. What is different in Africa is that there is relatively little manufacturing or “dirty” generation of electricity. To see this, we consider PM2.5 concentrations in other countries. We find that in China, India and the United States, pollution gradients with density are strongly positive. In China, PM2.5 levels for the top population density decile amount to 68 $\mu\text{g}/\text{m}^3$, almost seven times the WHO recommended threshold; the

Table 7
Pollution differences by density quartile.

	Population Density Quartile			
	Q1	Q2	Q3	Q4
Nitrogen Dioxide	0.26	0.26	0.25	0.26
PM2.5	19.45	6-7	7-9	7-9
		20.24	18.55	18.15
PM2.5 (dust and sea salt removed)	5.84	8-5	10-5	11-7
		5.81	5.79	5.84
		8-8	9-9	9-8

Note: This table reports pollution exposure quartiles across our set of 22 countries listed in Table A.1. 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 use within-country quartiles rather than across-country quartiles to ensure that all countries have observations across quartiles.

lowest population density decile has a level of 16 $\mu\text{g}/\text{m}^3$. In India, the top decile has a level of 48 $\mu\text{g}/\text{m}^3$, more than four times the WHO recommended threshold, compared to 24 $\mu\text{g}/\text{m}^3$ in the lowest decile. These positive gradients are very much what we might expect of a world in




Fig. 6. Nitrogen dioxide concentrations by population density quartile. *Note:* This figure plots average concentrations of nitrogen dioxide in the atmosphere in the least dense and most dense areas. Each dot represents one country.

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 et al., 2016); 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 all dimensions of pollution exposure. At a 10km resolution, our measures are spatially rather coarse. Local effects, such as proximity to roads, may matter significantly.¹⁴ Moreover, satellite-derived measures reflect the column of pollution as observed from space, rather than the concentration experienced on the ground. Nevertheless, as the data from India, China, and the United States illustrate, our metrics seem reasonable and appropriate. What emerges from this analysis is that African cities are neither large enough nor industrialized enough to create large clouds of pollution, and background non-anthropogenic pollution is high. This combination produces different pollution gradients from those observed in more industrialized parts of the world.

We note 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 health measures – which are on average better in urban areas – are quite sensitive to certain types of water pollution and sanitation problems. One important and related – but separate – issue is indoor air pollution. This issue is closely associated with the types of cooking materials that households rely on. However, the nature of this problem is that it is effectively driven by private consumption choices. For that reason, we covered this issue in Section 4.

8. Indirect measures of amenities

The descriptive statistics of outcomes across population densities in the previous sections of this paper suggest that there is no easily observable measure of consumption or amenity that is decreasing with population density. We turn now to two indirect measures of urban amenities.

¹⁴ For example, Kinney et al. (2011) find average PM_{2.5} concentrations at four traffic sites in Nairobi, between 7.30am and 6.30pm, that fall between 58.1 and 98.1 $\mu\text{g}/\text{m}^3$; by contrast the maximum multi-annual average PM_{2.5} concentration for Kenya in our sample is 13.9 $\mu\text{g}/\text{m}^3$ – and this pixel is at Lake Turkana, the world's largest desert lake.

The first is subjective well-being; the second is migration.¹⁵ If urban disamenities literally balanced the higher private consumption levels, so that we were in a (static) spatial equilibrium, we should observe no systematic differences in subjective well-being for people in rural areas and urban areas, and we should also observe no systematic patterns of migration between rural and urban areas. We might still expect to see some movement of individuals across space, to account for sorting and selection based on individual tastes, preferences, and talents; but we would expect these movements more or less to balance. In this section, we review the micro data to ask whether this is the case.

8.1. Subjective well-being

Measures of subjective well-being are inevitably difficult to interpret, and comparisons across countries may be particularly problematic because of differences in social norms. However, differences between urban and rural areas within a country seem to offer a plausibly valid comparison. Our analysis draws on data from the World Values Survey, and follows Glaeser (2012) in comparing subjective well-being in rural and urban areas. We concentrate on two variables: satisfaction, reported on a scale of 1 (least satisfied) to 10 (most satisfied), and happiness, which we measure as the fraction of individuals reporting that they are “quite happy” or “very happy.” In both cases we compute the average for people aged 15 and over by urban and rural areas.

Table 8 reports the average life satisfaction rates and percent of people that report being happy in the seven African countries for which we have data on both urban and rural households. In this data set, we lack geo-references for individual households, so we are obliged to use the urban-rural classification of the survey. The results are striking, however. In all seven countries, urban households report higher levels of average life satisfaction than rural households; in five of the countries, the differences are statistically significant. Urban households are happier in six of the seven countries, with the exception being Rwanda, where both rural and urban households report extremely high levels of happiness. Four of the seven countries have statistically significant differences between rural and urban households; in all four of these, the difference is that urban households report being happier. Overall, the evidence for this subset of our countries is consistent with the finding of Glaeser (2012) that residents of urban areas in the developing world

¹⁵ We also looked at secondary amenities such as trust and anxiety; please see Appendix H for the detailed results.

Table 8
Subjective well being.

	Average Life Satisfaction			Percent “Happy”		
	Urban	Rural	Difference	Urban	Rural	Difference
Burkina Faso (2007)	5.7	5.5	0.3**	81.0	78.4	2.6
Ethiopia (2007)	5.0	4.4	0.6*	63.6	58.3	5.2
Ghana (2007)	6.5	5.8	0.7***	82.1	75.5	6.5***
Nigeria (2011)	6.3	6.1	0.2**	89.2	74.7	14.4***
Rwanda (2012)	6.5	6.4	0.1	90.2	91.2	-1.0
Uganda (2001)	5.7	5.6	0.1	80.1	78.2	1.8*
Zimbabwe (2011)	6.3	5.6	0.6***	43.4	37.9	2.8*

Note: This table reports average life satisfaction, and the fraction of adults reporting that they are “very happy” or “somewhat happy” (rather than “not very happy” or “not happy at all”), by urban and rural areas. The data come from the World Values Surveys. ***, **, * mean statistically significant at the one-, five- and ten-percent levels.

Table 9
Net migration.

	Rural-to-Urban Percent of Adults	Urban-to-Rural	Difference
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***
Sierra Leone (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.

are more likely to report being satisfied with their lives, and happy, than those living in rural areas.

8.2. 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.¹⁶

Table 9 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 the countries but one have significantly greater rural-to-urban migration than urban-to-rural. (Liberia is the exception; rural-to-urban migration is still higher, but not significantly so.)¹⁷

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, at least in a simple static sense. 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 distribution of population. For these reasons, we find it useful to emphasize that the net flow of people in these economies is clearly single-directional.

¹⁶ Recent literature, e.g., Ingelaere et al. (2018), has drawn on quantitative and qualitative longitudinal evidence to emphasize the importance of migration from small towns to larger towns and secondary cities. Many individuals begin the migration process by moving from rural areas to nearby towns, and then gradually over a period of years make successive moves to larger cities. Our data are too crude to allow us to trace these flows. For those individuals who reported that they previously lived in a town, we do not have confidence that we can assign them to a high- or low-density area. Thus, we omit them from our migration calculations, although we keep them in the sample to compute the appropriate population-weighted cut-off for density quartiles.

¹⁷ Our finding is not in contrast to Young (2013), 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.”

9. Why are African cities different?

African patterns of rural-urban migration are not substantially different from the patterns that characterized the historical cities of Europe and the Americas. Across time and space, where economies are growing, rural people have tended to leave farms and towns to move towards cities. What is different, perhaps, is that urbanization processes in earlier times were held in check by the clear negative features of urban life: disease, death, and horrifying living conditions. Those conditions have not disappeared in today's African cities – where tens of millions of people live today in dismal slums – but the reality is that on many dimensions, Africa's cities benefit from a set of historical and social circumstances that make them highly attractive relative to historical cities in Europe and the Americas.

Why are African cities different? Why do they appear to lack the massive disamenities and disadvantages that characterized urban cities in other times and places? Why do the trade-offs of urbanization seem less acute? And if the trade-offs are absent, what has prevented urbanization in Africa – rapid though it has been – from simply exploding? There are no obvious answers here, but we can offer some speculation.

One important factor that has benefited urban life in Africa is that today's sub-Saharan cities are closely connected to a globalized world. For all the problems associated with development assistance and modernization efforts, today's African residents are beneficiaries of multiple advances in the global knowledge frontier with respect to sanitation, health, and medical care. Improvements in water supply and sewage systems have lessened the burden of diseases like typhoid and cholera, which were widespread killers in earlier eras. Vaccination campaigns, pharmaceuticals, and antibiotics have reduced the burden of infectious disease. Simultaneously, increases in global agricultural productivity, combined with improvements in food safety and quality, have made food relatively affordable even for poor urban dwellers in developing countries. In this sense, Africa's cities have benefited from having access to frontier technologies that were developed elsewhere – precisely in response to the problems of urbanization in today's rich countries. These technologies have provided benefits in Africa's cities by offsetting some of the most serious disamenities of urban life.

To some degree, perhaps, urbanization in sub-Saharan Africa has also benefited from being a voluntary process in many countries. This contrasts with the period of rapid urbanization in European history, which coincided with the displacement of many rural people who were forced off agricultural land through processes such as enclosure and the consolidation of land holdings (Lambert, 1963; Allen, 1982). Similar processes have taken place in parts of Latin America, in conjunction with the development of large-scale agriculture and *latifundia* (Conning, 2002). Push migration of this kind tends to create large classes of migrants living in desperate conditions in cities.

Although any generalization is necessarily crude, the experience in most African countries has been somewhat different. In much of west and central Africa, agricultural land has remained relatively abundant, and land distribution has supported smallholder farmers, so that rural livelihoods have remained viable (though low in productivity) (De Janvry et al., 2001). Other parts of the continent have had different experiences. In southern Africa and parts of eastern Africa, European settlers claimed vast swathes of productive farmland and used state-sanctioned violence to force the previous inhabitants into homelands, “native reserves”, or communal lands (Werner, 1993; Boisen, 2017). But even in these countries, rather than forcing rural people into cities, colonial governments exercised tight controls on rural-urban migration, through the imposition of movement restrictions and residence permits (such as the apartheid system of South Africa or the Rhodesian “pass laws”). Cities were kept artificially small, although in the second half of the twentieth century, informal townships sprang up on the periphery of South Africa's cities. Various coercive measures were used to supply labor for the mines of southern Africa, but mining camps and compounds were typically located outside major cities, so these processes did not contribute much

to urbanization. And although living conditions in mining communities were poor (especially in the notorious “closed compounds”), governments and mine owners arguably had some incentive to prevent mass death and disease.

Taken together, these factors may help explain why urban living standards in today's sub-Saharan countries are generally high relative to rural living standards, whereas the opposite was true in the historical cities of Europe and the Americas.

What about the other puzzle – why African cities have not simply exploded in population? Arguably, the premise of this question is false. People are indeed moving to cities – and rapidly. Urban growth rates in sub-Saharan Africa are the highest in the world, and United Nations projections suggest that the share of people in sub-Saharan Africa living in urban areas will rise from 41.4 percent in 2020 to 58.1 percent in 2050 (United Nations, 2018).

10. Conclusion

In this paper, we examine the spatial distribution of amenities in 20 African countries using new disaggregated evidence. We find that *both* real measures of private consumption *and* real measures of amenities seem higher in urban areas. More generally, these measures improve with population density, more or less monotonically. It is of course possible that we may be missing some intangible characteristics of urban areas that make them undesirable places to live. Our search for disamenities cannot, of course, be exhaustive. For instance, some people might simply prefer to live in rural areas, viewing urban life itself as a disamenity. We cannot rule out such heterogeneity in preferences. Rural areas may also promote greater social connectedness and better networks for friendship and community engagement. We recognize that these may be important contributors to well-being. Nevertheless, our analysis has looked at three of the most likely candidate disamenities, and also at measures of subjective well-being. If indeed some consequential urban disamenity exists, we have not succeeded in identifying it.

But even in a world where utility levels in cities are consistently and substantially higher than in rural areas, mobility frictions of various kinds may limit the speed at which people move towards places with higher living standards. This is the view of Chauvin et al. (2017), who suggest that in low-income countries, widespread frictions prevent utility from being fully equalized across locations. Frictions might arise from high migration costs, as in the dynamic models of Eckert and Peters (2018), Morten and Oliveira (2018), and Lagakos et al. (2020), and consistent with the ideas of Topel (1986). Other sources of frictions might include the lack of information in rural areas about employment opportunities and living conditions in urban areas, as in the work of Baseler (2019). Migration frictions might also arise through rigidities in land markets (de Janvry et al., 2015), or in insurance markets (Morten, 2019), or simply through the loss of social ties for migrants.

The simplest explanation of all might be that there are frictions associated with the life-cycle: rural-urban migrants tend to be young, and they often move before they are married. These life-cycle forces create a natural inertia that slows the process of urbanization, as pointed out in the earlier literature by Yap (1977) and Stark and Lucas (1988).

Understanding these frictions – and the broader spatial disparities in living standards – is not just a matter for academic disputation; it is ultimately an important issue for policy. Different interpretations of spatial disparities encourage us to think about different policy responses. For instance, heterogeneity in preferences would not call for any policy intervention. By contrast, in a world where urban disamenities are substantial, policies might focus on increasing well-being in urban areas – for example, by improving roads or transportation networks. 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 make utility across locations more equal; if our

analysis is correct, such investments would tend to *exacerbate* spatial disparities in utility.

An alternative approach might focus on reducing frictions to movement – perhaps by addressing the social disruptions and fixed costs associated with mobility. If a model with frictions can better match the data, policies of this kind might serve to equalize utility across locations and to allocate labor to its most productive uses. Although our analysis cannot conclusively identify the reasons for persistent spatial disparities in income, mobility frictions could certainly shape urbanization processes in our set of African countries. Rural-urban migration is taking place, but it has seemingly not equalized income or utility across locations.

To conclude, urbanization in today's African countries does not appear to be associated with the negatives that characterized the historical cities of Europe and North America. To be sure, governments and international development organizations face challenges in assuring the adequate provision of water and sanitation and other services in crowded urban conglomerations. The global COVID-19 pandemic brings home the potential vulnerability of urban life. But African cities benefit from being visible to national governments and the international community. Their problems receive attention from a global technocracy that wields impressive scientific understanding. By nearly all measures, it is in rural areas of the developing world – where poverty is less visible and the poor are less politically powerful – that problems are most acute. Rural lives and livelihoods in sub-Saharan Africa are challenging even under the best circumstances, and the growing threats posed by climate change will exacerbate the vulnerability of farmers and rural communities. Pressure on land and water resources will increasingly undermine the ability of the rural poor to prosper. Although the challenges of urban centers and urban people are real, the data remind us that the challenges facing rural people may be even more pressing and urgent.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at [10.1016/j.jue.2020.103301](https://doi.org/10.1016/j.jue.2020.103301)

References

- Adhvaryu, A., Kala, N., Nyshadham, A., 2020. The light and the heat: productivity co-benefits of energy-saving technology. *Review of Economics and Statistics* 102 (4), 779–792.
- Allen, R.C., 1982. The efficiency and distributional consequences of eighteenth century enclosures. *Economic Journal* 92 (368), 937–953.
- Banzhaf, H.S., Walsh, R.P., 2008. Do people vote with their feet? an empirical test of tiebout's mechanism. *American Economic Review* 98 (3), 843–863.
- Baseler, T., 2019. Hidden income and the perceived returns to migration: Experimental evidence from Kenya. Unpublished Working Paper, University of Rochester.
- Bates, R., 1981. *Markets and States in Tropical Africa: The Political Basis of Agricultural Policies*. University of California Press.
- Bishop, K.C., Murphy, A.D., 2011. Estimating the willingness to pay to avoid violent crime: a dynamic approach. *American Economic Review Papers and Proceedings* 101 (3), 625–629.
- Boisen, C., 2017. From land dispossession to land restitution: European land rights in south africa. *Settler Colonial Studies* 7 (3), 321–339.
- Boo, K., 2014. *Behind the beautiful forevers: Life, death, and hope in a mumbai undercity*. Penguin Random House.
- Bryan, G., Morten, M., 2019. The aggregate productivity effects of internal migration: evidence from Indonesia. *Journal of Political Economy* 127 (5), 2229–2268.
- Center for International Earth Science Information Network, 2015. *Gridded Population of the World, Version 4 (GPWv4), Preliminary Release 2* (2010). Columbia University. Accessed 23 September 2015; available: <http://www.ciesin.columbia.edu/data/gpw-v4>
- Chaudhury, N., Hammer, J., Kremer, M., Muralidharan, K., Rogers, F.H., 2006. Missing in action: teacher and health worker absence in developing countries. *Journal of Economic Perspectives* 20 (1), 91–116.
- Chauvin, J.P., Glaeser, E., Ma, Y., Tobio, K., 2017. What is different about urbanization in rich and poor countries? cities in Brazil, China, India and the United States. *J Urban Econ* 98, 17–49.
- Cohen, M.A., Rust, R.T., Steen, S., Tidd, S.T., 2004. Willingness-to-pay for crime control programs. *Criminology* 42 (1), 89–110.
- Combes, P.-P., Duranton, G., Gobillon, L., 2008. Spatial wage disparities: sorting matters!. *J Urban Econ* 63 (2), 723–742.
- Conning, J., 2002. *Latifundia Economics*. Technical Report. Hunter College Department of Economics.
- Cook, P.J., Ludwig, J., 2000. *Gun Violence: The Real Costs*. Oxford University Press, Oxford, United Kingdom.
- Costa, D.L., Kahn, M.E., 2006. Public health and mortality: what can we learn from the past? In: Auerbach, A.J., Card, D., Quigley, J.M. (Eds.), *Public Policy and the Income Distribution*. Russell Sage Foundation.
- 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., Hanushek, E.A., Kahn, E.M., Neidell, M., Rivkin, S.G., 2009. Does pollution increase school absences? *Review of Economics and Statistics* 91 (4), 682–694.
- Currie, J., Walker, R., 2011. Traffic congestion and infant health: evidence from e-zpass. *American Economic Journal: Applied Economics* 3 (1), 65–90.
- Cutler, D., Miller, G., 2005. The role of public health improvements in health advances: the twentieth-century United States. *Demography* 42 (1), 1–22.
- Davis, M., 2006. *Planet of Slums*. Verso, London.
- De Janvry, A., Gordillo, G., Sadoulet, E., Platteau, J.-P., 2001. *Access to land, rural poverty, and public action*. Oxford University Press.
- Deaton, A., Dupriez, O., 2011. Purchasing power parity exchange rates for the global poor. *American Economic Journal: Applied Economics* 3 (2), 137–166.
- Demombynes, G., Özler, B., 2002. Crime and local inequality in South Africa. World Bank Policy Research Working Paper 2925.
- van Donkelaar, A., Martin, R.V., Brauer, M., Boys, B.L., 2015. Use of satellite observations for long-term exposure assessment of global concentrations of fine particulate matter. *Environ. Health Perspect.* 123 (2), 135–143.
- Doxsey-Whitfield, E., MacManus, K., Adamo, S.B., Pistoles, L., Squires, J., Borkovska, O., Baptista, S.R., 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., Peters, M., 2018. Spatial structural change. Unpublished Working Paper, Yale University.
- Fafchamps, M., Moser, C., 2003. Crime, isolation and law enforcement. *J Afr Econ* 12 (4), 625–671.
- Ferré, C., Ferreira, F.H.G., Lanjouw, P., 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., Martin, R.V., Boys, B.L., van Donkelaar, A., 2016. Long-term trends worldwide in ambient no₂ concentrations inferred from satellite observations. *Environ. Health Perspect.* 124 (3), 281–289.
- Glaeser, E.L., 2012. *Triumph of the City*. Penguin Books, New York.
- Glaeser, E.L., Gottlieb, J.D., 2009. The wealth of cities: agglomeration economies and spatial equilibrium in the United States. *J Econ Lit* 47 (4), 983–1028.
- Glaeser, E.L., Maré, D.C., 2001. Cities and skills. *J Labor Econ* 19 (2), 316–342.
- Glaeser, E.L., Sacerdote, B., 1999. Why is there more crime in cities? *Journal of Political Economy* 107 (S6), S225–S258.
- Gollin, D., Jedwab, R., Vollrath, D., 2016. Urbanization with and without industrialization. *Journal of Economic Growth* 21, 35–70.
- Gollin, D., Lagakos, D., Waugh, M.E., 2014. The agricultural productivity gap. *Quarterly Journal of Economics* 129 (2), 939–993.
- Graff Zivin, J., Neidell, M., 2012. The impact of pollution on worker productivity. *American Economic Review* 102 (7), 3652–3673.
- Graff Zivin, J., Neidell, M., 2013. Environment, health, and human capital. *J Econ Lit* 51 (3), 689–730.
- Herrendorf, B., Schoellman, T., 2018. Wages, human capital, and barriers to structural transformation. *American Economic Journal: Macroeconomics* 10 (2), 1–23.
- Hnatkovska, V., Lahiri, A., 2016. Urbanization, structural transformation and rural-urban disparities in China and India. Unpublished Working Paper, University of British Columbia.
- Ingelaere, B., Christiaensen, L., De Weert, J., Kanbur, R., 2018. Why secondary towns can be important for poverty reduction: A Migrants perspective. World Bank.
- de Janvry, A., Gonzalez-Navarro, M., Emerick, K., Sadoulet, E., 2015. De-linking land rights from land use: certification and migration in Mexico. *American Economic Review* 105 (10), 3125–3149.
- Jedwab, R., Vollrath, D., 2019. The urban mortality transition and poor-country urbanization. *American Economic Journal: Macroeconomics* 11 (1), 223–275.
- Kahn, M.E., Walsh, R., 2015. Cities and the environment. In: Duranton, G.J., Henderson, V., Strange, W.C. (Eds.), *Handbook of Regional and Urban Economics*. Elsevier, pp. 405–465. doi:10.1016/B978-0-444-59517-1.00007-6. *Handbook of Regional and Urban Economics*, Vol. 5 of, chapter 7.
- Kesztenbaum, L., Rosenthal, J.-L., 2016. The democratization of longevity: how the poor became old in Paris, 1880–1913. In: Ramiro Fariñas, D., Oris, M. (Eds.), *New Approaches to Death in Cities during the Health Transition*. Springer International Publishing, pp. 137–154.
- Kinney, P.L., Gichuru, M.G., Volavka-Close, N., Ngo, N., Ndiba, P.K., Law, A., Gachanja, A., Gaita, S.M., Chillrud, S.N., Sclar, E., 2011. Traffic impacts on PM_{2.5} air quality in Nairobi, Kenya. *Environmental Science and Policy* 14 (4), 369–378.
- Lagakos, D., Mobarak, A. M., Waugh, M. E., 2020. The welfare effects of encouraging rural-urban migration. Unpublished Working Paper, Boston University.
- Lagakos, D., Waugh, M.E., 2013. Selection, agriculture, and cross-country productivity differences. *American Economic Review* 103 (2), 948–980.
- Lambert, A.M., 1963. Farm consolidation in western Europe. *Geography* 48 (1), 31–48.
- Lee, V.E., Zube, T.L., Ross, K.N., 2005. School effectiveness in 14 sub-Saharan African countries: links with 6th graders' reading achievement. *Studies in Educational Evaluation* 31 (2–3), 207–246.
- Lipton, M., 1977. *Why Poor People Stay Poor: A Study of Urban Bias in World Development*. Temple Smith; Australian National University Press.

- Ludwig, J., Cook, P.J., 2001. The benefits of reducing gun violence: evidence from contingent-valuation survey data. *J Risk Uncertain* 22 (3), 207–226.
- Marx, B., Stoker, T., Suri, T., 2013. The economics of slums in the developing world. *Journal of Economic Perspectives* 27 (4), 187–210.
- Michalopoulos, S., Papaioannou, E., 2013. Precolonial ethnic institutions and contemporary african development. *Econometrica* 36 (3), 113–152.
- Morten, M., 2019. Temporary migration and endogenous risk sharing in village india. *Journal of Political Economy* 127 (1), 1–46.
- Morten, M., Oliveira, J., 2018. The effects of roads on trade and migration: Evidence from a planned capital city. Unpublished Working Paper, Stanford University.
- Munshi, K., Rosenzweig, M., 2016. Networks and misallocation: insurance, migration, and the rural-urban wage gap. *American Economic Review* 106 (1), 46–98.
- Nunn, N., Wantchekon, L., 2011. The slave trade and the origins of mistrust in africa. *American Economic Review* 101 (7), 3221–3252.
- Perez-Heydrich, C., Warren, J.L., Burgert, C.R., Emch, M.E., 2013. Guidelines on the Use of DHS GPS Data. Technical Report. Demographic and Health Surveys.
- Pope, C.A., Dockery, D.W., 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.
- Ravallion, M., Chen, S., Sangraula, P., 2007. New evidence on the urbanization of global poverty. *Popul Dev Rev* 33 (4), 667–701.
- Riley, J.C., 2001. *Rising Life Expectancy: A Global History*. Cambridge University Press.
- 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: Mieszkowski, P., Straszheim, M. (Eds.), *Current Issues in Urban Economics*. Johns Hopkins University Press.
- Stark, O., Lucas, R.E., 1988. Migration, remittances, and the family. *Econ Dev Cult Change* 36 (3), 465–481.
- Tatem, A.J., Gething, P.W., Smith, D.L., Hay, S.I., 2013. Urbanization and the global malaria recession. *Malar. J.* 12 (1), 133.
- Thomas, V., 1980. Spatial differences in the cost of living. *J Urban Econ* 8, 108–122.
- Topel, R.H., 1986. Local labor markets. *Journal of Political Economy* 94 (3), S111–S143.
- United Nations, 2018. *World urbanization prospects: The 2018 revision*, online edition. United Nations, New York.
- Werner, W., 1993. A brief history of land dispossession in namibia. *J South Afr Stud* 19 (1), 135–146.
- Williamson, J.G., 1982. Was the industrial revolution worth it? disamenities and death in 19th century british towns. *Explor Econ Hist* 19, 221–245.
- World Health Organization, 2006. *Air Quality Guidelines: Global Update 2005: Particulate Matter, Ozone, Nitrogen Dioxide, and Sulfur Dioxide*. Technical Report.
- World Health Organization, 2014. *Indoor Air Quality Guidelines: Household Fuel Combustion*. Technical Report.
- World Health Organization, 2015. *Nutrition Landscape Information System (NLIS): Children 6–23 months of Age Who Receive a Minimum Acceptable Diet*. Technical Report.
- Yap, L.Y., 1977. The attraction of cities: a review of the migration literature. *J Dev Econ* 4, 239–264.
- Young, A., 2013. Inequality, the urban-rural gap, and migration. *Quarterly Journal of Economics* 128 (4), 1727–1785.