

Urbanization in the Developing World: Too Early or Too Slow?

J. Vernon Henderson and Matthew A. Turner

Most regions of the world seem fully urbanized. North America, Europe, Latin America and the Caribbean, and West Asia all have shares of the population living in urban areas over 68 percent, with most regions near 80 percent. They also have small annual growth rates in this share, all under 0.62 percent a year and most near 0.25 percent (United Nations 2018). East Asia still has rapid urbanization, but its population is now over 60 percent urbanized and should soon top 70 percent, as in more developed regions. North Africa is only just over 50 percent urban, but that number is stable with little further urbanization. The global frontier of rising urbanization is sub-Saharan Africa (urbanization rate of 40 percent, annual growth rate of 1.4 percent), South Asia (urbanization rate of 36 percent, annual growth rate of 1.2 percent) and South-East Asia (urbanization rate of 49 percent, annual growth rate of 1.3 percent). Urbanization in these regions, and in sub-Saharan Africa in particular, will be the focus of much of our attention.

■ *J. Vernon Henderson is School Professor of Economic Geography, London School of Economics, London, United Kingdom. Henderson is also a Research Fellow of the Centre for Economic Policy Research. Matthew A. Turner is Professor of Economics, Brown University, Providence, Rhode Island. Turner is also Research Associate, National Bureau of Economic Research, Cambridge, Massachusetts, and Senior Research Fellow, Property and Environment Research Center, Bozeman, Montana. Both authors are Research Affiliates, International Growth Center, London, United Kingdom. Their email addresses are J.V.Henderson@lse.ac.uk and matthew_turner@brown.edu.*

For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at <https://doi.org/10.1257/jep.34.3.150>.

To understand the forces driving urbanization in developing countries, we begin by documenting key patterns and puzzles about urbanization and population density that have emerged in the literature. The classic economic model of urbanization is a story of technological change and structural transformation. Over generational time scales, people move from rural farms to urban factories in response to higher productivity in cities (for a review of this literature, see Desmet and Henderson 2015). East Asia and regions that urbanized prior to the late twentieth century seem to follow this path. However, sub-Saharan Africa is different. There, many countries are urbanizing “early”—that is, urbanizing at levels of per capita income generally far lower than when previous regions urbanized. Moreover, many cities in sub-Saharan Africa are growing without the expected increase in manufacturing or decline in agriculture. Perhaps related, many more farmers are living in urban areas in sub-Saharan Africa than we would predict from observing cities in other places and times. For a review on sub-Saharan Africa urbanization, see Henderson and Kriticos (2018).

Next, we offer evidence about costs and benefits of living in urban and rural areas in developing countries. In the first part of the paper, we rely on the dual sector model and its division into urban and rural. However, in the presentation of facts about costs and benefits, as in models in the modern literature (for example, Michaels, Rauch, and Redding 2012), we treat space as a continuum. We explore how various outcomes change with population density estimated at the level of a global grid of one-kilometer squares, using the Global Human Settlements Layer (GHSL) data. We show that significant fractions of the urban population in developing countries live at high densities that are practically nonexistent in the developed world. This suggests that the impacts of such densities can only be studied by looking at the developing world.

For data on a variety of outcomes, we use three geocoded surveys: the World Bank Living Standards and Measurement Survey (LSMS) for data on income and wages; the Demographic and Health Surveys (DHS) for data on a variety of amenities; and the Afrobarometer surveys for data on crime. We find that the high densities of developing world cities, in Africa and South Asia in particular, are associated with many benefits, including higher incomes and access to electricity, clean water, and inoculations. However, they also entail costs, including higher incidence of lifestyle diseases, poorer child health outcomes, and greater exposure to crime.

Finally, we turn to a discussion of spatial equilibrium and the differences between rural and urban life that inform the rational choice of location. The classic Roback (1982) model of spatial equilibrium suggests that people will move between rural and urban locations until they have equalized utility of living in the two areas. However, our results tend to confirm earlier findings that incomes and wages are far higher in urban than in rural areas of developing countries (for earlier work, useful starting points are Gollin, Lagakos, and Waugh 2013; Chauvin et al. 2017). Like earlier work, we also find that many urban amenities in rapidly urbanizing countries dominate those in rural areas (for example, see Gollin, Kirchberger, and Lagakos 2017). In short, utility levels seem higher in urban areas of developing countries.

To account for this, the classic model of spatial equilibrium has been modified in various ways. For example, structural modelling now incorporates moving costs or affinities for particular locations (for example, Tombe and Zhu 2019; Balboni 2019; Bryan and Morten 2019). Perhaps people are so attached to rural locations, or the rural-to-urban move is so costly, that the large apparent benefits of urban life are still not large enough to motivate migrating. It may also be that certain negative aspects of urban living may play a larger role in people's decision-making than previously recognized.

Much remains to be understood about how the drivers of urbanization in sub-Saharan Africa and other developing countries differ from the classic model of rural migrants heading for urban manufacturing jobs. But looking at the urban/rural gaps in income and amenities, we ask whether the true puzzle is not whether urbanization is happening too early, but rather, why it is not happening even faster.

Some Distinctive Patterns of Urbanization in Developing Countries

Early Urbanization?

Many low-income countries today are urbanizing “early”—that is, at historically low levels of income—with the prime examples being countries in sub-Saharan Africa (Lall, Henderson, and Venables 2017; Bryan, Glaeser, and Tsivanidis 2019). The nations of sub-Saharan Africa surpassed 40 percent urban share in 2010 at a GDP per capita of \$1,481. For comparison, Latin America passed the 40 percent mark in 1950 at a GDP per capita of \$2,500, while East Asia surpassed a 40 percent urban share in 2000 at a per capita GDP of \$5,451.¹ For reference, in 1900, per capita in Western Europe was at least twice that of sub-Saharan Africa today.

Why might urbanizing while poor matter? Early urbanization poses enormous challenges in governance. Poor countries cannot afford the ideal investments required to deal with the negative externalities of dense cities and are always playing a game of catch-up with rapid industrialization. Clustering of employment requires expensive transportation infrastructure to allow large numbers of workers to reach firms in city centers or peripheral industrial and commercial zones as well as to allow firms to get their goods to markets (Fujita and Ogawa 1982; Heblich, Redding, and Sturm 2018; Akbar et al. 2018; Tsivanidis 2019). The sewer systems and safe water supplies required to improve health and reduce mortality from disease (Kappner 2019) at high population densities are also expensive.

The problem goes beyond a simple lack of funds. Cities require institutions to collect taxes, keep order, and govern land. It is natural to suspect that the institutions

¹All GDP numbers here are expressed in 1990 dollars at the purchasing power parity exchange rate, based on Bolt, Timmer, and van Zanden (2014). In 2010, the comparable number for South and South East Asia was \$3,537, with South Asia still well under 40 percent urbanized today and South East Asia only having passed that mark in about 2005.

Table 1

Share of Manufacturing in GDP by Region and Year

<i>Region</i>	<i>1990</i>	<i>2000</i>	<i>2010</i>	<i>2017</i>
East Asia	24.6	25.2	27.6	27.4
South East Asia	22	24.8	22.6	20.9
Latin America and Caribbean	20.7	17.9	15.7	15.2
North Africa	17.6	17.9	16	16
Europe	17.5	15.3	11.9	11.8
South Asia	15.9	15.6	16.1	14.4
West Asia	14.4	13.2	12.1	13.8
Sub-Saharan Africa	13.8	11.6	8	9

Note: Data from the World Development Indicators 2018 are organized by UN regions. The table reports regional weighted averages using weights based on country share of regional GDP in 2017. Data cover 126 countries in a consistent sample over time. The Middle East is part of West Asia (not North Africa). Oceania is excluded.

and state capacity in these newly urbanizing areas reflect the lower income and education of the population.

What Is Driving Developing Country Urbanization?

The classic dual sector model of urbanization predicts that urban populations arise as farmers move to cities to work in factories making manufactured goods. Countries like Brazil and Argentina each had about 30 percent of GDP in manufacturing in 1980 even as urbanization was starting to slow, while China was over 40 percent in 1979 as urbanization was just taking off (World Bank 1981). Table 1 shows regional patterns for 1990 onwards where we have a large enough sample of countries reporting data for all listed years. As of 2017, East Asia has 27 percent of GDP from manufacturing, China about 29 percent, and South East Asia 21 percent. East and South East Asia maintained high manufacturing shares over the whole 1990 to 2017 time period. Latin America's manufacturing share started at over 20 percent in 1990 and declined to just over 15 percent.

In contrast, the 33 countries of sub-Saharan Africa that our data describe (excluding South Africa) have the lowest regional share of manufacturing world-wide in 1990—a share that has only declined over time. While other regions have experienced declines in manufacturing share, they tend to be countries with high income levels that are deindustrializing in favour of traded services. In general, most of sub-Saharan Africa has never developed a manufacturing sector beyond production of traditional goods for within-country consumption.

In short, sub-Saharan Africa and parts of South Asia have relatively few manufacturing employees and their numbers are growing slowly. So what is driving urbanization in these regions? We consider several possibilities, but there is no agreed-upon answer.

One possibility is that the current wave of developing country urbanization is being led by consumption opportunities, including urban amenities rather than production. A literature on “consumer cities” began with Glaeser, Kolko, and Saiz

(2001) and was extended to developing countries by Gollin, Jedwab, and Vollrath (2016). The latter paper demonstrates that in Asia and Latin America, there is a strong positive correlation between urban share and the GDP share of manufacturing and services. However, no such correlation exists in Africa and the Middle East. They conclude that urbanization in sub-Saharan Africa and the Middle East is driven by rents from natural resource exports, which they conjecture are distributed in cities. Such rents then can fund civil servants and urban services as well generate demand for urban private services. Also, one may draw a connection between natural resources rents and low manufacturing, based on the argument that revenues from natural resource exports affect exchange rates and wage costs, crowding out manufacturing and its technological spillover benefits (Sachs and Warner 2001; Ismail 2010; Alcott and Keniston 2017).

Henderson and Kriticos (2018) reexamine the consumer city argument. While they confirm the finding that urbanization in sub-Saharan Africa is not correlated with the manufacturing and services share of GDP, they find that increases in natural resource rents are also not associated with increased urbanization in Latin America and sub-Saharan Africa. More generally, countries without natural resource rents are urbanizing too. Simply put, variation in urbanization within sub-Saharan Africa is not well explained by GDP shares in manufacturing, services, and resource extraction. Perhaps future research will find that the lack of definitive patterns reflects measurement error or outliers in data from sub-Saharan Africa. But if urbanization in sub-Saharan Africa is not a consequence of traditional structural transformation and is not well related to natural resource rents, then what are other possibilities?

Perhaps urbanization in Africa is not so much about the benefits of urban density in Africa but more a consequence of especially low rural productivity and offerings of services. Agricultural productivity in Africa is low by global standards, reflecting low irrigation rates, low fertilizer usage, and an attachment to old seed technologies (Ray et al. 2012; Sánchez 2010). Cereal yields in sub-Saharan Africa are half those of South Asia, which in turn are half those of high-income countries, and well below East Asia and Latin America (Henderson and Kriticos 2018). Low rural productivity helps to explain why urban incomes are comparatively so much higher than rural incomes in Africa, conditional on education, age, gender, and the like (Henderson, Kriticos, and Nigmatulina 2019).

A seeming oddity is that sub-Saharan cities house a surprising number of farmers. Table 2 reports results for a set of twelve countries with a total population of 220 million for which there is relevant data in the Integrated Public Use Microdata Series (IPUMS). The first row shows for different spatial entities, the fraction of workers who report the industry in which they primarily work as agriculture, while row 2 does the same for manufacturing. Thus, for example, in row 2 and column 2, less than 2 percent of workers living in rural areas report the main industry in which they work as being manufacturing. In columns 1 and 2, we report these fractions for all workers living in census-defined urban versus rural areas. The remaining columns isolate the primate (largest) cities in each country, and then those in the top 25 percent by size within each country (excluding the primate),

Table 2

Farmers in African Cities by City Size

<i>Spatial scale</i>	<i>All urban</i>	<i>All rural</i>	<i>Primate city</i>	<i>Secondary cities (top 25%)</i>	<i>Tertiary cities (50–75%)</i>	<i>All others</i>
Percentage of workers reporting agriculture as main industry	20.5	88	8.5	23.8	38.6	41.3
Percentage of workers reporting manufacturing as main industry	10.6	<2	12.4	10	8.3	7.3

Source: Henderson and Kriticos (2018) Figure 3 and Supplemental Figure 2.

Note: Data from IPUMS for the most recent census for Ethiopia, Tanzania, Uganda, Mozambique, Ghana, Cameroon, Mali, Malawi, Zambia, Sierra Leone, Liberia, and Botswana. Small cities are in the bottom 50 percent of cities by size and tertiary cities are in the 50–75th percentiles. Cities are defined by night-light boundaries to which population is assigned.

and those in the 25–50 percent, 50–75 percent, and the bottom 50 percent by city size. The share of agriculture in city employment by city size type ranges from 9 to 41 percent and averages 20.5 percent. In fact, in the bottom 75 percent of cities by size, the share of agriculture in urban employment averages 40 percent in this sample. In contrast, in Brazil, India, and Malaysia, shares of urban farmers are all under 7.5 percent. Table 2 also shows that in these sub-Saharan African countries, 88 percent of rural sector employment is in farming. This is far higher than other countries, where rural services, construction, and even manufacturing employment are more important. Finally, note the especially small manufacturing share in smaller cities and towns. Most likely, any manufacturing in these places is traditional food processing, non-metallic minerals, locally made furniture, weaving, and the like for local consumption.

The table tells us African cities are home to a substantial number of farmers. Why do farmers move to cities and by inference commute out to farms? One answer may be better access to amenities and public services as well as consumer services. This in turn may be related to the absence of almost anything but farming in rural areas, which may also reflect a lack of rural infrastructure and institutions. Another answer may be the better employment opportunities in cities for other family members, both in terms of hours worked and the diversity of occupations available (Henderson et al. 2019). Moreover, the large number of urban residents who report their primary occupation as farming may also work in the off-season in other occupations.

Apart from so many farmers living in cities, there is a literature suggesting that Africa may bypass the development of modern manufacturing. The papers in Newfarmer, Page, and Tarp (2018) suggest that African development may rely more on tourism and aspects of information technologies as well as work related to farming such as food processing and horticulture. Henderson and Kriticos (2018) show for a sample of five countries that tradable urban services, like finance, are

growing at extraordinary rates, albeit from a very low base. With that being said, evidence for the current level and trajectories of urban employment by sector is fragmentary. Understanding how sub-Saharan Africa is urbanizing remains a subject of debate, and one that would benefit from more and better data.

Density and Population

Up to this point, our discussion has used definitions of *urban* that are based on host-country specific definitions and implemented using data that may also vary qualitatively from country to country. Unsurprisingly, these definitions are not consistent across countries and may involve subjective assessments like whether an area has certain public facilities, administrative responsibilities, or has a central economic core. Moreover, the extent of metropolitan areas is typically based on the boundaries of country-specific administrative units (like counties in a US context). In some cases, national definitions, especially for capital cities, have tended to severely restrict official urban area size on the basis of historical criteria like a defined national capital zone (as is the case for Jakarta).

One way to avoid such classification problems is to focus instead on population density. Density can be used to define urban areas based on density cut-off points as for example in the Global Human Settlement data (GHS-SMOD L1). However rather than use arbitrary cut-off points and, as noted above, to be consistent with modern modelling, we allow density to vary continuously across space. This is also in line with evidence for developing countries that agglomeration economies arise from density rather than absolute labor market size (Chauvin et al. 2017; Combes et al. 2020; Quintero and Roberts 2018; Henderson, Kriticos, and Nigmatulina 2019).

Which data sets give us finely gridded densities? Perhaps the best-known is the Gridded Population of the World version 4 (GPWv4; CIESIN 2018), which uses population data for country-specific administrative or enumeration units used in their census. GPW sets up the world in grid cells of (approximately) one kilometer. However, GPW then has to map the census unit data into these grid squares, where census enumeration units may be larger or smaller than these grid cells. The census units for which data are released may be quite large administrative units, such as a county or even a province. In these cases, GPW prorates enumeration unit population to grid cells by assuming population is spread uniformly over each reported unit. While high-income countries like the United States often release population data on a fine spatial scale, that is not the case in developing countries. For example, of the 12.9 million polygon-shaped administrative units that form the basis for population estimates in the global GPW, only 2.4 million are from outside the United States.

Rather than use the GPW, we use the European Union's Global Human Settlements population layer (GHS-POP from Schiavina, Freire, and MacManus 2019; Freire et al. 2016). The GHSL again allocates GPWv4 population estimates across one-kilometer grids, but instead of assuming that population is evenly distributed across a polygon-shaped enumeration area, it allocates people according to the

spatial distribution of the footprint of built cover within each area. “Built cover” is based on the EU’s specific processing of Landsat data 30-meter resolution satellite data circa 2015 (Corbane et al. 2018, 2019).² In the rare cases where there is no built cover in an enumeration polygon, it reverts to the GPWv4 estimates. More information about the GHS data can be found in Florczyk et al. (2019).³

Based on these data, we present information about population density per square kilometer for grid squares whose size is one square kilometer at the equator. We compare Europe, North America, sub-Saharan Africa, Latin America, and South Asia. We pool East and Southeast Asia together to improve the legibility of our figures.

In Figure 1a, we graph the cumulated share of population by density. Clearly, North America and then Europe have the highest accumulated shares of population at low densities. In America and Europe, less than 10 percent of the population lives at densities above 10,000 people/square kilometer. Sub-Saharan Africa, along with East and South East Asia, have the lowest accumulated shares at low density, or equivalently, the highest degree of density inequality. In sub-Saharan Africa and in East and Southeast Asia, 30-40 percent of the population lives at densities above the 10,000 threshold, while for Latin America and South Asia, it is about 20 percent.

To improve legibility, our graphs stop at 20,000 people/square kilometer. In Southeast and East Asia, 18-20 percent of the population lives at densities above 20,000/square kilometer. In the developed world, the proportion of people living at such densities is tiny. For the purpose of understanding density and its implications, the developed world probably cannot teach us much about the very high densities experienced by a significant portion of the developing world’s population.

Figure 1b graphs the corresponding cumulated share of population by area. It shows only the 3 percent of regional area that is most densely populated. Looking


²Landsat data is an alternative source of gridded population data. These data rely on a proprietary algorithm to construct population estimates based on higher resolution satellite imagery than Landsat and information on airports and rails (Rose and Bright 2014). The algorithm is not publicly documented and changes from year to year. Moreover, the estimates are for ambient population averaged throughout the day, whereas GHS-POP is for the nighttime (residential) population. We choose the GHS data because it is consistently defined over time and the algorithm is public. One issue concerns how all these data sets deal with the vast number of grid squares with very low or zero population worldwide. The GHSL, Landsat, and GPW handle the problem very differently. However, the densities we look at for our purposes (say, above 8 people per square kilometer), based on other work in progress, the distributions are quite close.

³As noted in the text, the GHS Settlement Model (GHS-SMOD L1) also attempts to define city status based on density and population cut-offs. Starting from gridded population data, “cities” are defined as sets of contiguous one-kilometer grid cells having density over 1,500 people and summing to over 50,000 people. The settlement model also constructs “towns” and suburbs,” defined as sets of contiguous pixels with density and size thresholds of 300 people per square kilometer and 5,000 in total. This approach has the advantage of avoiding administrative boundaries for classifying urban areas, but it also seems arbitrary. For example, it would be hard to argue that an agglomeration of people satisfying such a definition would accurately describe the actual labor market or commuting zone of a city. Indeed, appropriate density cut-offs would probably vary by country and region of the world. Our discussion in the text focuses only on density and ignores the definitions that would emerge from using these population cut-offs.


Figure 1

Population Density Gradients by Region

A: Cumulative share of population by density



B: Cumulative percentage of population by land area in the region

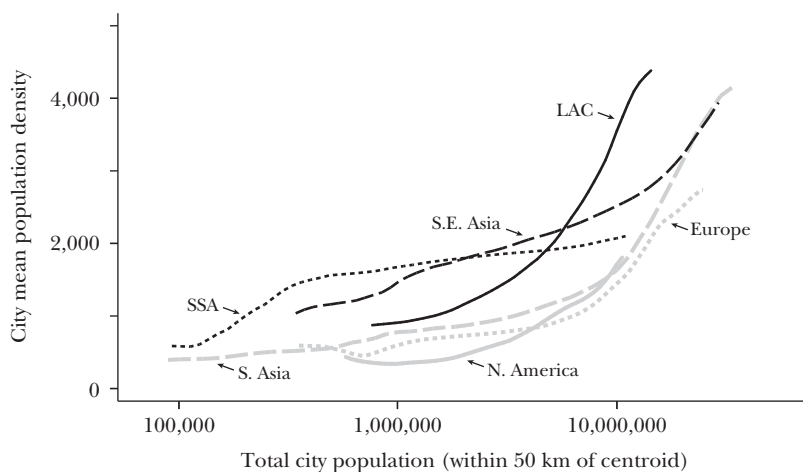


Source: Based on population data from GHS.

Note: SSA—sub-Saharan Africa, LAC—Latin America and the Caribbean.

at the y-axis in this figure, we see that about 25 percent of the population of South and South East Asia occupies the 97 percent of regional area that is the least densely populated, while in other regions that population share is small—especially in North America and Europe. As a result, almost everyone lives on less than 3 percent of the land area. The more widespread occupation of land in South and South-East Asia reflects two factors: 1) a larger fraction of Asian land employed in labor-intensive agriculture; and 2) in much of Asia there is a relatively high ratio of national population to land area, forcing use of a greater proportion of land. But even in South and South East Asia, there is still a lot of room for people to live at lower density: just 25 percent of the population occupies 97 percent of the land area.

Figure 2

City Population Density by Total City Population

Source: Based on population data from GHS.

Note: Vertical axis is mean population density from GHS in a 50km radius disk centered on the centroid of each of the 657 UN world cities. Horizontal axis is total population in the same disk.

Combining the results of the two parts of the figure, we note that North Americans live at relatively low densities but are endowed with a large land area, so most land is very sparsely populated. In contrast, many Africans live at high densities as we saw in Figure 1a, while in Figure 1b, most land is also sparsely populated. There are many other factors apart from regional land availability per person determining the patterns of population density and land use that we see in Figures 1a and b.

Figure 2 describes the relationship between city size and urban population density but with a common geographic measure of a “city.” To construct Figure 2, we use the 657 cities described by the UN World Cities data. These are cities that housed more than 300,000 people at any time between 1950–2010. For each city, the UN World Cities data reports the latitude and longitude of the center of the city. We draw a 50-kilometer radius around each such centroid and sum the population in the gridded population squares within this disc. To measure population density for these cities, we calculate the person-weighted density of grid cells within each city’s disk. Figure 2 presents local polynomial regressions of the relationship between city population and density by region.

In most regions of the world, higher densities are associated with larger city populations. This pattern is clear in North America. However, in Africa the relationship is weak. For most regions of the world, mean population density in the 50-kilometer disk rises quickly with total population above about 450,000, although in North America, the take-off point is near 750,000. Africa differs. Below 450,000 and above about 2 million, African cities have higher densities than in other regions. As city sizes increase in Africa, density rises relatively slowly.

Evidence on How Living Conditions in Developing Countries Vary across Density

When studying agglomeration economies and diseconomies, with a typical focus on cities, researchers have employed different scale measures such as total city employment or population or a measure of density. When looking at total city population, the researcher is largely constrained to accept an administrative or other boundary of an urban area, which then makes an implicit assumption that any resulting agglomeration economies operate uniformly within that boundary. Here, we will instead consider how a variety of outcomes vary continuously with population density. This not only allows for within urban area variation but treats space as a continuum (as in Michaels, Rauch, and Redding 2012; Desmet and Rappaport 2017).

Our empirical methodology is straightforward. To learn more about how peoples' lives change with population density, we measure density using the Global Human Settlements data described in the previous section. For measures of outcomes, we turn to three sources noted in the introduction: the Living Standards and Measurement Survey for data on income and wages in six sub-Saharan African countries; the Demographic and Health Surveys for data on female outcomes, child outcomes, infant mortality, household utilities, schooling, and adult lifestyle outcomes for 40 countries in Latin America, South East Asia, sub-Saharan Africa and South Asia (with a focus on the last two regions), generally conducted from 2010–2016; and Afrobarometer for data on crime in 24 sub-Saharan African countries. Details on the DHS and Afrobarometer are in Henderson et al. (2020). As is common, these surveys use a “cluster” approach of questioning randomly selected people near a smaller number of randomly selected “cluster” points and assign all such respondents the location of the cluster.⁴ With our gridded data, we can draw a five-kilometer disc surrounding each cluster, and in this way, we can match geocoded individual-level surveys to population density. Thus, we are able to examine how survey responses describing income, health, education, public health, and public goods vary with nearby density in a large sample of developing world countries.⁵

⁴For the Afrobarometer and Living Standards and Measurement Survey, clusters are generally located at the centroid of a small administrative unit, such as the finest census enumeration unit. To protect respondent privacy, clusters in the Demographic and Health Surveys are displaced by up to two kilometers for urban respondents and up to five kilometers for rural respondents. This introduces some error into the respondent relevant measure of population density. We truncate respondents in five-kilometer radiuses with population densities less than 7.4 people per square kilometer because we are suspicious of the accuracy of GHSL estimates at low population densities. We also observe dramatically wider confidence bands around non-linear regressions of outcomes on log density below this threshold.

⁵Our approach is conceptually similar to that of Gollin, Kirchberger, and Lagakos (2017), who look at the relationship between various outcomes reported by the Demographic and Health Surveys and population density in an area around clustered survey respondents. They find that survey respondents living at the 80th or 90th percentile of the set of DHS cluster densities typically have better amenities than those for people living at the 20th or 10th percentile. These results are interesting and important, but somewhat difficult to interpret. As we saw earlier in Figures 1a and 1b, population is highly concentrated into very small, very dense regions in sub-Saharan Africa and South and South East Asia. The 80th or

To illustrate our results, we focus on figures constructed using the “binscatter” methodology described in Cattaneo et al. (2019). In our figures, we show the outcome for an (endogenous) number of equal size bins. The confidence bands describe the region around local polynomial regressions in which we expect a local polynomial regression line to lie with 95 percent probability. In each figure for the left-panel non-parametric estimates, we do not include control variables. In the right-panel semi-parametric regressions, we include country fixed effects and a range of control variables, which differ somewhat by outcome according to various factors like whether the outcome in question reflects a household, person, or child-level outcome, or what was included in the survey instrument. Broadly, the control variables reflect the education, gender, and age of the household, person, or child whom the survey response describes.

The figures also report a line of best fit and its slope coefficient. If this best-fit line falls outside the confidence band from the binscatter, then a linear relationship can be rejected—at least locally. While generally the regression lines lie within the confidence bands, the illustrative graphs on which we report show very different widths for confidence bands. In an online Data Appendix available with this paper at the journal’s website, we offer a detailed presentation of ordinary least squares regressions, with and without control variables, as well as the list of specific controls, the countries, and the like.⁶ We note that the density gradients we report can be based on quite different samples of countries, and so some caution is required in comparing regression results across outcomes.

This methodology has well-known weaknesses, and the evidence presented here should be viewed as a suggestive first pass in analyzing multiple data sets, which deserve more in-depth work. Our results are associations and do not give magnitudes of true causal effects. For example, while we will find a rapid rise in income with density even with control variables, omitted variables such as ability and ambition are surely also important, and may influence how people sort across rural and urban locations. Part of the association of higher incomes with density could be that, conditional on education, higher ability people may live at higher densities. Of course, higher ability people may benefit more from higher densities, so density effects are heterogenous. As we will note below, controlling for education in income regressions may lead us to understate some benefits of density for those with lower incomes, in terms of facilitating better schooling.

Resolving these inference problems is difficult and beyond the scope of this project. A few experiments have sought to induce random variation in subject

even 90th percentile of DHS clusters by density is not very dense, especially given the enormous rural over-sampling in the DHS. Implicitly, the Gollin, Kirchberger, and Lagakos (2017) methodology is telling us about the distribution of amenities across places rather than across people according to how they live. Given the small proportion of the landscape occupied by cities, it is hard to interpret these findings in terms of a difference between the rural and urban experience.

⁶We note the similarity between results presented here and those in Henderson et al. (2020), which examines how outcomes differ across the discrete urban-rural classifications given in the GHS Settlement Model (GHS-SMOD L1) noted earlier.

locations, but most are within city or involve refugee and other programs applied to very special populations (reviewed in Bryan, Glaeser, and Tsivanidis 2019). Extending these experimental and quasi-experimental methods to the larger set of outcomes that we consider is an obvious area for further research.

Incomes

The LSMS provides survey data on income and wage (for hourly workers) for six African countries—Ethiopia, Ghana, Malawi, Nigeria, Tanzania, and Uganda—with a combined population of over 400 million people. Household income is constructed from several LSMS questions. It includes all wage income and business receipts (including farm), less business expenses per month (for details of variable definitions, see Henderson and Kriticos 2018).

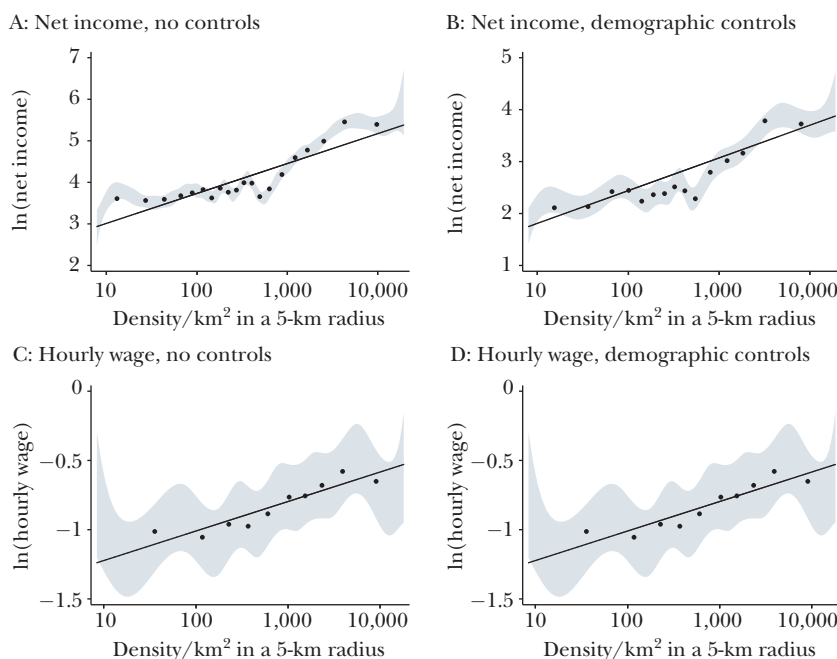
The top two graphs of Figure 3 show the binscatter plot of (the log of) net household income against (the log) of density, with and without control variables. The bottom two graphs show a similar plot for wage data. The estimated elasticities from the best-fit line reported in the figure are high. Doubling density increases household net income by about 32 percent and hourly wages by about 5 percent, with controls. At 5 percent, the density elasticity of hourly wages exceeds those typically found in developed countries but is in the range of estimates in recent work on other developing regions and countries (for example, Quintero and Roberts 2018; Duranton 2016; Combes et al. 2020; Chauvin et al. 2017). Yet the income elasticity is probably more important. After all, it is families that migrate permanently to cities. The fact that the density elasticity of net income is a multiple of that for wages likely reflects both an increase in hours worked and varieties of job opportunities for family members, as analyzed in Henderson et al. (2019). We know of no comparison for the density elasticity of net income in the literature.

While the figures suggest that a linear fit is reasonable, the graph also suggests a potential non-linearity. The density gradient is flatter from about 8 to 550 people/square kilometer. This is well below the average density of African cities (shown earlier in Figure 2). We do not think this flat portion has to do with measurement of income. Incomes appear to be measured as well in this low-density part of the graph as other parts, given the detail and high standards of the LSMS. After this point, the gradient increases sharply such that a household moving from a density of 550 to 8,100 people/square kilometer shows about a four-fold increase in income. While the LSMS reports respondents at densities near 20,000 people/square kilometer, such respondents are rare and so our estimates of income at these high densities are imprecise. The corresponding plots for hourly wages are similar but with a less steep slope and modestly wider confidence bands. In all, these estimations indicate that African wages and income are sensitive to density and suggest that moving to denser locations may have a high return for African families.

Utilities and Schooling: Public Goods Strongly Influenced by Policy

Household access to utilities and schooling depend in greater part on public sector provision of, for example, water mains, reservoirs, schools and teachers. The

Figure 3

Log of Household Net Income and Hourly Wage versus log Population Density


Note: Binscatter plots of LSMS net income of respondent household and of hourly wage, against the log of GHS population density in a 5km disk around the survey respondent. Log population density is censored below at about $8/\text{km}^2$. Left panels have no controls. Right panels include demographic controls and country fixed effects. Shading indicates 95 percent confidence band. Income includes wage income, net farm income, and net business income. For a small number of observations expenses exceed (monthly) incomes. We drop these observations to permit logarithmic scaling. LSMS survey countries are listed in table A2. Linear regression based on results in table A1a, which provides more details about the sample. Slope coefficients and standard errors of best linear fits are; (a) 0.313 (0.016) (b) 0.317 (0.014) (c) 0.118 (0.015) (d) 0.049 (0.009). Details in online Appendix.

Demographic and Health Surveys ask questions about electricity, safe drinking water, improved sanitation, and educational attainment. The questions about water and sanitation are nuanced and tailored to allow an evaluation of whether the UN's sustainable development goals are attained. For example, "safe" water can be quite different than piped water. In cities in sub-Saharan Africa and in South Asia (as defined in the GHS settlement model) about 40 and 80 percent of people have access to "safe" water, but only about 8 and 25 percent respectively have water piped into their dwelling unit (Henderson et al. 2020). Toilets flushing into a central sewer system are rare in these cities.

The top two graphs of Figure 4 present a scatterplot plot where the outcome variable is the indicator for improved sanitation. Even after controlling for household demographic characteristics, we see a rapid and precisely estimated increase in

Figure 4

Access to Improved Sanitation and Probability of Children Receiving Eight Years of School versus log Population



Note: Binscatter plots of a DHS indicator variable that is one if a respondent household has access to improved sanitation and of an indicator that is one if a household child 16 years old completed eight years of school, against the log of GHS population density in a 5km disk around the survey respondent. Log population density is censored below at about 8/km². Left panel is unconditional. Right panel includes demographic controls and country fixed effects. Shading indicates 95 percent confidence band. DHS survey countries are listed in table A2. Linear regression based on results in table A1a, which provides more details about the sample. Slope coefficients and standard errors of best linear fits are; (a) 0.083 (0.001) (b) 0.063 (0.001) (c) 0.050 (0.001) (d) 0.016 (0.001). Details in online Appendix.

access to improved sanitation with density. As in the earlier case of net income, we also see a slow increase in access to improved sanitation at lower densities and more rapid increase at higher densities. Going from 550 people to 8,100 people/square kilometer raises the likelihood of improved sanitation from under 25 percent to over 50 percent. There is also evidence of a downturn at very high densities. This may reflect a decline in services to high-density slums, but our limited sample of very high-density respondents does not allow precision at this tail.

A figure for access to safe water looks similar, including the non-linearity at high densities. For electricity, the fit with control variables included is very tight, although the rise is more linear. With mean outcomes of 0.5 to 0.7 for these three utilities, a one-standard deviation increase in the log of density increases outcomes from 0.075 to 0.11. We believe these differences are supply-driven and reflect lower

costs of service provision in dense areas, perhaps along with political considerations. Denser areas may be more favored in the political arena as in the classic urban bias literature (for example, Ades and Glaeser 1995; Davis and Henderson 2003).

The bottom two graphs of Figure 4 show a relationship between density and schooling. In our estimates, the schooling outcome is for 16-year-olds and is an indicator variable that takes value one in the event that a household 16-year-old has completed at least 8 years of schooling, and zero otherwise. For the best-fit line in the left-hand figure, increasing density by one log point increases the share of 16-year-olds completing eight years of schooling by about 0.050. Controls reduce this effect by two-thirds to 0.016 in the right-hand figure, so that a one standard deviation increase in $\ln(\text{density})$ raises the probability by 0.027 for a sample mean of 0.61.

Density effects for schooling are smaller than utilities. However, why should schooling attainment, after controlling for family characteristics, be affected at all by density at all? One possibility worth exploring is a more reliable supply of schooling and teachers in denser areas.

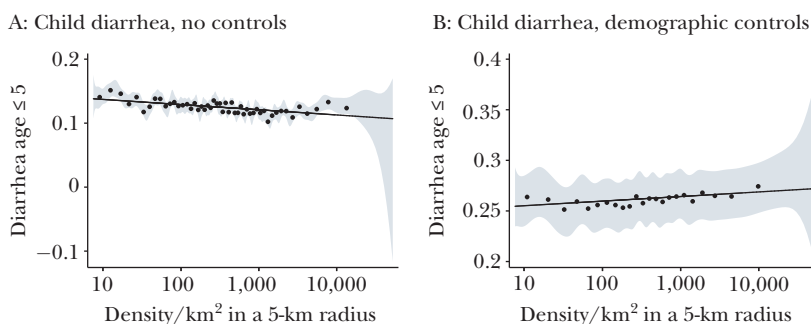
This raises an important issue. In examining the income and wage returns to density, we tried to control for sorting by controlling for education. However if higher density in developing world cities has a causal effect on human capital accumulation for families who move to density, education ought to be seen as part of the benefit of density, not as a sorting effect that should be held constant when looking at higher urban incomes. How to separate out these components is a subject for future research.

Female, Child, and Birth Outcomes

The DHS also reports on a variety of indicators related to the status and well-being of women and children. For example, indicator variables include: the use of modern contraception for sexually active women ages 20-40 who are not pregnant and do not want to have a child in the next two years; reporting an affirmative response by females to the question “is wife beating justified for any reason?”; and if a woman reports having ever experienced spousal household violence. The data also include the total number of births in the last three years to each woman age 15–49, and whether each child born from three months to three years ago survived at least three months. For each household child, there are indicator variables reporting whether the child has had the third and final DPT3 immunization shot by two years of age, whether the child has had diarrhea in the last two weeks, and whether each child age five and under has had a cough in the last two weeks.

As one illustration of the general findings, Figure 5 presents a binscatter plot of the relationship between the incidence of childhood diarrhea and density. First, as in this figure, best-fit lines for these outcomes indicate small marginal effects of density, but generally, incidence rates are also low. For example, here the (significant) slope in the left-hand graph is -0.0035 for an average incidence of 0.125. Second, all unconditional outcomes improve as density rises, except for being a victim of spousal abuse and whether a child five-or-under has had a cough recently.

Figure 5

Diarrhea Last Two Weeks for Children Five and under versus log Population Density

Note: Binscatter plots of a DHS indicator that is one if a child five or under had diarrhea in the past two weeks, against the log of GHS population density in a 5km disk around the survey respondent. Log population density is censored below at about 8/km². Left panel is unconditional. Right panel includes demographic controls and country fixed effects. Shading indicates 95 percent confidence band. DHS survey countries are listed in table A2. Linear regression based on results in table A1b, which provides more details about the sample. Slope coefficients and standard errors of best linear fits are; (a) -0.004 (0.0005) (b) 0.003 (0.0004). Details in online Appendix.

Third and most critically, using control variables changes the picture considerably. In a number of cases, demographic controls reduce density coefficients by well over 50 percent. However, most critically, in the case of diarrhea as illustrated, along with cough and infant mortality, effects are actually reversed, and being a victim of spousal abuse, having diarrhea, having cough and infant mortality increase significantly with density once controls are added. After controlling for demographic characteristics, one standard deviation increase in density is associated with an increase in domestic violence, diarrhea, cough, and infant mortality of 3.5 to 5 percent relative to their means.

Finally, with the addition of controls, the confidence bands expand dramatically, as illustrated in the right-hand panel in Figure 5. This huge widening of confidence bands once controls are added applies to most of the outcomes in this sub-section (with the exception of fertility and spousal abuse). This means that we can have less confidence in the local precision of marginal density effects for most outcomes discussed in this section despite significant slopes to best-fit lines. To put it another way, the relationships among density, demographic controls, and these outcomes need much more investigation and may be more subject to unobserved features of the local environment.

The finding that diarrhea may rise with density may seem at odds with the finding in the previous subsection that that safe water and improved sanitation both improve with density. One possible interpretation is that, as density rises, the increased access to safe water and improved sanitation is not enough to offset the effect of increased crowding on contamination of food and water. Another

possibility is that the UN sustainable development goals are setting too low a bar: what is being counted as safely managed water and improved sanitation is not clean enough.

Lifestyle Diseases and Crime

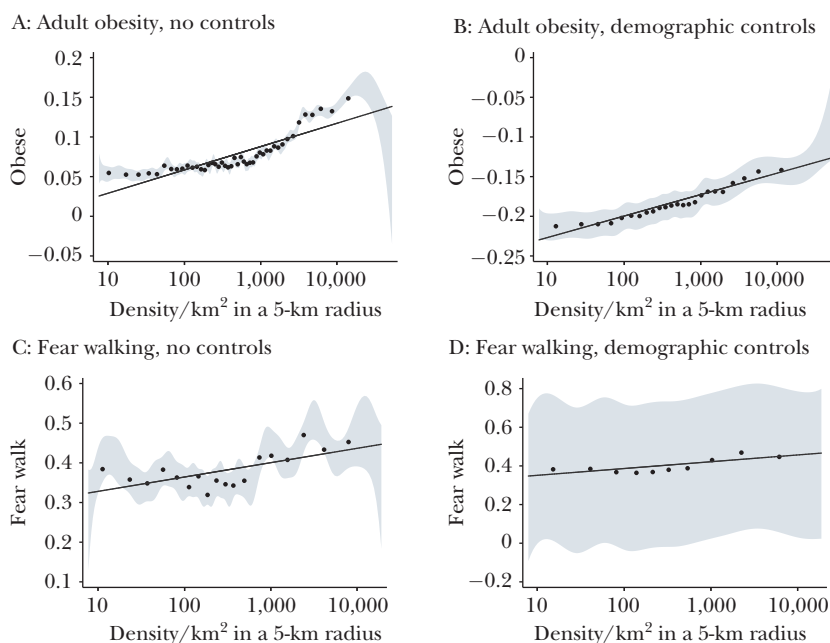
The DHS data allows us to report on the relationship between population density and four lifestyle diseases for adults age 20–49. Obesity data exist for all our countries. For India and Nepal, there are data on the incidence of measured high blood pressure, and for India alone, there are data on self-reported asthma and diabetes. These are to some extent lifestyle diseases. They reflect at least in part patterns of diet, exercise, work intensity, and stress, which may be affected by higher density. We can imagine that stress might come from long commutes or hours worked, smaller social networks, changed family circumstances, and crowding.

The top panel of Figure 6 shows the scatterplots for obesity, which is defined here as having a Body Mass Index (BMI) above 30. (The formula for BMI is weight in kilograms divided by height in meters squared.) The incidence of obesity, high blood pressure and diabetes all increase with density. Controls have a small impact on marginal outcomes, either raising or lowering them by less than by 25 percent. In the right-hand graph, we see that the slope of the best-fit line is 0.010 for a mean of 0.077, so that an increase of one standard deviation in density (1.7) is associated with a 22 percent increase in obesity from the mean.

We note that asthma does not respond to density. This perhaps surprising result would seem to be consistent with results in Aldeco, Barrage, and Turner (2019). Using global data, that paper finds the relationship between population density and the concentration of airborne particulates to be unambiguously positive but quite small: that is, air pollution is worse in cities than in rural areas, but not much worse.

The Afrobarometer survey collects data from 26 African countries about four sets of feelings or outcomes about crime: whether the survey respondent reported being fearful while walking outside in their neighborhood; whether the respondent reported being fearful of crime at home; whether the survey respondent's home has been robbed in the past year; and whether anyone in the household has been attacked by an outsider in the past year. Results are fairly similar for all these outcomes, in terms of marginal density effects relative to average incidence with effects rising with density. For illustration, the bottom panel of Figure 6 shows a binscatterplot for being fearful while walking outside, where a one standard deviation (1.8) increase in density is associated with a 0.029 increase in fear, for an average incidence of 0.38. While the left-hand graph suggests a sharp rise in fear in the higher density ranges, in the right-hand graph, the confidence intervals on local marginal effects with controls are very large. Finally, we note that, for actually being attacked in the past year, the slope of the best-fit line is insignificant and the incidence at 0.10 is much smaller than those for other outcomes including, from above the fear of being attacked. Perhaps those with a greater fear take more precautions to avoid actual incidence.

Figure 6

Adult Obesity and Self-Reported Fear of Walking Outside versus log Population Density

Note: Binscatter plots of a DHS indicator that is one if the survey respondent is obese or reported being afraid for their safety while walking outside, against the log of GHS population density in a 5km disk around the survey respondent. Log population density is censored below at about 8/km². Left panel is unconditional. Right panel includes demographic controls and country fixed effects. Shading indicates 95 percent confidence band. DHS survey countries are listed in table A2. Linear regression based on results in table A1b, which provides more details about the sample. Slope coefficients and standard errors of best linear fits are; (a) 0.013 (0.0005) (b) 0.010 (0.0003) (c) 0.016 (0.004) (d) 0.016 (0.003). Details in online Appendix.

Summary

While incomes, wages, access to utilities, schooling, number of births, and use of contraception and vaccinations all improve with density, we see declines with density in child and adult health outcomes, including infant mortality, obesity, domestic violence, and fear of crime. Of course, some people will worry more about lifestyle diseases or crime than others, although child health may be harder to put aside. Regardless, those who place a heavier weight on these factors may be less likely to migrate to urban areas.

The Roback Model Meets Current Patterns of Urbanization

The Roback (1982) model is the workhorse model for thinking about spatial equilibrium. In the original model, people are identical and move across space to

equalize utility levels. An important innovation in the recent literature has been to introduce moving costs and different forms of individual heterogeneity. With these additions, agents no longer move across locations to equalize utility levels. Rather, all agents choose their favorite location, taking into account the cost of moving there from their starting point. To illustrate, suppose a continuum of people choose between an urban and rural location. In the most general set-up, people receive a location and individual specific income, an “amenity” that is valued similarly by all agents (for example, up to income effects), and an “affinity” draw that is person- and location-specific. Amenities represent location specific attributes like the availability of safe water, the prevalence of crime, or the difficulty of commuting. Affinities reflect things like a personal taste for local weather or landscape or the presence of family members or roots in a home location.⁷ Finally, to move between locations, agents must pay a migration cost.

In a static spatial equilibrium, no one can gain by moving, at least not after accounting for the costs of moving. The notion of spatial equilibrium provides us with a powerful framework for organizing our ideas about what causes different people to arrange themselves across the landscape in the ways that we observe. At its heart, the model assumes that people act to arbitrage spatial differences in productivity and amenities by changing locations. Their ability to conduct such moves is hampered by frictions, moving costs, and idiosyncratic attachment to locations.

Consider the simple case in Moretti (2010), where there are no moving costs. All people *within* a region receive identical real incomes and amenities, and at the margin, real incomes are declining in population in each of the two regions. Then, there is a marginal person whose affinity draws make that person exactly indifferent between living in the two regions. For example, in the urban region, everyone has a weaker relative affinity for *the rural region* than the marginal agent who chooses the urban region. Note that in the resulting equilibrium, utility levels are not equalized across agents nor are real incomes equalized across regions, except in special cases.⁸

To illustrate these ideas, we have stated the model in a very simple form. We suspect that people are “more biased” towards the place they are born. To accommodate this, some formulations shift the distribution of affinity draws for the “birth location” to the right of the other locations. While this is intuitively appealing, it is practically similar to a change in moving costs in our formulation. While we assumed that moving costs are the same across people and independent of the direction of the move, this assumption is clearly incorrect in many contexts. In China, for example,

⁷In practice, these different draws are typically imagined to arise from an econometrically convenient distribution, most often an extreme value distribution.

⁸If the two regions offer identical amenities, endowments, and technologies, agents are identical absent their affinity draws, and the distribution of the differences in affinity draws is symmetrical about zero, then real incomes will be equalized and the marginal person will have equalized affinity draws. However, if the urban region has superior endowments or technologies, then we generally expect an equilibrium where real incomes are higher in the urban region and the marginal person has a greater affinity draw in the rural than urban location.

one would expect moving costs to vary on the basis of *hukou* status (whether a person is registered to live as a citizen in an urban or a rural place) and the direction of the move.

Note that without restrictions on moving costs or idiosyncratic affinities, the model has no content. If moving costs are sufficiently high, we can rationalize any observed outcome. People could be like trees: they stay where they are planted no matter how much their wages might increase if they moved over the hill. Similarly, we can always choose affinity draws such that everyone will want to stay where they are born.

We are just beginning to learn about the importance for mobility of frictions like moving costs and affinities. In a static model, Tombe and Zhu (2019) find enormous moving costs for Chinese migrants from rural farms to urban factories. Moving within a province costs migrants over half of their real income (at destination) and moving across provinces increases this to more than 90 percent. Similarly, Bryan and Morten (2019) on Indonesia argue that moving 1,000 kilometers from the place of birth costs 40 percent of real income and moving 200 kilometers costs 20 percent. Note that these two papers, like ours, rely on a static model, while “migration” is explicitly a dynamic concept. Working with models of spatial equilibrium with dynamics is difficult but is an active area of research (for example, Balboni 2019; Caliendo, Dvorkin, and Parro 2019; Ahlfeldt et al. 2020). Finally, there are other considerations of incomplete markets and risk raised by the experiments conducted in Bangladesh in the context of round trip or seasonal migration, which also find high disutility or lack of affinity from migration, as reviewed in Lagakos, Mobarak, and Waugh (2018).

Putting aside the relative lack of empirical evidence on attachment and moving costs, if we are willing to assume that moving costs are not too large and differences in affinities are limited, then we should not observe the case where both amenities and incomes are much higher in one populated place than another.

The theoretical model thus suggests that the “puzzle of early urbanization” might be rephrased as the “puzzle of too-slow urbanization.” Recent empirical work and our own results indicate wages and household incomes in developing world cities are dramatically higher than in the countryside, even after we condition on individual age, gender, and education. Moreover, the data show clearly that access to safe water, electricity, and modern sanitation improve rapidly with urbanization.

These observed patterns suggest that something is slowing down a faster pattern of urbanization that would otherwise be happening: in other words, it suggests that mobility costs and spatial attachment matter. That said, our more exhaustive accounting suggests that, while much about urban life is better than rural life, at least some things are worse, such as adult and child health outcomes and crime. Therefore, together with mobility costs and spatial attachment, if people also trade off the costs and benefits of urban life at plausible rates, current rates of urbanization in developing countries can be consistent with the spatial arbitrage that is the foundation of models of spatial equilibrium.

Conclusion

The new metropolises of the world are being built in sub-Saharan Africa, South Asia, and South East Asia. However, the mechanism that seems to have driven urbanization in much of the rest of the world—the decline of labor productivity in agriculture relative to manufacturing—may not always be at work. In some regions of the developing world, and in sub-Saharan Africa in particular, people are moving to cities when they are poorer and less productive than were their nineteenth and twentieth century counterparts in developed countries. In addition, population densities in many urban areas of South and South East Asia and sub-Saharan Africa are also much higher than what we observe in developed countries.

We also presented evidence, confirming findings of earlier research, that incomes and wages increase rapidly with density. Moreover, in spite of the “earliness” of developing world urbanization, many important aspects of life improve rapidly with density: access to electricity, safe water, modern sanitation, schooling, and inoculations for children. In seeking to understand how these factors and patterns interact, we turn to a variant of the classic Roback (1982) model of spatial equilibrium. Its basic intuition is that people will move to exploit utility differences across space. However, the benefits of urbanization seem large both economically and econometrically. Against these benefits, the costs of density seem more modest. We consider possible additions to the basic model—like costs of moving or affinities for certain traits of urban or rural areas—that might help to explain why rural-to-urban migration in developing countries is not even higher than we currently observe. Finally, our results suggest that reductions in urban crime and public health interventions that target lifestyle diseases, child health outcomes, and crime may be important tools for policymakers who would like to facilitate rural to urban migration.

■ *We are grateful to Geetika Nagpal, Vivian Liu, and Julia Lynn for excellent research assistance on this project. We thank Sebastian Kriticos and Jamila Nigmatulina for their preparation of the wage and income data which was used in the Kriticos and Henderson (2018) and Henderson, Kriticos, and Nigmatulina (2019) published papers and Cong Peng and Vivian Liu for their preparation for the DHS and Afrobarometer data used in the Henderson et al. (2020) European Union report. We thank the editors for all their very useful comments and suggestions, and additionally, Timothy Taylor for all his editorial work.*

References

- Ades, Alberto F., and Edward L. Glaeser. 1995. "Trade and Circuses: Explaining Urban Giants." *Quarterly Journal of Economics*, 110 (1): 195–227.
- Ahlfeldt, Gabriel M., Fabian Bald, Duncan Roth, and Tobias Seidel. 2020. "The Stationary Spatial Equilibrium with migration Costs." Unpublished.
- Akbar, Prottoy A., Victor Couture, Gilles Duranton, Ejaz Ghani, and Adam Storeygard. 2018. "Mobility and Congestion in Urban India." The World Bank.
- Aldeco, Lorenzo, Lint Barrage, and Matthew A. Turner. 2019. "Equilibrium Particulate Exposure." Working Paper. Brown University.
- Allcott, Hunt, and Daniel Keniston. 2017. "Dutch Disease or Agglomeration? The Local Economic Effects of Natural Resource Booms in Modern America." *The Review of Economic Studies* 85 (2): 695–731.
- Balboni, Clare Alexandra. 2019. "In Harm's Way? Infrastructure Investments and the Persistence of Coastal Cities." PhD dissertation, The London School of Economics and Political Science.
- Bolt, Jutta, Marcel Timmer, and Jan Luiten van Zanden. 2014. "GDP per Capita since 1820." In *How Was Life?: Global Well-being since 1820*, edited by Jan Luiten van Zanden, Joerg Baten, Marco Mira d'Ercole, Auke Rijpma, Conal Smith, and Marcel Timmer, 57–72. Paris: OECD Publishing.
- Bryan, Gharad, and Melanie Morten. 2019. "The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia." *Journal of Political Economy* 127 (5): 2229–68.
- Bryan, Gharad, Edward Glaeser, and Nick Tsivanidis. 2019. "Cities in the Developing World." NBER Paper 26390.
- Caliendo, Lorenzo, Maximiliano Dvorkin, and Fernando Parro. 2019. "Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock." *Econometrica* 87 (3): 741–835.
- Caselli, Francesco, and Wilbur John Coleman II. 2001. "The US Structural Transformation and Regional Convergence: A Reinterpretation." *Journal of Political Economy* 109 (3): 584–616.
- Cattaneo, Matias D., Richard K. Crump, Max H. Farrell, and Yingjie Feng. 2019. "On Binscatter." arXiv preprint arXiv:1902.09608. <https://arxiv.org/abs/1902.09608>. (accessed May 31, 2020).
- Center for International Earth Science Information Network (CIESIN), Columbia University. 2018. "Documentation for the Gridded Population of the World, Version 4 (GPWv4), Revision 11 Data Sets." NASA Socioeconomic Data and Applications Center (SEDAC). <https://sedac.ciesin.columbia.edu/data/collection/gpw-v4/documentation>. (accessed May 31, 2020).
- Chauvin, Juan Pablo, Edward Glaeser, Yueran Ma, and Kristina Tobio. 2017. "What Is Different about Urbanization in Rich and Poor Countries? Cities in Brazil, China, India and the United States." *Journal of Urban Economics* 98: 17–49.
- Combes, Pierre-Philippe, Sylvie Démurger, Shi Li, and Jianguo Wang. 2020. "Unequal Migration and Urbanisation Gains in China." *Journal of Development Economics* 142 (January 2020).
- Corbane, Christina, Aneta Florczyk, Martino Pesaresi, Panagiotis Politis, Vasileios Syrris. 2018. "GHS-BUILT R2018A - GHS built-up grid, derived from Landsat, multitemporal (1975–1990–2000–2014)." European Commission, Joint Research Centre, JRC Data Catalogue.
- Corbane, Christina, Martino Pesaresi, Thomas Kemper, Panagiotis Politis, Aneta J. Florczyk, Vasileios Syrris, Michele Melchiorri, Filip Sabo, and Pierre Soille. 2019. "Automated Global Delineation of Human Settlements from 40 Years of Landsat Satellite Data Archives." *Big Earth Data* 3 (2): 140–69.
- Davis, James C., and J. Vernon Henderson. 2003. "Evidence on the Political Economy of the Urbanization Process." *Journal of Urban Economics*. 53 (1): 98–125.
- Desmet, Klaus, and J. Vernon Henderson. 2015. "The Geography of Development within Countries." In *Handbook of Regional and Urban Economics*, vol. 5, edited by Duranton, Henderson, and Strange, 1457–1517. Amsterdam: Elsevier.
- Desmet, Klaus, and Jordan Rappaport. 2017. "The Settlement of the United States, 1800–2000: The Long Transition towards Gibrat's Law." *Journal of Urban Economics* 98: 50–68.
- Duranton, Gilles. 2016. "Determinants of City Growth in Colombia." *Papers in Regional Science* 95 (1): 101–31.
- Florczyk, Aneta J., Christina Corban, Daniele Ehrlich, Sergio Manuel Carneiro Freire, Thomas Kemper, Luca Maffenini, Michele Melchiorri et al. 2019. "GHSL Data Package 2019." Technical Report EUR 29788 EN, Publications Office of the European Union.
- Freire, Sergio, Kytt MacManus, Martino Pesaresi, Erin Doxsey-Whitfield, and Jane Mills. 2016. "Development of New Open and Free Multi-emporal Global Population Grids at 250 M Resolution."

Proceedings of the AGILE.

- Fujita, Masahisa, and Hideaki Ogawa.** 1982. "Multiple Equilibria and Structural Transition of Non-monocentric Urban Configurations." *Regional Science and Urban economics* 12 (2): 161–96.
- Edward L. Glaeser, Jed Kolko, and Albert Saiz,** 2001. "Consumer City." *Journal of Economic Geography*, 1(1): 27–50.
- Gollin, Douglas, David Lagakos, and Michael E. Waugh.** 2013. "The Agricultural Productivity Gap." *The Quarterly Journal of Economics* 129 (2): 939–93.
- Gollin, Douglas, Martina Kirchberger, and David Lagakos.** 2017. "In Search of a Spatial Equilibrium in the Developing World." NBER Working Paper 23916.
- Gollin, Douglas, Remi Jedwab, and Dietrich Vollrath.** 2016. "Urbanization with and without Industrialization." *Journal of Economic Growth* 21: 35–70.
- Heblich, Stephan, Stephen J. Redding, and Daniel M. Sturm.** 2018. "The Making of the Modern Metropolis: Evidence from London." NBER Working Paper 25047.
- Henderson, J. Vernon, Vivian Liu, Cong Peng and Adam Storeygard.** 2020. *Demographic and health outcomes by Degree of Urbanisation: Perspectives from a new classification of urban areas*. Brussels: European Commission.
- Henderson, J. Vernon, and Sebastian Kriticos.** 2018. "The Development of the African System of Cities." *Annual Review of Economics* 10: 287–314.
- Henderson, J. Vernon, Sebastian Kriticos, and Dzhamilya Nigmatulina.** 2019. "Measuring Urban Economic Density." *Journal of Urban Economics* ISSN 0094–1190.
- Ismail, Kareem.** 2010. "The Structural Manifestation of the Dutch Disease: The Case of Oil Exporting Countries." IMF Working Paper 10–103.
- Kappner, Kalle.** 2019. "'Cholera Forcing' and the Urban Water Infrastructure: Lessons from Historical Berlin." EHES Working Paper 0167.
- Lagakos, David, Ahmed Mushfiq Mobarak, and Michael E Waugh,** 2018. "The Welfare Effects of Encouraging Rural-Urban Migration." NBER Working Paper 24193.
- Lall, Somik Vinay, J. Vernon Henderson, and Anthony J. Venables.** 2017. *Africa's Cities: Opening Doors to the World*. The World Bank.
- Michaels, Guy, Ferdinand Rauch, and Stephen J. Redding.** 2012. "Urbanization and Structural Transformation." *The Quarterly Journal of Economics* 127 (2): 535–86.
- Moretti, Enrico.** 2010. "Local Labor Markets." In *Handbook of Labor Economics*, edited by Orley Ashenfelter and David Card, 1237–1313. Amsterdam: Elsevier.
- Newfarmer Richard S., John Page, and Finn Tarp.** 2018. *Industries without Smokestacks: Industrialization in Africa Reconsidered*. Oxford: Oxford University Press.
- Quintero, Luis E., and Mark Roberts.** 2018. "Explaining Spatial Variations in Productivity: Evidence from Latin America and the Caribbean." Policy Research Working Paper WPS 8560.
- Ray, Deepak K., Navin Ramankutty, Nathaniel D. Mueller, Paul C. West, and Jonathan A. Foley.** 2012. "Recent Patterns of Crop Yield Growth and Stagnation." *Nature Communications* 3 (1293).
- Roback, Jennifer.** 1982. "Wages, Rents, and the Quality of Life." *Journal of Political Economy* 90 (6): 1257–78.
- Rose, Amy, and Eddie Bright.** 2014. "The Landsan Global Population Distribution Project: Current State of the Art and Prospective Innovation." Paper presented at the Annual Meeting of the Population Association of America, Boston, MA, May 1–3.
- Sachs, Jeffrey D., and Andrew M. Warner.** 2001. "The Curse of Natural Resources." *European Economic Review* 45 (4–6): 827–38.
- Sánchez, Perdo.** 2010. "Tripling Crop Yields in Tropical Africa." *Nature Geosciences* 3: 299–300.
- Schiavina, Marcello, Sergio Freire, Kytt MacManus.** 2019. "GHS-POP R2019A - GHS population grid multitemporal (1975–1990–2000–2015)." European Commission, Joint Research Centre, JRC Data Catalogue. https://ghsl.jrc.ec.europa.eu/ghs_pop2019.php (accessed September 1, 2019).
- Tombe, Trevor, and Xiaodong Zhu.** 2019. "Trade, Migration, and Productivity: A Quantitative Analysis of China." *American Economic Review* 109 (5): 1843–72.
- Tsivanidis, Nick.** 2019. *The Aggregate and Distributional Effects of Urban Transit Infrastructure: Evidence from Bogotá's TransMilenio*. University of California, Berkeley Haas School of Business.
- United Nations.** 2018. "2018 Revision of World Urbanization Prospects." World Urbanization Prospects 2018. <https://population.un.org/wup/> (accessed September 1, 2020).
- World Bank.** 1981. *World Development Report 1981*. Washington, DC : World Bank Group.