Chapter 10

# Spatial economics for low- and middle-income countries☆

**Gharad Bryan[a], Kyra Frye[b], and Melanie Morten[b,c]**

[a]*London School of Economics, London, United Kingdom,* [b]*Stanford University, Stanford, CA, United States,* [c]*NBER, Cambridge, MA, United States*

## Contents



☆ The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. We thank, without implicating, the editors Dave Donaldson and Steve Redding, our discussant Nina Harari, as well as Pierre-Philippe Combes, Don Davis, Doug Gollin, Vernon Henderson, and Adam Storeygard for comments. We also thank audiences at UCLA, The Barcelona Institute of Economics Workshop on Urban Economics, African Econometric Society meetings, and the Princeton Handbook Conference. Special gratitude is given to the students in Econ 215 at Stanford for reading the manuscript draft. Frye acknowledges that this material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-2146755. Any errors are our own.

**Handbook of Regional and Urban Economics, Volume 6, ISSN 1574-0080**
https://doi.org/10.1016/bs.hesreg.2025.06.001

**Abstract**
Research at the intersection of development and spatial economics is increasingly important to address pressing issues in rapidly-urbanizing cities in low- and middle-income countries. This handbook chapter presents the canonical spatial model and then explores it through the lens of development economics, pointing out the "on-the-ground" facts of missing markets, frictions, and context-specific parameters that are often absent in applications of the model. We then discuss what the existing literature and the spatial model tell us about the optimal allocation of labor across space. We close by highlighting exciting possibilities for future work that integrates the spatial model with the reality of low- and middle-income countries.

**Keywords**
Spatial, Urban, Rural, Poverty

**JEL classification**
O17, O18, R14, R23, R41

## 1 Introduction

In 2017, the Ethiopian government explored a work guarantee program offering 60 days of employment to residents of its capital city, Addis Ababa. Eager to understand the potential impact, they turned to development economists who proposed an innovative research approach. The economists convinced the government to randomize program implementation across different city neighborhoods, creating a unique opportunity to study policy effects. Yet traditional randomized control trials would miss crucial nuances. Consider this: if program participants withdraw from the labor market, wages for other city workers might actually increase, including for workers who might live in control neighborhoods. These complex, indirect effects are fundamental to understanding true policy impact. Recognizing this challenge, the researchers developed a spatial framework that modeled interactions within the city, tracking how neighborhoods connect through commuting patterns. The results were striking. By accounting for neighborhood linkages through commuting, researchers discovered the program's impact was six times larger than a simple treatment-control comparison would suggest (Franklin et al., 2024). This finding underscores a critical insight for policymakers: in densely populated areas, economic spillovers can dramatically reshape policy outcomes. If researchers and governments want to understand which policies to prioritize in the presence of these spillovers, we need to bring together this type of data collection and research methodologies of development economists with the understanding of spatial interactions from spatial economics. This chapter lays out opportunities and challenges of the emerging research agenda of spatial economics in low- and middle-income countries.

This research agenda is particularly crucial given the central role that cities play in development outcomes and poverty reduction. Cities represent opportunities for growth and sustainability, and it is typically believed that density increases productivity, giving city development a key role to play in economic development. Cities may also be a safe harbor from climate change, which threatens the lives and livelihoods of many. But cities won't play this role without management. Millions more people are predicted to live in the developing world's cities by the turn of the century, and dense living comes with both the benefits of agglomeration and the cost of congestion. Understanding how to respond to growing urban populations requires answering a series of important positive and normative questions: for instance, how will infrastructure affect the distribution of populations across space, what role do credit constraints play in restricting migration to cities, or how can unsafe informal housing be efficiently and fairly converted to more productive uses?

Development and urban economists respond to these questions in different ways. Development economists see a world of market failures and prioritize cleanly identified causal estimates of the impact of policies targeting those frictions. But, while there are important exceptions, that work is often focused in rural areas, struggles to account for equilibrium effects, and, at its heart, relies on the ability to randomize multiple, non-interacting units to treatment and control, something that is hard to do within a city where everyone interacts. Urban economics, on the other hand, is often more model-based, organized around the principle of spatial equilibrium. The growth in spatial quantitative modeling over the past 10 years is testament to the success of this approach, which directly addresses equilibrium responses, enables making predictions about future policies, and gives a more holistic measure of welfare. But, despite substantial advances in recent years, these structural approaches are stylized, rarely adapted to their specific settings, and do not yet capture the array of market failures and frictions that development economists believe characterize the world.

Both approaches have costs and benefits, but answering our important questions almost surely requires more work that combines the benefits of both. For instance, failing to account for the frictions that characterize labor and housing markets in Africa's megacities is likely to lead to spatial models that mispredict the impact of productivity shocks, such as climate change, and failing to account for the in-migration that is bound to follow redevelopment of informal housing is likely to bias estimates from an RCT. Finding precise answers to these types of policy questions with research work at the intersection of development and urban economics is particularly important since much of urbanization today is driven by low-income countries. Fig. 1 shows levels of urbanization by region, using data from each country's national office and compiled by the UN as detailed in Appendix A.2 (in Section 4 we explore how these patterns change depending on the definition of "urban" used). The trend in Fig. 1 is clear and shows the standard narrative: in the past 75 years, Africa and Asia have rapidly urbanized, but in 2025 are still significantly behind the rest of the world. They are predicted to continue to urbanize very rapidly, much faster than Europe and the Americas at any point in their recent history, but they are not projected to catch up until after 2050. In the coming decades, the pressing issues facing cities in the low-income world will likely only become exacerbated as population growth and rural-to-urban migration push the urban population levels higher. There is genuine interest in both policy and academic communities in producing research that addresses these concerns.

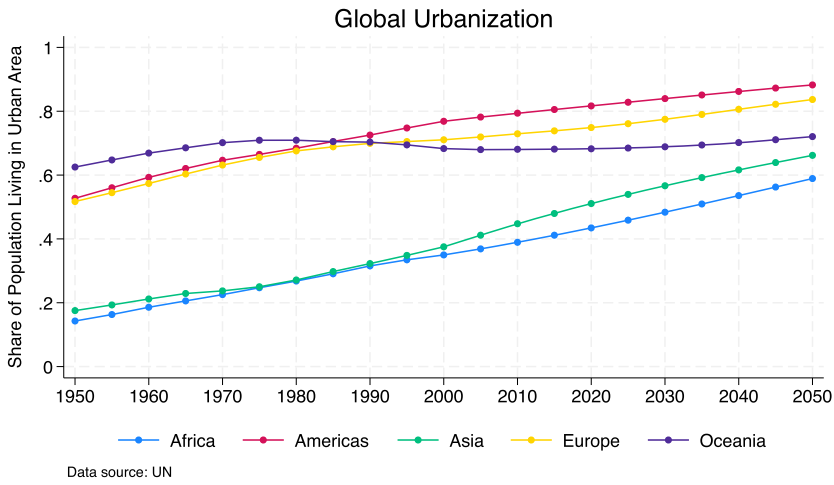


**FIGURE 1** Urbanization rate over time.

Urbanization is also linked to the economic growth and development process. Fig. 2a shows the correlation between levels of urbanization and GDP per capita. There is a strong relationship: more urbanized countries have higher incomes. The correlation implies that the two might be connected, with the potential for urbanization to lead to increased productivity (and incomes) through density-based agglomeration externalities. Urbanization is also related to structural transformation, the process of shifting labor force employment out of agriculture and into industry or services. Fig. 2b shows the relationship between the level of urbanization and the share of the labor force employed in the agricultural industry, according to ILO estimates. As people migrate to cities, they often leave behind their agricultural jobs and find employment in new industries.[1] The relationship in the figure is strong and provides evidence for the connection between urbanization and structural transformation.

[1] Note, however, that this may not always be the case in developing countries. Chris Udry's 2024 Kuznets lecture (Udry, 2024), for instance, discusses the large share of urban residents in low-income countries who also participate in agricultural activities.

The link between what the economy produces — the transition from producing agricultural goods to manufacturing and services — is often closely linked to where people live — whether in rural villages or urban centers, and the framework we present in this chapter can be adapted to include sectoral choice as well as location choice. Many of the same issues — a need for context-specific modeling and parameter choice — clearly also apply to the study of structural transformation. However, the structural transformation literature is vast and it is not our intention to review it in this handbook chapter. We refer the interested reader to a recent special issue that reviews the literature on structural change and development (Gollin and Kaboski (2023) and included papers).[2]

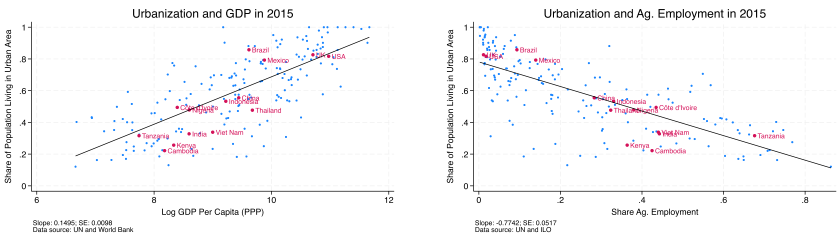


(a) GDP (b) Agriculture employment

**FIGURE 2** Urbanization rate is correlated with GDP, fewer people working in agriculture.

Although the relationships between a country's level of urbanization and its income are strong, the *within*-country heterogeneity in income is also important. For example, an RCT that subsidized poor rural workers in Bangladesh to move to the city found that their incomes went up by 30% (Bryan et al., 2014). This is striking: the individuals did not change anything other than the location where they were working. The fact that wages are higher in cities is true for many countries in the developing world, leading to the rural-urban wage gap puzzle: if there is the potential for increased earnings in cities, why don't more people move? One explanation is that migration costs might be so prohibitive that people choose not to migrate. However, data from Africa show that there are high rates of migration, including to both urban and rural areas. Fig. 3 shows the rate of migration for heads of household in several countries in Sub-Saharan Africa, with migration defined as leaving the region of birth. Rates of migration are high, including close to 50% in Malawi. We discuss this puzzle further in Section 4 and take as our starting point the spatial model.

2 The relationship between geographic space and structural transformation has been explored in the literature. For example, Fajgelbaum and Redding (2022) provide theory and evidence from historical Argentina that population density, the share of services, and the share of the urban population varies systematically even within a country based on connectivity to external trading partners. Eckert and Peters (2024) find that most urbanization in the historical United States occurred within, rather than across, counties, again highlighting the role of space in the urban development process.

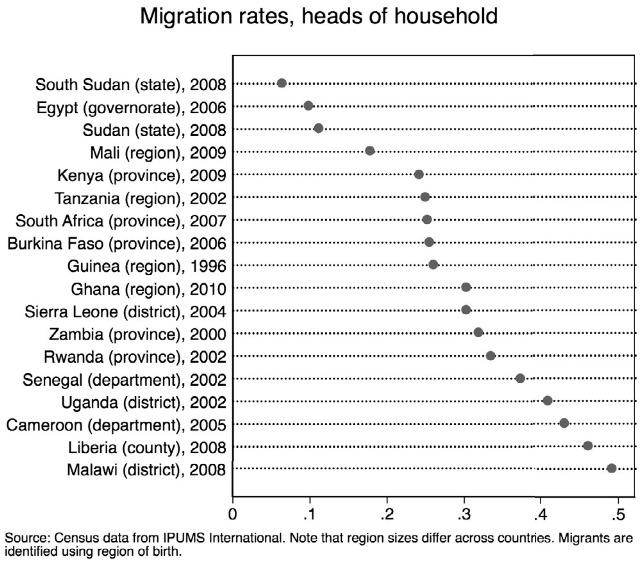


**FIGURE 3** Migration Rates in Sub-Saharan Africa.

The canonical spatial model gives a very useful framework for thinking about how to model spillovers and migration decisions; using it allows development economists to address the sort of urban problems that cannot be addressed by the standard treatment-control RCT approaches. We begin this chapter in Section 2 by laying out a stylized model of spatial equilibrium where workers make a migration decision, deciding which district to live in, and then a commuting decision, deciding where within a district to live and work. The model is a simplified version of that laid out in much more detail in Redding (2025). We use this model for several purposes. First, we discuss how the model is used, highlighting how it can resolve important problems faced by development economics. The model can accommodate spillovers, make comprehensive welfare statements, and make predictions on the impact of policy scale-up. We also highlight some exciting applications of the model in developing countries.

Next, in Section 3 we explore the canonical spatial model through the lens of development economics. We point out that the basic structure of the model is fairly frictionless. While there are congestion and agglomeration externalities, labor markets are integrated in the city, housing production responds to the arrival of migrants, commuters face little friction in traveling to work, and migrants don't like living away from home. We compare this relatively frictionless economy to the reality on the ground in developing countries. We look within each element of our basic spatial model: labor markets, housing markets, amenities, the shape of the utility function, commuting costs, and migration costs. We highlight frictions in each section and encourage urban economists working in low-income countries to understand the missing markets and elasticities that differ by context.

In Section 4 we explore the debate about the optimal allocation of people across space and offer suggestions for how the spatial model can contribute to understanding the problem. First, we provide evidence for the rural-urban income gap and ask whether it implies that people are allocated inefficiently within countries; we add a discussion on the many frictions to movement that could potentially inhibit migration. Second, we explore a classic 20th-century finding from the spatial literature: that there are fewer cities than optimal and that existing cities are too large. We discuss how this conclusion comes from a basic spatial model with no added frictions or migration costs, which may not be the empirically-relevant assumptions to make. Finally, we turn to measuring the current population allocation between rural and urban areas and revisit the common story that the developing world is currently less urbanized than the developed world. We show that this conclusion depends on the definition of "urban" used and that under some definitions the reverse is true: Asia and Africa are actually more urbanized than Europe using a definition of urbanization based on population density.

We close in Section 5 by providing suggestions for future work at the intersection of urban and development economics. We highlight three avenues that will bring the literature forward: research that incorporates the market frictions common in the developing world into spatial models, research that is able to obtain cleaner identifications for context-specific elasticities in the model, and research that finds and identifies novel data sources. For each strand, we mention existing work that has done a good job incorporating the spatial model with these important advances.

Much remains to be done, but the spatial modeling that has flourished in recent years forms a flexible framework that can incorporate the reality of low-income countries. We are excited about the possibilities for future work in this area and hope this chapter gives some inspiration to others to pursue research that uses the spatial model in context-specific settings.

## 2 Canonical spatial model

The RCT revolution has changed how development economists collect data and evaluate economic policy. But development economists have faced two key methodological questions over the past ten years: how to deal with spillovers and how to estimate effects when programs are scaled up from smaller pilots.

To understand the challenges more clearly, consider a policy change that improves amenities — perhaps paving the roads — in some informal neighborhoods in the city. The newly-paved roads could lead to migration into the updated slum. However, as people move in, rents may rise. All else equal, paying higher rents makes the net utility of the improved roads lower. A researcher evaluating the program using an RCT might naturally compare welfare of those living in the treated slums to those living in control ones. This approach could result in identifying a lower (or even zero, if the offsetting was perfect) treatment effect. The researcher might mistakenly conclude that the program did not work. However, this conclusion would be incorrect; the program did work, but the benefits were distributed across the entire region rather than concentrated solely in the treated area.

Notice that the spillover problem persists regardless of how the slums targeted by the program were selected. Randomizing the slum redevelopment program does not eliminate the spillover issue. How then can careful identification from RCTs be harnessed to evaluate larger-scale programs, especially ones that are very likely to generate spillovers? And, if there are spillovers that mean that control groups are also affected by the program, how can economists gain estimates of the aggregate effect of the program and not just the differential effect between treatment and control?

The second issue, the scale-up issue, is related. Many RCTs are piloted at a small scale, where spillovers will be minimized. However, the eventual goal may be to roll out the policy to the whole region. As a program is rolled out at scale, are there larger general equilibrium effects that may undo the benefits of the program at a smaller scale? If so, how can we model these effects so that can be predicted given the available data?[3] For example, it seems likely that encouraging 1,000 students in a city to complete high school would increase those students' wages. But what about encouraging 1,000,000 more high-school graduates? Is it possible that the wage returns to a high school education will eventually fall if the supply of high-school-educated people increases? Depending on how much wages change, a program that seemed a very good investment in partial equilibrium may no longer be as such at scale. Of course, there may be other general equilibrium effects in play as well: perhaps human capital agglomeration effects kick in and the returns actually increase, rather than decrease. The point is that it is not immediately obvious how to move from the elasticities estimated in small-scale research programs to those of a policy implemented by government partners, and one way forward is to think carefully about how to model the relationship between endogenous prices and quantities and then what elasticities would be needed to estimate these general equilibrium effects.

The canonical spatial model condenses complicated spatial and market interactions into a tractable framework, and when appropriately calibrated allows researchers to provide better answers to important questions and to deal with the issue of spillovers and scale-up.

The starting idea for the spatial model is that people choose where to live by comparing benefits and costs. For example, if a worker chooses to move from a rural village to the capital city, they must believe the benefits outweigh the costs of the move. The benefits could be broad, including a higher income or better access to cultural centers, but these come at a cost, usually higher

---

[3] We flag here the research program Yale Research Initiative on Innovation and Scale (Y-RISE), https://yrise.yale.edu/, which is tackling these important questions.

rents, perhaps more pollution, and for some an undesirable distance from family who stay at home. If our worker decides to move to the city they will enter the labor force, look for somewhere to live, and take a place for their children in a local school. All these actions will have equilibrium effects at the destination, perhaps lowering wages, increasing rents, and crowding classrooms. Or perhaps our worker bumps into a like-minded individual, opens a successful business, and local wages increase. A similar set of changes occur at the origin; perhaps our worker's departure leaves his sister with a larger farm, enabling her to adopt modern mechanized production. All these changes will then have ripple effects on others, some of whom will then make other changes to where they live and how they work. The spatial model assumes that people who would benefit move, taking into account all the equilibrium changes that would occur if they did move, and that people who will not benefit don't move.

The spatial equilibrium is pinned down by a marginal worker who is indifferent between locations. The spatial equilibrium assumption, and the models that use it, enable an analyst to capture spillovers and equilibrium effects, measure welfare, and make predictions about counterfactual events. The spatial model is the workhorse tool to study these interactions and understand where people will choose to live and how these choices change in response to shocks such as growing productivity, the building of infrastructure or the advent of climate change.

The relationships are complex, and the beauty of the spatial model is that it gives a parsimonious and tractable model that can answer important questions. Of course, this requires making simplifying assumptions that the analyst hopes are *inconsequential* for the answers the model gives. For example, in the specification we present below, there are no liquidity constraints that stop migrants from paying for a bus to the city. Assumptions are of course necessary, and much of the work of economists in developing these models should be seen as conceptual — an attempt to push modeling forward — with the results taken with a pinch of salt. But, the questions that motivate this chapter are pressing, and we need to have answers in a timely fashion. The worry is that this time pressure leads to answers based on incorrect and *consequential* assumptions. There are three broad concerns: the model may make incorrect positive or normative predictions, the model may misinterpret data leading to biased estimates, and the model may hide important areas of policy action, missing markets and market failures, giving a biased view of the actions that can alter outcomes and welfare.

In this section, we present a stripped-down version of the baseline spatial model. The goal is to focus on the mechanisms, showing how the model works and what it can achieve. After presenting the canonical model we highlight some recent work at the intersection of development and spatial economics that illustrates how the model can be operationalized. In Section 3 we then discuss each component of the spatial model through the lens of development economics to understand how the model may work for contexts with missing markets or where key elasticities are context-specific. Our goal is not to criticize the model (we are taking it directly from our own work), but to ask how it needs to be improved.

## 2.1 Baseline (static) spatial model in partial equilibrium

We start by presenting the model for migrating across space.[4] Individuals are each born in an origin $o$ and decide which destination $d$ to live in.[5] If they live in $d$ they receive indirect utility $V_d$. For now we will keep $V_d$ general, but it could include wages, the cost of living, rental costs, amenities, and other things that vary by location. Living away from home decreases the value of indirect utility. This cost of moving between the origin $o$ and destination $d$ is given by $c_{od}$. This could capture direct cost of moving, such as bus fare, but also ongoing costs such as the unpleasantness of being away from home.

To introduce heterogeneity across individuals, each individual also receives an idiosyncratic draw for each location, $\epsilon_d$, which could represent how much they like the location or how productive they are in a location. The shock $\epsilon_d$ acts to disperse people across space, whether it is a preference or productivity shock, and captures the obvious fact that we do not all agree on where is best to live. If the shock $\epsilon_d$ is interpreted as a productivity shock, it also captures the obvious fact that different people have different levels of productivity, that some people are better suited to the industry in some locations, and that we do not all earn the same amount.

Each individual chooses to live in the destination $d$ that maximizes their utility:

$$\max_d \frac{V_d}{c_{od}} \epsilon_d^i. \tag{1}$$

We would like to analyze the chosen migration decisions in Eq. (1). However, if $\epsilon_d$ is a random variable, then the maximand of Eq. (1), $\frac{V_d^*}{c_{od}^*}\epsilon_d^*$, is also a random variable, challenging our analysis. For this reason it is common to assume that the error term $\epsilon_d$ is drawn from an Extreme Value Distribution because the maximum of a shock drawn from an extreme value distribution is itself distributed extreme value (i.e., the distribution is closed under maximization), allowing us to derive closed-form solutions for the chosen migration decision.[6] A common approach is to model $\epsilon_d$ as either drawn from the Fréchet distribution (Type II extreme value) or the Gumbel distribution (Type I extreme value). In what follows, we assume that $\epsilon_d$ is distributed Fréchet with shape parameter $\theta$, but an alternative model could be constructed with a Gumbel shock, which is the log of the Fréchet. As we show below, the Fréchet gives a gravity migration equation that is linear in logs, and is thought to fit the data better.

The Fréchet distribution is given by the CDF:

$$P(X \leq x) = \exp^{-x^{-\theta}} \tag{2}$$

[4] For simplicity we shut down trade in goods. However, many papers present this model with both migration and trade present. See, for example, the review chapter Redding and Rossi-Hansberg (2017).

[5] We assume a discrete set, $N$, of location that can be either origins or destination.

[6] This approach comes from the work of McFadden (e.g., McFadden (1974) and Eaton and Kortum (2002)).

The parameter $\theta$ is the shape parameter of the distribution. This parameter is approximately proportional to the inverse of the variance of the distribution. A distribution with a large value of $\theta$ will have a small dispersion. In other words, if $\theta$ is high, everyone has a similar realization of the idiosyncratic shock and so the idiosyncratic shock is less important in determining where people choose to live: people all agree on which location is the best. If people have similar preferences then a shock to the returns of living in one place will cause many people to move as no one is very attached to NYC over Los Angeles. On the other hand, a small $\theta$ means the distribution is more disperse — some people will love the beaches of Los Angeles; others the dense city of NYC. The people who really like NYC don't want to move to Los Angeles even if average wages in Los Angeles increase and this will tend to generate less migration for the same-sized shock compared to the case where $\theta$ is larger. The size of $\theta$, often referred to as the migration elasticity, is a key empirical elasticity that needs to be estimated in order to understand migration responses.

With the Fréchet assumption, the probability of individual $i$ migrating to location $d$ and, by the law of large numbers, the share of people from origin $o$ migrating to location $d$, $\pi_{od}$, is[7]:

$$\pi_{od} = \frac{(V_d/c_{od})^\theta}{\sum_{d'} (V_{d'}/c_{od'})^\theta} = \frac{(V_d/c_{od})^\theta}{\Phi_o}. \tag{3}$$

The denominator is constant within origin $o$, so we can define it as $\Phi_o$. Taking logs of the migration probability yields the common "gravity" form:

$$\log \pi_{od} = \theta \log V_d - \theta \log c_{od} - \log \Phi_o. \tag{4}$$

This gravity equation tells us how people are distributed across space for given values of the endogenous variables ($V_d$ and $\Phi_o$).

The term $\theta \log V_d$ determines the value of living in location $d$. The larger it is, the more people will be drawn to $d$ from all origins. The strength of this effect, however, is determined by the dispersion force coming from the idiosyncratic draws. A higher $\theta$ (the Fréchet shape parameter) means that these draws are less variable, so more people agree that places with a high $V_d$ are the best places to live, and thus more people will live there relative to the case where $\theta$ is lower.

The second term, $-\theta \log c_{od}$, says that fewer people will migrate between $o$ and $d$ if the cost of migrating is larger, which makes intuitive sense. It also shows that the elasticity of migration to migration costs is given by $\theta$, for the same reason as above.[8]

---

7 The derivation is given in Appendix B.1.

8 Researchers often use a proxy for migration costs. For example, Morten and Oliveira (2024) look at the roll-out of new highways in Brazil where part of $c_{od}$ is the estimate of travel time between $o$ and $d$, $\log c_{od} = \beta \log \text{travel time}_{od}$. In that case, the estimated gravity elasticity of migration to travel time will be the product of both the migration elasticity $\theta$ and the elasticity of migration costs to e.g., travel time, $\beta$. In Section 3.6 we illustrate this method where we back out the implied $\beta$ from gravity regressions after assuming a value for $\theta$.

The final term, $\Phi_o$, has the flavor of market access. $\Phi_o$ is defined as $\sum_{d'} (V_{d'}/c_{od'})^{\theta}$, i.e., it is the sum of the cost-adjusted indirect utility option set available to someone who is born in location $o$. It summarizes how much access those born in origin $o$ have to destinations where indirect utility is high. It is possible to show that with Fréchet-distributed shocks the expected utility — accounting for both indirect utility $V_d$ and the experienced realization of the idiosyncratic shock $\epsilon_d$ — for someone born in location $o$ is proportional to $\Phi_o$, making $\Phi_o$ a natural measure of the welfare of those who are born in location $o$.[9]

Changes in $\Phi_o$ are also a straightforward way to measure the implied impacts of a change in some parameter of the model on all people in the economy. For example, suppose a new transport system within destination $d$ increases the indirect utility of living in $d$. If this effect is large enough, it will not only affect those who were already in $d$, but also those who are in other locations, and the change in $\Phi_o$ measures these impacts for each location in the economy.

The usefulness of having an endogenous object that contains all equilibrium changes — in this case $\Phi_o$ — is clear when we consider how to estimate treatment effects in the presence of spillovers. The spillover problem connects to the foundational assumption in the potential outcomes framework: the Stable Unit Treatment Value Assumption (SUTVA) (Rubin, 1980). SUTVA has two components: "no interference", meaning the treatment status of one location does not affect the outcomes in another location, and "no hidden treatments", meaning the treatment is consistently defined and implemented across all treated units. Consider the example of studying a new bus line in a city and asking whether it led to an increase in employment. Analyzing the treatment effect of living within 100 meters of the new bus line will likely face two types of SUTVA violations. If people near the bus line are more likely to work downtown, the increased labor supply could affect wages for all downtown workers, including those who don't live near the bus. This is a spillover (interference) effect. Furthermore, living within 100 meters of the bus stop is not a uniformly defined treatment: the impact of proximity to a bus line may vary depending on whether it is the only transportation option or part of a larger network. Defining treatment as living within 100 meters of the bus stop thus introduces heterogeneous (hidden) treatments. A key insight from the spatial model is that treatment can

[9] To see this, first note that under the Fréchet distribution the expectation of the shock for someone who chooses to move to location $d$ is given by $E(\epsilon_d|\text{choose d}) = \overline{\Gamma}\pi_{od}^{\frac{-1}{\theta}}$, where $\overline{\Gamma}$ is a constant related to the Gamma distribution. Therefore, the expected utility of someone who migrates to $d$ is given by $E(\frac{V_d}{c_{od}}\epsilon_d|\text{choose d}) = \overline{\Gamma}\frac{V_d}{c_{od}}\pi_{od}^{\frac{-1}{\theta}} = \overline{\Gamma}\left(\frac{V_d}{c_{od}}\right)^{\theta} = \Gamma\Phi_o$. Therefore, in the special case of the Fréchet distribution the expected utility of going to $d$ is only dependent on origin-specific factors and not destination-specific ones. The intuition for this surprising result is that a higher value of the indirect utility at the destination is exactly offset by people with lower realizations of the idiosyncratic shock moving (a negative selection effect) so the combined effect of the indirect utility and idiosyncratic term cancel out. Because the expected utility of going to $d$ is given by $\overline{\Gamma}_d\Phi_o$, which is not destination-specific, expected utility (conditional on choosing the destination) over all destinations is also equal to $\overline{\Gamma}_d\Phi_o$.

often be redefined in how much it changes market access. A unit more of market access is a consistently defined treatment. Market access solves the spillover effect by explicitly modeling how spillovers occur. Market access can also solve the scale-up problem. Because it contains all general equilibrium effects, it thus can be used to estimate aggregate, and not just partial, effects of policies, an approach pioneered by Donaldson and Hornbeck (2016).

Returning to Eq. (4), because the first and third terms are constant with destination and origin, they can be substituted for origin and destination fixed effects $\alpha_d$ and $\alpha_o$, yielding the commonly-seen regression equation:

$$\log \pi_{od} = \alpha_d - \theta \log c_{od} + \alpha_o + \eta_{od}, \tag{5}$$

where we assume an error term $\eta_{od}$ that means the model does not perfectly fit the data (perhaps those from origin $o$ have a strong connection to living in destination $d$). This simple functional form invites a straightforward regression approach to understanding the desirability of locations, and the costs of migrating across space. If we assume that migration costs are symmetric so $c_{od} = c_{do}$, then running the implied regression on data that measure the proportion of people who were born in $o$ and move to $d$ allows for recovery of the relative $V_d$ for each location in the dataset, and $c_{od}^{\theta}$ for each pair. Intuitively, within a pair $od$, $V_d$ is higher when most of the flow of migrants is toward $d$, and if this is consistent across all pairs, then $V_d$ is high relative to all pairs. Migration costs $c_{od}$ are high if people are too likely to stay home given what we have learned by the differences in $V_d$ across space. Through the lens of the model, this approach can give a clear snapshot of how would-be migrants see relatively desirability and the extent to which they are hindered in their flows between locations. These facts give a starting point to understanding where people would choose to move in response to any changes.

With a small change of notation, the same model can be used to think about commuting within a destination city. Assume that an individual lives in $l$ and works in $w$. Indirect utility becomes $V_{lw}$ (instead of indirect utility $V_d$), the idiosyncratic shock becomes $\epsilon_{lw}$ (instead of $\epsilon_d$), and $c_{lw}$ is the cost of commuting between two locations. Then, substituting the indices $lw$ for $d$ yields the same gravity equation but now in terms of commuting instead of migrating:

$$\log \pi_{lw} = \theta \log V_{lw} - \theta \log c_{lw} + \log \Phi_l.$$

If one is prepared to assume indirect utility is separable into a piece that is common to the live location and a piece that is common to the work location ($V_{lw} = V_l V_w$), then substituting in fixed effects and adding an error term yields:

$$\log \pi_{lw} = \alpha_l - \theta \log c_{lw} + \alpha_w + \eta_{lw}, \tag{6}$$

giving a commuting gravity equation that can be used in the same way as the above migration gravity equation — we can use it to learn which places people feel are best to live (a combination of market access and local amenities), which they feel are best to work, and the costs that stop them from separating their live and work locations.

## 2.2 Closing the model: baseline model in general equilibrium

The migration probability in Eq. (3) determines the distribution of people, holding constant all endogenous components of the model such as wages, rents, and the cost of living. To study equilibrium responses we need to specify what is in $V_d$ and how markets clear and prices change. It is common for indirect utility to consist of amenities ($B_d$), wages ($w_d$), rents ($r_d$), and the cost of living ($P_d$). To set ideas, assume

$$V_d = f(B_d, w_d, r_d, P_d)$$

We present simple market-clearing conditions in this section to illustrate how closing the model could work in an extreme case with perfectly competitive markets. However, the hallmark of development economics is that many markets are missing or imperfect. Section 3 studies the endogenous outcomes in turn to consider how each market functions in low-income countries.

### *2.2.1 Wages and productivity*

Assume output in a location is determined by a Cobbs-Douglas technology that combines productivity ($A$), land ($T$), and labor ($L$). Because land is a fixed factor, the parameter $\alpha$ allows for decreasing returns to scale in the variable factor $L$, and it would be standard to assume $\alpha \in [0, 1]$. Then output is given by the equation:

$$Y_d = A_d T_d^{\alpha} L_d^{1-\alpha}.$$

Of course, the production function may not be the same between urban and rural areas, urban production may be relatively free from land constraints, and even within rural or urban areas there is debate over the correct parameterization. For example, authors have interpreted the well-known negative correlation between land holding and yield as evidence of decreasing returns to land, and historically there was an argument that there was surplus labor in rural locations (Lewis, 1954), which would be consistent with either a high $\alpha$ or very large levels of labor relative to land. It is a little unclear what Lewis envisaged caused this surplus labor, or the source of the correlation between yield and plot size, and it could be a fact about the aggregate production function or a market failure. If the latter it is important to be reminded of the peril of aggregating when markets do not work well (Banerjee and Duflo, 2005). A growing literature attempts to gain greater clarity on this issue, a point we will return to below (e.g., Foster and Rosenzweig (2022); Adamopoulos et al. (2022)).

In order to study agglomeration, the literature tends to endogenize the level of productivity, $A_d$. $A_d$ comprises a fundamental level of productivity, $\overline{A}_d$, and an externality coefficient $\gamma$, likely above zero (more people increase the productivity of everyone) that measures the agglomeration benefit of having more people in a location:

$$A_d = \overline{A}_d L_d^{\gamma}.$$

Although again, the agglomeration effects of labor may plausibly differ between rural and urban areas. For example, the arguments of Boserup (1975) are sometimes interpreted as arguing that labor density above a certain level is detrimental to the transition from subsistence to modern agriculture, and hence to productivity. Additionally, it is challenging to find credible variation to estimate the actual value of the agglomeration parameter, and most existing estimates are from high-income countries. We know little about the shape of agglomeration in low-income countries.

For simplicity, we shut down trade and assume each location makes a homogeneous good that is freely traded at normalized price $P = 1$.[10] If firms pay workers their marginal product, which is not an innocuous assumption as we discuss in the next section, then first-order conditions imply that the wage is given by:

$$\begin{aligned} w_d &= (1-\alpha) A_d T_d^{\alpha} L_d^{-\alpha} \\ &= \overline{A}_d T_d^{\alpha} L_d^{\gamma-\alpha}, \end{aligned}$$

where the strength of the agglomeration parameter relative to the decreasing returns to scale in the labor parameter determines the net effect of an additional person on the wage rate.

If we assume that the idiosyncratic shock is a productivity shock, so that wages in destination $d$ for individual $i$ are $w_d \epsilon_d^i$,[11] then we can use the fact that with the Fréchet distribution the expectation of the idiosyncratic shock conditional on choosing location $d$ is $\overline{\Gamma} \pi_{od}^{\frac{-1}{\theta}}$ (see the discussion in Footnote 9), and so average expected wages for people from location $o$ working in $d$ are given by $\overline{\text{wage}}_{od} = w_d E(\epsilon_d | \text{choose d}) = \overline{\Gamma} w_d \pi_{od}^{\frac{-1}{\theta}}$, where $\overline{\Gamma}$ is a constant. This allows us to derive a wage gravity equation:

$$\log \overline{\text{wage}}_{od} = \underbrace{a}_{\log \overline{\Gamma}} + \log w_d - \frac{1}{\theta} \log \pi_{od}. \tag{7}$$

To analyze further, we can substitute in for $\pi_{od}$. This requires first specifying $V_d$. It is common to assume that $V_d$ is the product of its components ($V_d = \frac{B_d w_d}{r_d}$, accounting for the fact that high rents reduce, not increase, utility). Substituting

10 Refer to other chapters in this handbook that deal with trade. Additionally, Bryan and Morten (2019) have a model where each location produces an independent good that consumers then consume with CES preferences. As a result, the price enters the wage and there is a standard additional general equilibrium channel that as more output increases, the price of the good falls, depressing wages.

11 Bryan and Morten (2019) show it is easy to add in human capital differences at the origin, in which case wages at the destination for individual $i$ are given by $w_d \epsilon_d^i q_o$, where $q_o$ is a measure of the quality of human capital generation (schooling) at the origin $o$.

$\pi_{od} = \frac{(V_d/c_{od})^\theta}{\Phi_o}$ yields

$$\log \overline{\text{wage}}_{od} = \underbrace{a}_{\log \overline{\Gamma}} + \log r_d - \log B_d + \log c_{od} + \frac{1}{\theta} \log \Phi_o, \tag{8}$$

where $\overline{\text{wage}}_{od}$ is the average wage of those born in location $o$ who live in destination $d$. This formula has some strong implications. First, average wage within origin does not depend on destination productivity $w_d$. This is a result of two offsetting effects: places with higher productivity pay higher wages for a given person, but they also attract more people who are less well suited to the location, a negative selection effect. Second, along with classic urban models that do not feature heterogeneity, there are compensating differentials. Within origin, locations with better amenities pay lower wages. This reflects the fact that high amenity creates the same negative selection as productivity, but does not have the same direct positive effect. A similar story applies for migration costs — the higher the costs, the fewer people migrate and the more positively selected they are. Finally, across origins it is market access ($\Phi_o$) that matters for earnings.

If there are no migration costs then the market access term collapses to a constant, rather than an origin-specific term ($\Phi = \sum_{d'} V_{d'}^\theta$). This has two implications. First, without migration costs, expected utility (measured by $\Phi$) is constant across space. Second, the wage gravity equation becomes:

$$\log \overline{\text{wage}}_{od} = \underbrace{a}_{\log \overline{\Gamma} + \frac{1}{\theta} \log \Phi} + \log r_d - \log B_d. \tag{9}$$

That is, if migration costs are zero, earnings differences can only survive in equilibrium if there are compensating differentials through either rents or amenities. The mechanism here with heterogeneity is somewhat different from models without heterogeneity. Without heterogeneity, amenities are low when productivity is high because they must endogenously decrease to stop the flow of migrants in — congestion effects must offset agglomeration effects. In the model with heterogeneity the offsetting forces may occur to some extent, but there is also a selection force: a place with endogenously poor amenities will only attract those who are very productive in that location.

### 2.2.2 Amenities

As more people move into a location, amenities in a location may become congested — more traffic, more people jostling for space in the park. The congestion is modeled as an underlying value, $\overline{B}_d$, and an endogenous component determined by the externality parameter $\lambda$, likely less than zero (i.e., more people reduce the value of the amenity)

$$B_d = \overline{B}_d L_d^\lambda.$$

It is not always clear what should fall into this amenity term. For example, one might think that access to publicly-provided goods like schools and hospitals are appropriately placed in the amenity term. But, while the provision of these services are surely subject to congestion, they may well be easier to provide in a dense market where fixed costs can be spread over more people, a fact that seems to us more relevant in developing countries where rural communities are more remote and state capacity is lower.

### 2.2.3 Rents

The cost of housing is often modeled as the equilibrium in housing demand and housing supply.[12] Assuming that there is a housing supply curve with elasticity $\eta$, where $\eta > 0$, and that each person demands one unit of housing, implies that rents are a function of the endogenous labor force and potentially some baseline level of rent, $\overline{r}_d$,

$$r_d = \overline{r}_d L_d^{\eta}.$$

In this formulation, rents act isomorphically to a congestion amenity and so models do not always separate the two.

### 2.2.4 Definition of equilibrium

Given the location fundamentals, $\overline{B}_d$, $\overline{A}_d$, $\overline{r}_d$, the initial allocation of people across space, $L_{-1,d}$, and the costs of migration, $c_{od}$, for $o = \{1, ..., N\}$, $d = \{1, .., N\}$, the equilibrium allocation of people across space $L_d$ is determined by the following set of equations:

1. Labor is paid its marginal product: $w_d = \overline{A}_d T_d^{\alpha} L_d^{\gamma-\alpha}$
2. Rents are determined by housing supply elasticity: $r_d = \overline{r}_d L_d^{\eta}$
3. Congestion amenities are determined endogenously: $B_d = \overline{B_d} L_d^{\lambda}$
4. Labor supply is given by the migration rule: $L_d = \sum_o \pi_{od} L_{-1,o}$, where $\pi_{od} = \frac{(V_d/c_{od})^{\theta}}{\Phi_o}$

## 2.3 Spatial equilibrium

How does an economy adjust to spatial equilibrium? Early models assumed that individuals were identical (i.e., no idiosyncratic shocks) and migration costless, and hence equilibrium required that indirect utility was equalized everywhere across space (see, e.g., Donaldson and Hornbeck (2016)).

How would indirect utility equalize? To set ideas, consider a productivity shock in location $a$. Holding all other endogenous variables constant, the productivity shock increases wages, attracting in-migrants. As people move in, there could be upward pressure on housing rents. Higher rents reduce the utility

12. Technically, the rental price of land is pinned down between the marginal return to land as a productive use and its return as a residential use. See Redding (2025) in this handbook for a formalization of the land market in this way.

from higher wages. People continue to move until the value of living in location $a$ is the same as where they were living before. When individuals are homogeneous and there are no migration costs, this equalization occurs endogenously as migration into a location may push up earnings through agglomeration externalities, but will also push down amenities and increase rents through congestion externalities. So long as the latter effect eventually dominates there will be an equilibrium in which indeed utility is equalized across space and more than one location is populated. In this equilibrium, because everyone is identical, the equilibrium condition implies that the average migrant is indifferent to moving.

The introduction of migration costs and heterogeneity changes the equilibrium condition. With heterogeneity, the relevant utility payoff now includes the idiosyncratic return as well as the indirect utility. The idiosyncratic component of utility acts as a dispersion force. As a result, indirect utility is no longer necessarily equalized across space as in the case where all agents are homogeneous and there are no migration costs. Instead, the spatial equilibrium condition requires only the marginal, and not the average, migrant to be indifferent across locations. An increase in productivity in location $a$ has the same effect of increasing wages, but with heterogeneity the idiosyncratic utility of the migrant moving in will likely be lower than existing residents (if not, they would have already chosen to live in $a$ before the wage increase). The equilibrium condition is that the next person to move — the person whose idiosyncratic shock is such that their utility would be the same in either their former or new location — is indifferent. Note that with heterogeneity, indirect utility can differ across locations in equilibrium. However expected utility (accounting for the indirect utility and the idiosyncratic shock) will be on average the same across all destinations for people who live in the same origin (as discussed in Footnote 9). This dispersion force lessens the need for congestion to dominate agglomeration in order to have an interior equilibrium. Details can be found, for example, in Allen and Arkolakis (2014).

## 2.4 Some exciting applications of the model

Here we highlight two uses of the model from recent literature, emphasizing how much development economics has to gain from incorporating these models. We start by returning to the example of the workforce program in Addis Ababa we discussed in the introduction. We then discuss a paper that estimates the global aggregate economic effects of climate change through combining non-experimental data with the spatial model.

As mentioned in the introduction, Franklin et al. (2024) gives one example of how to combine exogenous variation in policy with a spatial model to identify both direct and indirect effects of a policy, explicitly accounting for spatial spillovers. The authors study the labor market implications of a large-scale workfare program exogenously rolled out in Addis Ababa, Ethiopia, starting in 2017. The program provided 60 days of guaranteed work for up to 4 household members in eligible households for up to 3 years and completed activities such as street sweeping and small-scale neighborhood building projects. Initial roll-out was randomized at the woreda (neighborhood) level, with 35 of 90 eligible woredas being randomly chosen to receive treatment in the first year of the program and the remainder receiving the program from the second year on. The randomized setup allows for the use of data from the first year to estimate the impact of the program on household outcomes. Franklin et al. (2024), however, show that, compared to control, treatment areas not only increased labor applied in the program, but decreased labor in the private market. They use this fact to motivate three problems with the simple treatment versus control approach. First, if the withdrawal of labor supply in the treatment areas led to an increase in wages in general, affecting people who live in control areas, then causal estimates of the program on the treated will be biased down due to a breach of the SUTVA assumption. Second, if those in control areas see an increase in wages, then those are an important welfare gain from the program that we would like to know. Finally, if the treatment affected labor supply and private wages in the first year, then when the program was rolled out across the city it would likely have even larger wage effects. What we would really like to know is the impact of the entire program once it has been rolled out, not just the small-scale randomized introduction.

Franklin et al. (2024) show how a version of the within-city aspect of the canonical spatial model presented above can solve all these problems. First, the model can be used, along with data on commuting patterns, to derive a market-access-like measure of exposure to the program for each woreda in the city. This measure allows control locations that had more commuting links to treated locations (for example, because of direct transportation links) to be more affected by the program than control locations with fewer commuting flows to treated locations. The SUTVA assumption then becomes that all locations who had the same commuting-weighted exposure to the program (rather than simple treated and control) were affected the same way by the program. This method essentially replaces the usual SUTVA assumption that treatment does not affect control with the assumption that places in the control group where commute costs are too large would not be affected. Second, the same method can be used to estimate the causal impact of the program on those woredas that were in the control group, but indirectly affected. Finally, they estimate the key parameters of the model, including local productivities, amenities, commute costs, and the Fréchet parameter and use the calibrated model to give estimates of the likely impact of the program when rolled out across the whole city. The headline result of all this is that the impact of the program is six times larger than would be thought using just the treatment and control comparison. The paper demonstrated clearly the importance of spillovers within cities that are incorporated in the canonical model, the difficulties that standard approaches to RCT data have with these spillovers, and the utility of the model for providing an alternative approach.

A second example of the value of combining microdata with spatial models is the essential practical question for development of understanding climate change. What will happen to the people living in the developing world, how important are different dimensions of adaptation and where will welfare losses be worst and the need for assistance largest? A key challenge to answering these questions is that the direct impact of climate change will vary substantially across space, and people are able to adapt, in part by moving away from those locations, and in part by changing how they live their lives and generate their incomes. The basic model presented above captures these two key aspects of climate change, and a series of papers have adapted the basic spatial framework to estimate the likely impact of climate change, with the emphasis on the developing world (e.g., Conte (2022), Cruz and Rossi-Hansberg (2024)). Cruz and Rossi-Hansberg (2024) provides perhaps the most up-to-date attempt, using a much expanded dynamic version of the canonical spatial model with endogenous technical change, a carbon cycle and costly migration to study the impact of climate change on different geographies. The results of the project show much larger expected losses for poorer countries, with losses of up to 20% of welfare in the poorer parts of Africa. The impacts of climate change are also predicted to be greatly impacted by the ability to migrate across space as a form of adaptation.

These papers demonstrate the value of the canonical model to allow researchers to use data to answer important questions relevant to development economics. They also show the adaptability of the model to different settings. One important caveat is, however, worth noting. The models are applied to developing countries, but are usually off-the-shelf versions of models that were developed for use in the world's wealthier countries. For example, both papers make use of assumptions of perfect labor markets, assumptions that migration costs will not adjust based on political economy factors, and Cruz and Rossi-Hansberg (2024) make use of parameters that are calibrated from studies in the developed world. This approach seems entirely appropriate for papers that are mostly trying to make a methodological point, but it raises some important questions if the approaches are used for real-world prediction and policy evaluation. It is important to ask whether the strong smoothness and well-functioning markets assumed in the canonical model are at least in part a source of the relatively small predicted losses from climate change, and if so whether markets, particularly in the developing world, work as the model assumes. This is what we turn to next.

## 3 The canonical spatial model through the lens of development economics

Development economics is characterized by the study of how economies operate under resource constraints and market imperfections. We want to encourage the urban economist to engage with the "on-the-ground" reality that is familiar to development economists and to ask how the parameters and modeling choices used in spatial models could be altered to better match the economic conditions in low-income countries.

Indirect utility, $V_d$, is usually modeled as a bundle of amenities $B$, wages $(w)$ and rents $(r)$,[13] for an individual who lives in location $d$, $V_d = f(B_d, w_d, r_d)$. Throughout this chapter, we argue that there are two sets of important considerations when applying spatial models to low- and middle-income settings. First, the need to estimate context-specific elasticities, and second, the need to consider which frictions may be present in the economy. This chapter steps through the main components of indirect utility — amenities, wages, housing costs — as well as commuting costs, migration costs, and the shape of the utility function itself (non-homotheticities) — to summarize some key facts that are important to model the respective markets. Throughout this section, we illustrate ideas with data from a representative low-income country using the World Bank Nigeria LSS sample from 2019 (National Bureau of Statistics (NBS), 2019).

Our view is that the spatial model itself is not restrictive, but rather, the spatial model should be developed for the context-specific environment in which it is being deployed. Of course, contextual knowledge is important whether the spatial model is employed to understand questions of rent control in the US or slum clearing in India. Our hope in discussing these issues is very much in line with Glaeser and Henderson (2017) who state, "It is time for urban economists to know as much about Dar es Salaam as about Detroit and as much as New Delhi as about New York."

Table 1 summarizes some key development facts alongside the standard modeling choice in the traditional spatial model.

## 3.1 Labor market

$$V_d = f(B_d, \mathbf{w_d}, r_d)$$

In the canonical spatial model, each location has a wage rate.[14] How the wage rate is determined can be specified by the model, but a common formulation is a fundamental level of productivity that combines with labor, human capital, and land to form output, with wages determined by marginal productivity, as outlined in Section 2. The adjustment of the wage rate as migration occurs is a key general equilibrium channel that restores spatial equilibrium. How well does the assumption of a perfect labor market map to the high levels of informality and self-employment in the low-income world?

There are several important differences in how income is earned in low- and middle-income cities. Table 2 illustrates common labor market patterns by tabulating labor force participation in the 2019 Nigeria LSS, splitting the sample

[13] The indirect utility could also include the cost of goods, $(P_d)$. Since we shut trade down in the simple model we omit this from the discussion below.

[14] The basic spatial model does not usually assume unemployment, although there may be a sectoral choice with "home production" as an option (see, e.g., Hsieh et al. (2019)).

**TABLE 1 Summary of development facts compared to spatial model.**

| Development facts | Traditional implementation of spatial model |
|---|---|
| *Wages* | |
| • Most people work informally<br>• Large share of employment in agriculture<br>• Wages may not be determined by marginal productivity | • Wages determined by production function specified<br>• Common to assume perfect labor market |
| *Housing* | |
| • Large share of people living in informal/slum housing<br>• Lack of formal property rights | • Rents determined by housing demand and supply<br>• Common to assume competitive housing market |
| *Amenities* | |
| • Urban areas have higher amenities and wages | • Spatial equilibrium implies compensating differentials |
| *Utility function* | |
| • Non-homotheticities in share of food consumption | • Homogeneous utility over food and non-food |
| *Commuting costs* | |
| • Large share of trips by foot<br>• Most common motorized transit informal minibuses<br>• Households may have budget constraints to accessing transport<br>• High congestion | • Commuting modeled as a time cost |
| *Migration costs* | |
| • Households may face credit constraints to migrating<br>• Other frictions as in Section 4 | • Migration costs modeled as an utility cost |

by rural and urban. The table shows that labor force participation is high — 74% and 78% in urban and rural, respectively. However, participation in formal employment is low — 20% of individuals in urban areas, and only 9% of individuals in rural areas. As a result, most employment occurs informally. Informal work is work that is not regular or set by a job contract — the day laborer who works for a different employer day-to-day, for a different number of days of work per week, or the small self-employed entrepreneur who sells vegetables on the side of the street. In urban areas, 47% of the population works in informal employment. In rural areas the rate is even higher: 67%. Informal employment means that average income may be highly variable over time.

The high level of informal work in the Nigerian LSS is consistent with other studies. For example, Meghir et al. (2015) find that 46% of employed, low-

**TABLE 2 Individual-level Characteristics.**

| | Urban | Rural |
|---|---|---|
| Labor Force Including Subsistence Ag. | 0.74 | 0.78 |
| Labor Force Not Including Subsistence Ag. | 0.70 | 0.58 |
| Formal Employment | 0.21 | 0.09 |
| Informal Employment Including Subsistence Ag. | 0.47 | 0.67 |
| Informal Employment Not Including Subsistence Ag. | 0.41 | 0.42 |
| Worked 7 Days Including Subsistence Ag. | 0.65 | 0.73 |
| Worked 7 Days Not Including Subsistence Ag. | 0.60 | 0.50 |
| Monthly Wage (Naira) | 51600.43 | 43030.59 |
| N | 17334 | 40471 |

*Notes:* Data source: Nigeria LSS Survey (2018-2019). Table sample is adults 18 years and older. Weighted at the household level. Variable definitions are in Appendix A.6.

education workers in Brazil are in the informal sector. The “Jobs of the World Project” Bandiera et al. (2022) (available at https://jwp.iza.org) is a massive database that harmonizes DHS and National Census datasets to measure job characteristics across locations. The database shows that the share of workers (as a share of population) in self-employment is negatively correlated with GDP. The poorest countries in the world have more than 60% of their population employed in self-employment, where richer countries have much less than 10%.

Ulyssea (2020) gives a good overview of the informality literature and emphasizes the view that informality reflects the interplay between the costs of formality, such as taxation, and the costs of informality, such as legal enforcement or lack of access to formal credit markets. Importantly, the cost of informality is thought to increase with firm size, consistent with larger firms being more likely to be formal. This seems to us to open a two-way interaction with the spatial model. First, cities in developing countries may be less dense than they appear, in the sense that they have less of the density that drives productivity. For example, Glaeser and Maré (2001) argue that the urban wage premium in the US is driven by greater increases in human capital over time, and Bobba et al. (2022) show that in Mexican data human capital increases more over time in the formal than the informal sector, suggesting that the productivity advantages of the city may be limited by informality. Relatedly, improved labor matching is thought to be one of the benefits of urban areas, but the existence of large numbers of informal workers may congest labor market search. For example, Meghir et al. (2015) use a structural model along with Brazilian data to argue that enforcement that leads to the closure of informal firms can lead to better matching between high skilled firms and workers, and hence greater productivity in general. Second, increasing effective density by helping people search for jobs may help reduce informality, leading to an additional productivity multiplier. Evidence consistent with the idea can be found in Zárate (2023), who shows how improved transport in Mexico City led to a decrease in informality. Overall, these effects caution against taking agglomeration parameters from other settings and applying them to developing country cities with large amounts of informality.

How wages are determined in rural, low-income, or highly informal markets is also important to consider. The assumption of perfect labor markets assumes that workers are paid their marginal product. The assumption of a concave production function implies that as labor leaves a location wages would increase, helping move toward wage convergence. However, a very large development literature, including seminal work such as Lewis (1954), has considered whether there is surplus labor in village economies and, indeed, whether a reduction in labor supply would increase the wage. Breza et al. (2021) empirically study whether reductions in labor supply result in increases in wages in the context of the daily agricultural day laborer market in village India. The authors exogenously generate a negative labor supply shock in the village by recruiting up to 25% of the workers who live in the village for a job that takes place outside the village. The recruited workers are thus removed from the village labor market. The field experiment was run throughout the year, and so the paper can study the impact of the labor supply shock on the village during different agricultural periods with more or less local availability of work. The paper finds important heterogeneity in the impact of the negative labor supply shock across the year. In the lean season, when work opportunities are scarce, the negative labor supply shock does not cause wage changes; removing 20% of the labor force does not increase the wage in the village and does not decrease the total employment in the village. Instead, the experiment has a positive spillover to the non-treated individuals in the village, who are more likely to be employed.[15] However, when the authors run the same experiment during the (shoulder) peak season they find the opposite effect. In this case, the negative labor shock increases the village wage and decreases the number of people employed in the village in total, as one would expect with a competitive labor market. This paper illustrates that, even within the same market, there may be periods of the year where wages adjust to a reduction in labor supply as a standard model would predict, but also periods of the year where the wage is not adjusting as expected. However, other literature has found evidence of wage adjustment in rural areas more consistent with what standard model would predict when labor supply reduces (e.g., Imbert et al. (2022) in China). The heterogeneity, and need to consider the sector of employment, reinforces that economists should carefully consider the assumptions of the spatial model and modify them as needed to fit the context in which they are working.

### 3.2 Housing market

$$V_d = f(B_d, w_d, \mathbf{r_d})$$

[15] The paper is very careful to consider and rule out other alternatives such as positive selection of the experimental sample (leaving behind negatively selected individuals in the village).

In a standard spatial model, the cost of housing is often modeled as an equilibrium outcome of a housing demand curve (driven by population) and a housing supply curve, potentially determined by the cost of construction, the interest rate, and available land as in Diamond (2016). In this setting, rents adjust as population increases and therefore is a mechanism for utility to converge across space.

How well do these assumptions map to the housing market in low-income countries? Again, there are important context-specific considerations. First, a large share of housing in low-income countries is slum housing. The UN-Habitat Urban Indicators meeting in 2002 defined a slum as "a contiguous settlement where the inhabitants are characterized as having inadequate housing and basic services. A slum is often not recognized and addressed by the public authorities as an integral or equal part of the city. It is an area which combines to various extents the following characteristics: insecure residential status, inadequate access to safe water, inadequate access to sanitation and other infrastructure, poor structural quality of housing, overcrowding" UN-Habitat (2002).[16] Slum housing is primarily a measure of the quality of housing. Table 3 reports the share of urban population living in slums in 2022, according to the UN Habitat Urban Indicators Database.[17] As the table shows, globally 25% of the urban population lives in slum housing. The share living in slums is much higher in the poorer regions of the world — 54% in Sub-Saharan African, 43% in Central/Southern Asia — than in Latin American/Caribbean (17%) or Europe/North America (0.7%).

The Nigeria data are consistent with these facts. Table 4 shows characteristics of households, split across rural and urban areas. On average, households in rural areas have slightly larger households, with 5.4 members (compared to 4.5 in urban areas). 39% of urban households live in a slum. Applying the same slum criteria to rural households, we find over 75% of rural households live in inferior living conditions.

Individuals can either own or rent their house. In the Nigeria LSS data, 33% of urban households (and 68% or rural households) own their own residence. It is important to emphasize that although they live in informal housing, most slum households still pay rent and participate in a housing market. In the Nigeria case, households spend just over 5% of their consumption on rent (accounting for imputed rent for those who own).[18] However, housing markets for slums may function differently than the expected perfectly competitive market. For instance, Marx et al. (2019), studying slums in Nairobi, Kenya, finds that when

---

[16] A distinct but related issue is whether inhabitants have formal property rights. The current formal definition of slums used by the UN-Habitat Urban Indicators Database does not include insecure residential status as criteria, and so a slum resident is a resident experiencing at least one of the other four deprivations.

[17] May 2024 version, https://data.unhabitat.org/pages/housing-slums-and-informal-settlements, downloaded 8/21/2024.

[18] This number is lower than the data we tabulate later on household budget shares in Table 6, where we find renters paying between 9-24% across a range of countries.

**TABLE 3 Share of urban population living in slums, 2022.**

| Region | Share |
|---|---|
| Global | 0.25 |
| Sub-Saharan Africa | 0.54 |
| Central Asia and Southern Asia | 0.43 |
| Eastern Asia and South-Eastern Asia | 0.25 |
| Western Asia and Northern Africa | 0.18 |
| Latin America and the Caribbean | 0.17 |
| Northern America and Europe | 0.007 |

*Notes:* Data from UN-Habitat Urban Indicators Database. Population weighted average produced by UNHabitat using reported/estimated national data points. Proportion of urban population living in slums or informal settlements. The estimates are based on the global methodology on household deprivations where the inhabitants suffer one or more of the following 'household deprivations': lack of access to improved water services, lack of access to improved sanitation facilities, lack of sufficient living area (each room is shared by no more than 3 individuals), and lack of housing durability. Underlying data for the slum/informal settlements components of the indicator is computed from censuses and national household surveys such as the Demographic and Health Surveys (DHS) and Multiple Indicator Cluster Surveys (MICS).

**TABLE 4 Household-level Characteristics.**

| | Urban | Rural |
|---|---|---|
| No. in HH | 4.53 | 5.44 |
| Slum | 0.39 | 0.77 |
| Own Residence | 0.33 | 0.68 |
| Imputed Monthly Rent (Naira) | 7137.26 | 3057.49 |
| Share Consumption on Food | 0.55 | 0.64 |
| Share Consumption on Rent | 0.07 | 0.05 |
| N | 6808 | 15302 |

*Notes:* Data source: Nigeria LSS Survey (2018-2019). Weighted at the household level. Variable definitions are in Appendix A.6.

a landlord and a local chief share the same tribal ethnicity, tenants pay higher rents. When tenants and the local chief share the same tribal ethnicity, tenants pay lower rents. These political economy considerations for rental price are outside the standard model. Elasticities may also differ in the presence of slums. Guedes et al. (2023) finds that slums increase the supply elasticity of housing in Brazil, since slums can develop without land use regulation constraints. Alves (2021) also finds different housing elasticities in Brazil for housing with basic services (sanitation and clean water) versus those without basic services. He finds that the elasticity of rent to demand is 0.37 for non-slum housing and 0.07 for slum housing.

In many cases, slum housing overlaps with a lack of formal property rights. The importance of property rights has been a fundamental consideration in economics, linking back to de Soto (De Soto, 2000). Early non-experimental work looked at natural situations where governments had changed housing conditions to understand how household investment changed. For example, Field (2007) finds that a large-scale urban land titling project in Peru had an indirect effect of increasing labor supply in slums by 10–15% as people felt that could work instead of being physically present to protect their homes. The lack of tenure security, and the subsequent requirement for guard labor, could be considered as part of commuting costs suggesting a need to model both markets jointly. Galiani and Schargrodsky (2010) find property rights led to improved investment in housing and education, as might be expected with lower expropriation risk. However, other studies have found smaller effects — Panman and Lozano Gracia (2022) highlights several studies where attempting to formalize land title backfired. One explanation is that in many areas locations often have complex systems of informal land titles that substitute for formal property rights (see, e.g., Bird and Venables (2020)). The overarching lesson here is that context is important: the absence of property rights may indicate a missing market, or there may be a second-best informal market that effectively substitutes for the formal one.

Another open question is whether slums are a temporary — a stop along the way to “modernization” — or a permanent phenomenon. Henderson et al. (2021b) present a dynamic model where rising house prices endogenously incentivize conversion of slum areas to formal-sector use.[19] However, empirically, evidence is more mixed. Marx et al. (2013) in their review paper find that residents in Kibera, a large slum on the outskirts of Nairobi, Kenya, had lived in Kibera for 16 years, suggesting a limited role for mobility out of the slum. For the neighborhood as a whole, only small improvements occurred over time: pit latrine use only fell 5 percentage points (from 82% to 77%) over the 10-year period 1999 to 2009 and the number of rooms per capita stayed fairly constant (0.68 in 1999 to 0.67 in 2009). Case studies in United Nations Human Settlements Programme (2003) also find long tenure rates — for example, 40% of slum residents in Kolkata have been slum dwellers for two generations or longer. The question of how, and whether, slum housing will convert to formal housing is an important policy debate and an active research area.

### 3.3 Amenities

$$V_d = f(\boldsymbol{B_d}, w_d, r_d)$$

19 See also, e.g., Frankenhoff (1967); Bank (2009).

As already noted, the simple spatial model outlined in Section 2 implies that absent migration costs, average wages will equalize across space unless there are compensating differentials: high wages imply low amenities. Do compensating differentials hold for low-income cities, or are both wages and amenities higher in cities than in rural areas?

Gollin et al. (2021) study this question in 20 African countries. They first show that earnings do indeed increase with density in their sample, and then focus on public services and environmental measures, such as access to electricity and water, air pollution, and crime as measures of amenities. The authors put together data from the DHS, Afrobarometer surveys, and geo-referenced pollution data. Table 5 pulls some representative results to show the relationships between measures of amenities and population-density quartiles. The key result is that for each of the amenity measures considered, with the exception of crime (where crime increases with density) and pollution (where there is little relationship between pollution and density),[20] there is a positive relationship between amenity and population density: amenities are higher in high-density, urban areas. For example, in the lowest-quartile density locations, 39% of areas on average have access to the electricity grid, but in the highest-quartile density locations, 72% of areas have access.

**TABLE 5 Private consumption, Public goods, Crime, and Air Pollution across density quartiles.**

| | Q1 | Q2 | Q3 | Q4 |
|---|---|---|---|---|
| *Private consumption* | | | | |
| Telephone | 0.41 | 0.49 | 0.60 | 0.83 |
| Finished roof | 0.41 | 0.5 | 0.67 | 0.88 |
| Child stunted (low height for age) | 0.4 | 0.4 | 0.38 | 0.29 |
| *Public goods* | | | | |
| Electricity grid | 0.39 | 0.42 | 0.48 | 0.72 |
| Health clinic | 0.59 | 0.58 | 0.62 | 0.73 |
| *Crime* | | | | |
| Property crime | 0.28 | 0.31 | 0.31 | 0.33 |
| Feel unsafe | 0.37 | 0.39 | 0.38 | 0.45 |
| *Air pollution* | | | | |
| PM2.5 | 19.45 | 20.24 | 18.55 | 18.15 |
| PM2.5 (removing dust and sea salt) | 5.84 | 5.81 | 5.79 | 5.84 |

*Notes:* Data compiled from tables in Gollin et al. (2021). The quartiles of population density are indicated by Q1-Q4, where Q1 is the lowest-density quartile and Q4 is the highest-density quartile.

[20] The pollution result may be specific to the sample of African countries studied. Gollin et al. (2021) note that they find strong pollution gradients with density in China, India, and the United States.

Henderson and Turner (2020) undertake a similar exercise as Gollin et al. (2021), showing how measures of amenities correlate with density, using data from Africa as well as Latin America, South-East Asia, and South Asia. Again the same pattern holds — locations that are denser have higher incomes, but also appear to have higher measures of amenities, with the possible exception of crime.

Is this result inconsistent with the spatial model? Eq. (7) above suggests three explanations for wage gaps: amenity (including rent) differences, migration costs, and origin-specific market access (which could also broadly include origin access to human capital as in Bryan and Morten (2019)). If amenities really increase with density then that would imply large migration costs or very strong differences in access to schooling must be stopping the inflow of migrants. We discuss this issue further in Section 4, but for the moment it is worth noting that it may well be easier to provide public goods in denser environments. For example, grid electricity has been shown to be subject to decreasing average costs (Lee et al., 2020), so its per unit cost will surely be lower in a dense environment, a fact that is probably more relevant in poor countries with tighter budget constraints in supplying public goods.

It is also worth noting that amenities may differ significantly within a city. There is already an established literature on neighborhood segregation in developed countries (e.g. Cutler and Glaeser (1997)), and an emerging one in the developing world. For instance, Asher et al. (2024) use data from India to show that public service provision is significantly lower in neighborhoods where marginalized groups — Muslims and Dalits (scheduled castes) — live. This difference is present across several types of public goods, including schools, health clinics, and water and electricity infrastructure. In addition to the differences in rent prices noted before, Marx et al. (2013), studying slums in Nairobi, Kenya, also find that investment in housing infrastructure changes depending on the ethnic similarities of the landlord, tenant, and local tribal chief. These effects imply that we shouldn't think of amenities as constant within a city, and instead be aware of context-specific heterogeneity to incorporate into the standard spatial model. The patterns could be consistent with the argument of Feler and Henderson (2011) that urban governments deliberately try to dissuade in-migration by restricting the supply of public goods. While in actuality a difference in amenity access, this would show up in our spatial model as a migration friction.

## 3.4 Non-homotheticities

The utility function itself may also differ between low- and high-income locations, or between low- and high-income people in the same location. For example, one relevant empirical fact is that in many poor countries people spend more than half of their income on food, limiting the amount of resources to be spent on other goods and services. As a result, it may be important to consider whether households have non-homothetic preferences and if so, whether this non-homotheticity is quantitatively important.

**TABLE 6 Household level expenditure shares across countries.**

| Share of Total Expenditure | Colombia (2010) | | Kenya (2005-6) | | Tanzania (2007) | | Uganda (2009) | | United States (2010) | |
|---|---|---|---|---|---|---|---|---|---|---|
| | (1) Below median | (2) Above median | (3) Below median | (4) Above median | (5) Below median | (6) Above median | (7) Below median | (8) Above median | (9) Below median | (10) Above median |
| Food | 0.37 | 0.25 | 0.51 | 0.33 | 0.50 | 0.41 | 0.48 | 0.39 | 0.16 | 0.10 |
| Housing | 0.24 | 0.19 | 0.09 | 0.12 | 0.12 | 0.11 | 0.17 | 0.14 | 0.25 | 0.23 |
| Transport | 0.09 | 0.08 | 0.02 | 0.07 | 0.06 | 0.07 | 0.04 | 0.05 | 0.10 | 0.09 |
| Other | 0.26 | 0.40 | 0.32 | 0.42 | 0.27 | 0.33 | 0.27 | 0.29 | 0.44 | 0.54 |
| GDP per capita | 12,639.40 | 12,639.40 | 3,227.81 | 3,227.81 | 1,925.62 | 1,925.62 | 1,892.77 | 1,892.77 | 56,693.22 | 56,693.22 |
| Median p.c. expenditure (day) | 10.32 | 32.20 | 5.78 | 16.62 | 3.25 | 7.41 | 4.36 | 9.88 | 42.57 | 112.87 |

*Notes:* The expenditure shares and total expenditure per capita per day were calculated using urban observations only. Expenditures made over a week, month or quarter were converted to annual expenditure. The housing category does not include housing expenditures other than rent for renters and imputed rent for owners and total expenditure includes imputed rent for owners. Households are considered to be renters only if they are renting privately and we have dropped observations that did not have rent data for renters. Data source: Colombia National Quality of Life Survey 2010, Kenya Integrated Household Budget Survey 2005/6, Tanzania Household Budget Survey 2007, Uganda National Household Survey 2009, and the U.S. BLS Consumer Expenditure Survey 2010 (using data from the Interview Survey only which accounts for 95% of total household expenditure). Numerical values are in 2019 $PPP.

Table 6 reports household expenditure shares across urban areas in five countries. Households with below-median consumption consume a larger share of their income on food than households who are richer — for example, in Colombia, 37% of consumption is on food for below-median households, compared with 25% for above-median households. In the three African countries in the table (Kenya, Uganda, and Tanzania) the below-median share on food is closer to 50%. Depending on the magnitude of the non-homotheticity in food, there may or may not be scope for additional non-homotheticities: an adding-up constraint may mean that if the non-homotheticity on food is large enough, non-homotheticity on other goods is smaller. So, while it is the case that in Colombia the poor spend more on both food and housing, in Kenya, the poor spend a smaller share of their income on housing than the rich (12% vs 9%).

Again, the lesson here is that specifying the correct utility function for the context studied is important so that the household consumption response will be correctly specified, especially if there is an interest in analyzing the distributional effects of a policy. A household that spends half its income on food is unlikely to have a large income share going to transport, and so may be less affected by any change in the cost of commuting, whereas a richer household may be more elastic to commuting cost changes because their basic needs are satisfied. More generally, non-homothetic preferences can play an important role in explaining larger patterns of structural transformation as households become rich enough to demand more than just food (Comin et al., 2021).

## 3.5 Commuting costs

One potential benefit of urban areas is density: thick labor markets lead to high worker-firm matches, a large pool of educated people lead to many people to learn from, and a large city generates a large market for goods. However, density may not deliver these benefits if congestion or lack of access, whether due to infrastructure or financial barriers, make it difficult for people to move around the city. The level of "effective density", accounting for the density people can access, may therefore be substantially lower, especially in low-income countries. As a result, the challenges facing transportation in low-income countries, especially in the context of rapidly-increasing population, often center around how to improve access — whether physical or financial — to transport. In the spatial model, lack of access and resulting low levels of commuting will be interpreted by the model as high commuting costs and a high elasticity of commuting costs to distance. The elasticity of commuting costs may therefore be highly context-specific.

How do people move around cities? One important fact about commuting in low-income countries, especially cities in Africa, is that a large share of trips are made by foot and informal transportation. Table 7 tabulates the share of trips made by different modes across 14 African Cities, replicated from Kumar and Barrett (2008). Walking is the most common mode of transport, at close to 40%

**TABLE 7** Share of trips by mode, 2008.

| Mode of transport | Share of trips |
|---|---|
| Walk | 37% |
| Minibus | 30% |
| Motorcycle | 12% |
| Private car | 12% |
| Taxi | 8% |
| Large bus | 7% |
| Other | 4% |

*Notes:* Data tabbed from Kumar and Barrett (2008). Data measured for 14 African Cities.

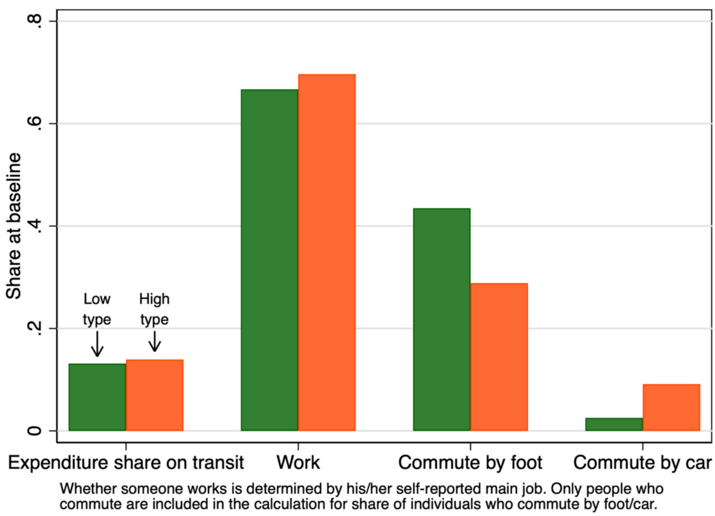


**FIGURE 4** Low-income households less likely to commute, walk more. Data from Dar es Salaam (Balboni et al., 2020).

of trips. The second-most common form of transport is the ubiquitous minibus (known as a "daladala" in Dar es Salaam, "matatu" in Kenya, and "tro-tro" in Ghana), a shared form of transit that carries up to 15 passengers. Motorcyles — where passengers sit behind a driver — are the next most common. More formalized forms of transport such as buses, cars, or taxis make up a much smaller share of transport. The large share of trips made by foot is also consistent with household data collected in Dar es Salaam, Tanzania by Balboni et al. (2020). The data from Dar es Salaam in Fig. 4 plots the share of people who work, commute by foot, commute by car, and the share of household consumption spent on transportation. The sample is split by above/below median consumption. The striking result is the very high share of low-income households that walk to work: conditional on those who work outside the house, 45% of low-income households (and 35% of high-income ones) commute by foot.[21]

One explanation for the high rates of walking is that few people in low-income countries can afford cars, and public transportation is expensive. Fig. 4 shows that households spend approximately 12% of their household budget on transport. Data sourced from household budget surveys, again for 14 major African cities, show similar levels of household expenditure. A daily round-trip commute on public transport (usually informal minibus) would be expected to cost between 5.1–27.5% of a household's consumption budget, with an average of 8%. For households in the poorest fifth of the income distribution, it would cost between 19–53% of their total household budget to pay for a daily round trip (Kumar and Barrett, 2008).[22] The high cost of public transport may limit the benefits of dense urban areas, particularly to the poor. A broader implication is that if people are not able to access the wider city, then the agglomeration benefits may be smaller. This may be one explanation for the fact that, despite high levels of density, many highly-urbanized low-income countries do not seem to be receiving the full benefit of urbanization.

However, even when households have access to transport, roads may be congested, increasing commuting costs. Akbar et al. (2023a) and Akbar et al. (2023b) study travel times across major cities in India and across countries. They find that there are large differences in travel speed across cities within India, and very large differences in travel speeds across countries. Both gaps — the within country and across country — are correlated with economic development. The authors find evidence that it is not congestion but rather that high commuting costs in this case may be more likely due to too few or not-good-enough roads, suggesting a potential role for infrastructure investment. Indeed, the World Bank is putting massive investments into improving transportation infrastructure. In 2019, transportation is the largest World Bank sector for lending, and represents at least 18% of its net commitments (Akbar et al., 2023a).

## 3.6 Migration costs

Recall that the spatial model implies that migration between origin $o$ and $d$ is given by:

$$\pi_{od} = \frac{(V_d/c_{od})^{\theta}}{\Phi_o}$$

Understanding the size of the migration cost $c_{od}$ is important for getting quantitative predictions from the canonical model, but it is also important to

[21] Similar rates of traveling by foot are found in other African cities: for example, 30–45% of trips (not just those for work) in Nairobi, Lagos, and Addis Ababa, and nearly 70% of all trips in Ugandan cities and Dar es Salaam occur by foot (Lall et al., 2017).

[22] This high cost of public transit is consistent with World Bank travel demand survey that found that the poorest households spend close to 40% of their household budget on public transport across Nigerian cities Lall et al. (2017).

understand the sources of the cost because different sources of cost will have different positive, normative and policy implications. We start here by comparing how migration costs correlate with distance across different countries but discuss migration frictions further in Section 4.

We start by taking the simple model seriously and applying it to cross-country data. We use a sample of countries in the IPUMS database where we have data on sub-region of birth and current sub-region of residence, defining migration as living away from where you were born. More information on the data is presented in Appendix A.1.

For each country, we calculate an implied migration cost between each pair of regions. The model shows how to do this.

If we assume that migration costs are symmetric, so $c_{od} = c_{do}$, and it is costless to stay at home, so $c_{oo} = c_{dd} = 1$, then we can take ratio of the migration probability from $o$ to $d$ with the migration probability from $d$ to $o$ and rearrange for $c_{od}$[23]:

$$c_{od} = \left( \frac{\pi_{od}}{\pi_{oo}} \times \frac{\pi_{do}}{\pi_{dd}} \right)^{\frac{-1}{2\theta}} .$$

Intuitively, there is evidence of movement costs if too few people move away from home, and this is symmetric so that it cannot be explained by differences in inherent productivity or amenity between location $d$ and $o$. Translating this effect into utils requires knowing the migration elasticity, which is given by $\theta$. We refer to this as a modified Head-Ries index given its similarity to the index commonly used to estimate trade costs in gravity models (Head and Mayer, 2013).

Bryan and Morten (2019) show that, if the idiosyncratic shock, $\epsilon$, is interpreted as a productivity shock, then $\theta$ can be estimated from the elasticity of average wages to the proportion of movers from a location, and apply this method to data from Indonesia and the US. The estimates show relatively large costs of moving across space, which are higher in the poorer country. In particular, Indonesian migrants who cross a stated boundary must be compensated with 39% higher earnings, while those in the US require only 15% higher earnings. While data on migration between districts to estimate $\pi_{od}$ is available for a range of countries, data on wages is less available in those same datasets. Hence, to compare migration costs across a broader sample of countries we assume a constant $\theta$ across countries, which we set equal to 3, close to the estimate from Bryan and Morten (2019). This method will allow for meaningful ranking across countries so long as $\theta$ does not strongly correlate with migration costs. Under this assumption, we calculate migration costs using the IPUMS data on share of migrants to and from each sub-region pair.

With migration costs in hand, we estimate the elasticity of cost to distance

$$\log c_{od} = \beta_0 + \beta \log dist_{od} + \epsilon_{od}.$$

[23] The derivation is given in Appendix B.2.

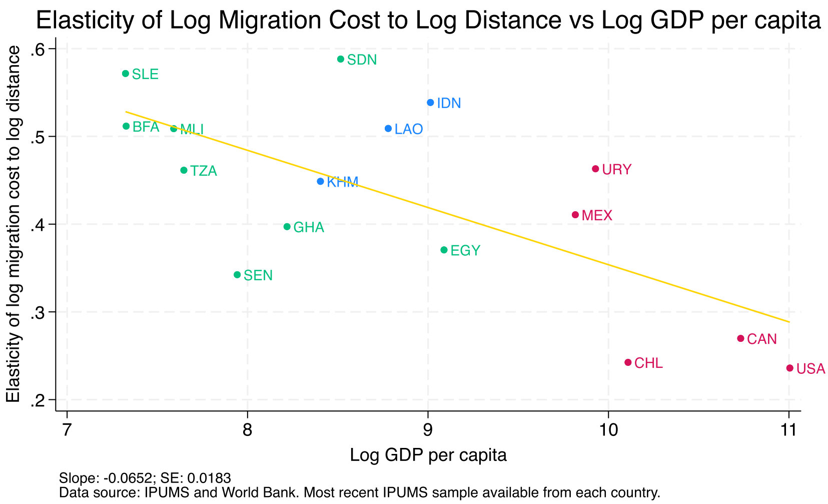


**FIGURE 5** Richer countries have lower migration costs.

A larger $\beta$ implies that farther regions have higher migration costs, while a low $\beta$ implies that there isn't a strong relationship between cost and distance. Fig. 5 plots these $\beta$'s against the log GDP per capita (PPP) for each country in our IPUMS data. The results clearly show that poorer countries have migration costs that increase more with distance than richer countries. In other words, poorer countries have larger implied migration costs to travel the same distance than richer countries.

The result suggests that migration costs are higher in poorer countries, but there are two limitations. First, the measure is built on a specific model of selection, and alternative models may interpret the evidence differently. Second, it is not clear the source of the costs, making the interpretation unclear, as noted above. We return to understanding the components of migration costs in Section 4.1.

## 4 Is labor allocation efficient across space?

In this section, we ask what the spatial model and related literature have concluded about whether the allocation of people across space is efficient. We first consider the rural-urban wage gap. Do the large wage gaps we observe across space reflect an exciting potential gain to increase incomes through better matching of workers to locations with high incomes? Or do they reflect returns to heterogeneous factors, such as differential human capital? If there are high migration costs, what is likely the underlying source of these costs? Second, we discuss a related literature in urban economics that asks whether cities are too large, a question that has obvious policy effects for the megacities of the low-income world. Third, we discuss how to measure the population allocation

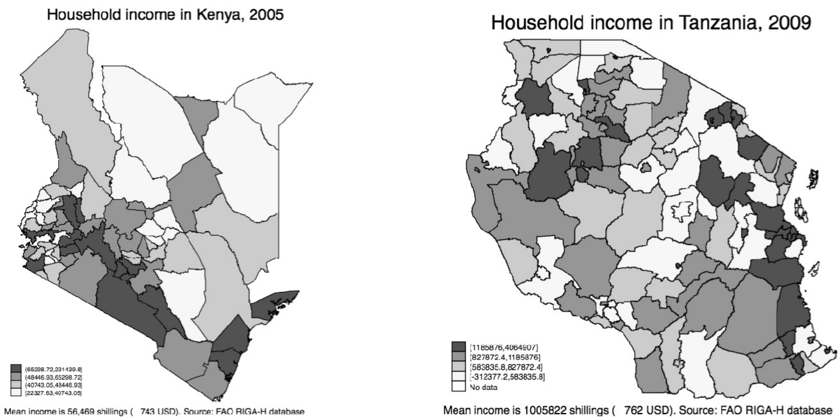


**FIGURE 6** Spatial Income Heterogeneity in Kenya and Tanzania.

across space and between rural and urban areas. We show that whether low-income countries *actually* have lower levels of urbanization than high-income countries depends on the definition of "urban" and the underlying data used.

## 4.1 What explains the rural-urban wage gap?

A key question for policymakers, and academics, is whether the current population allocation is efficient and, if not, if people should be encouraged to move to "better" locations. In this subsection, we first lay out some of the evidence for the rural-urban wage gap. We then review the literature on possible explanations for the gap that would imply that it is efficient: that in equilibrium wages should indeed be lower in rural areas. The evidence is inconclusive but it appears likely that there would be significant welfare and development gains to migration, i.e. that the population allocation in many developing countries is inefficient. Finally, we discuss many of the frictions that could prevent migration to cities, particularly in the context of low-income countries. These are frictions that exist on the ground but do not show up in the standard spatial model and reinforce the importance of adapting the canonical framework to specific contexts.

Evidence from many countries shows that there is significant heterogeneity in income within a country, with wages usually substantially higher in urban areas. Fig. 6 shows the average household income by county (Kenya) or district (Tanzania). The maps show a wide range of average incomes within each country, with households in some regions making significantly more than others. The regions containing the capital cities (Nairobi and Dar es Salaam, respectively) are in the highest bracket. Fig. 7 shows the distribution of log average monthly wage in Indonesia. Wages in urban areas are, on average, higher than in rural areas. In the Indonesia data it is also clear that there is a good deal of overlap in the distributions — wealthy people in rural areas earn a higher wage than

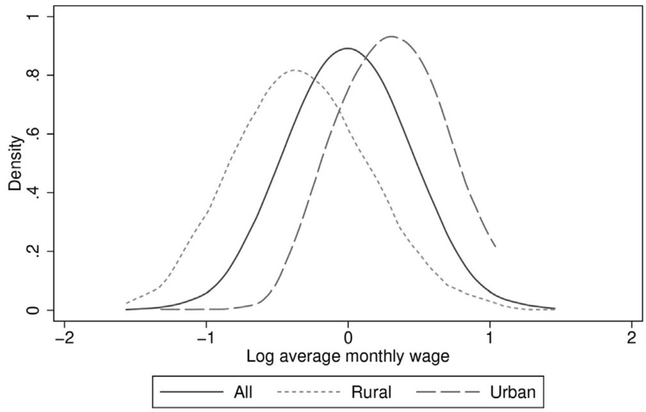


**FIGURE 7** Indonesia: the distribution of wages in urban areas is shifted to the right. Distribution of wages: regency level. Log average monthly wage is demeaned of year fixed effects. Unit of observation is the regency. Regency is defined as either rural or urban to match the national share of rural. Sources: 1995 SUPAS, 2011 SUSENAS, 2012 SUSENAS (Bryan and Morten, 2019).

poorer people in urban areas. Young (2013) documents a related fact using DHS data from 65 countries: on average, those living in urban areas consume the same amount as someone living in a rural area who has 9 more years of education, which implies a large potential gain to migrating to cities. More generally, Gollin et al. (2014) show that value added per effective worker in agriculture is about half what it is in other sectors, and that this number drops to one quarter for the 10 percent of poorest countries in the world.

These rural-urban gaps could be incongruent with the static spatial model presented in Section 2. In that model, with skills drawn from a Fréchet distribution, if there are no amenity differences, no differences in the availability of schooling across origins, and no migration costs, then earnings of workers should be equalized across space.[24] However, the model assumes that human capital has been adequately controlled for. In the model we presented, the only potential human capital differences were through the modeled idiosyncratic shock. If human capital is not adequately captured by the model then differences in selection — higher-educated people go to cities — could be part of the explanation of the rural-urban wage gap. Work exists exploring each of these explanations, with migration costs often treated as a residual explanation.

In Section 3 we showed that amenity differences often go in the opposite direction than would be needed to explain the higher wages in cities: instead of being lower in urban areas, which would balance out the increased earnings potential, measures of amenities are usually higher in cities (Gollin et al., 2021).

[24] Some additional challenges arise when trying to measure wages, especially for subsistence income earners, where further assumptions about the shape of the production function need to be made to compare the returns to labor across sectors. See the discussion in Gollin et al. (2014) for further information.

The literature on differences in schooling access across origins is small, but work evaluating the migration impacts of improving access to education does not find results consistent with the conjecture that differential access in schooling can explain the rural-urban wage gap (Khanna (2023); Hsiao (2023)). We then turn to the selection explanation before discussing migration.

The selection argument is important. If the gaps in earnings across space are in fact due to differential selection of the types of people who move to cities, then any efforts to encourage movement across space would only improve productivity if there were strong agglomeration externalities. The evidence in support of this argument is mixed. Young (2013), Hicks et al. (2021), and Lagakos et al. (2020), among others, have all used different assumptions to try to understand how much of the gaps might be explained by selection and appear to come to different conclusions. The most direct approach to isolating the role of selection involves subsidizing migration or looking for natural experiments that cause migration, and then asking if the returns to migration occur when selection is mechanically shut down by randomization. As an example of this, Bryan et al. (2014) study a conditional cash transfer given to poor households in northwestern Bangladesh in 2008. The cash transfer, which was about USD 8, was conditional on sending someone from the household as a migrant to seek work in one of the nearby cities during the lean season between planing and harvest. The results show a substantial gap in earnings between rural and urban areas for the same individuals and imply that selection alone cannot be the only reason for the rural-urban wage gap.

The last explanation for the rural-urban wage gap is the costs, frictions, and missing markets related to migration. These factors, which we have shown are higher in developing countries, are often missing from models, such as Young (2013), exploring the gap. But it is important to understand their source and, ultimately, the appropriate policy response. Here we review a non-exhaustive list of the sources of migration frictions.

*Credit constraints:* Migrating across space requires money, not only to pay for transport costs, but also to allow time to search for a job in the destination. There is evidence that credit constraints constrain movement. As noted above, the conditional cash transfer studied in Bryan et al. (2014) increased migration rates and is potentially evidence of a credit constraint.[25] Consistent with this, Banerjee et al. (2021) study the long-term impacts of a multifaceted "big push" program, which provided a large asset transfer to poor Indian households. They find that ten years after the program, household income and consumption remained higher, and that most of this gain occurred from migration, implying that saving for the cost of migration could be a barrier to movement.[26]

*Infrastructure and Roads:* Roads and other transport infrastructure are public goods, unlikely to be efficiently provided by the market. They also potentially

[25] Lagakos et al. (2023) offer a different interpretation of the costs in this paper.

[26] In the context of international migration out of Indonesia, Bazzi (2017) shows that income shocks increase out-migration from the poorest communities, but decrease out migration from richer communities, consistent with a credit constraint.

decrease both the upfront cost of migration and make it easier to return home, reducing the cost of being away from family. Morten and Oliveira (2024) find that roads built to accommodate the new capital city Brasilia in Brazil led to a large increase in migration between connected locations, with remote areas benefiting the most. Similarly, Asher and Novosad (2020) study rural road building in India, and find a large increase in employment out of agriculture, probably due to people finding wage work outside of the immediate village. The results suggest that roads are important for reducing the costs of moving out of rural villages.

*Land Markets:* Evidence suggests that land markets in the world's poorer countries do not often work well (e.g., Foster and Rosenzweig (2022)). This means two things. First, if land cannot be sold, or not for its full value, then migration away will be more costly due to the need to abandon income. This is especially true if property rights do not ensure that land cannot be seized. Second, in the presence of credit constraints, poorly functioning land markets mean that households cannot make use of their full wealth to fund migration. Consistent with this, de Janvry et al. (2015) show that a land registration program in rural Mexico significantly increased out-migration rates. Adamopoulos et al. (2024) take the argument one step further, using a quantitative model to estimate the fraction of migration frictions that are accounted for by land insecurity. Using data from China, their model suggests that as much as half of movement frictions are accounted for by land insecurity, and that removing insecurity would substantially increase agricultural productivity by reallocating people out of rural areas and agriculture.

*Information Frictions:* The canonical model assumes that potential migrants understand the job opportunities and amenities at the destination, but there are reasons to question whether this sort of information will flow smoothly. The experiment of Bryan et al. (2014) attempted to test for this friction by providing information on average earnings of recent migrants into nearby cities. They see no impacts of this intervention, but it is not clear if the information was understood and believed by the recipients. Baseler (2023) presents a more complete experiment in Kenya. He shows data strongly consistent with migrants hiding their income, with parents understating earnings relative to the migrants themselves, but not making similar mistakes for children who live at home. He also shows that an information intervention can lead to an increase in migration and earnings. This result is in-line with other papers that show information about opportunities from migration change behavior at the origin (e.g., Jensen (2012)).

*Risk:* Migration is also likely to be risky and insurance markets incomplete. Bryan et al. (2014) argue that this is part of the reason their migration subsidy is so successful at encouraging migration and show some evidence consistent with risk as a constraint from a separate experiment that provided rainfall insurance at one of the potential city destinations. More directly, Baseler et al. (2024) show that an intervention that reduces Indian households' belief that they will be able to receive food rations in the city if they migrate reduces out-migration rates. Munshi and Rosenzweig (2016) also find support for the role of risk and insurance. They argue that access to insurance between rural and urban areas is asymmetric if caste-based networks provide insurance that is limited to the origin. Consistent with this argument, they show that migration is more common for households that are *relatively* more wealthy than others in their caste network, and less common for those who face more income risk, and hence benefit more from network-based insurance. Morten (2019) and Meghir et al. (2020) find that the relationship between informal risk sharing and migration may inherently depend on how risky the destination is, and more broadly, informal insurance may be an important factor determining if, and when, people adopt new income-generating methods. All this research suggests that schemes that provide systematic insurance against loss of income, whether in rural or urban areas, could potentially reduce migration frictions.

*Labor markets:* Labor market frictions, both at the origin and destination, also have a potential role in dissuading migration. It is well documented that networks play an important role in getting jobs, particularly in larger urban labor markets. To the extent that migrants have access to smaller networks at the destination, their earnings may not be as high as those already in the market (Beaman, 2012; Tang, 2024). If this constraint reduces over time as migrants increase the size of their network, then it could account for the small fixed effect estimate of the return to migration.

*Policy:* Finally, policy differences across space is a potential source of migration costs. Development writers have long emphasized the possibility of urban bias in public policy making (e.g., Lipton (1977)). The particular form that this takes will determine whether it creates costs of movement. On the one hand, urban bias would lead to greater expenditure on public goods in urban areas, and hence encourage migration. On the other hand, urban bias may take the form of protecting existing urban populations from in-migration. In its most extreme form this would look something like China's Hukou system, which acts to restrict movement outside of a households registered location.[27] Far from urban bias, however, the development programs that have become the mainstay of analysis since the RCT revolution are predominantly rural. These programs, be they cash transfers, graduation, microfinance, fertilizer subsidies, or workfare programs are a potential movement friction if they are implicitly or explicitly conditioned on remaining in a rural area. For example, the world's largest social security program, India's National Rural Employment Guarantee Act (NREGA) provides guaranteed work and pay, but only for those who live in rural areas. Imbert and Papp (2020) study the impact of the program, and provide evidence consistent with the notion that it significantly reduces short-term rural-to-urban migration.

[27] More subtly, Feler and Henderson (2011) and Glaeser et al. (2005) argue that city government may have incentives to target public good provision and regulation in a way that deliberately excludes recent migrants.

Taking stock, this section has discussed the large wage gaps across space and the plethora of possible market failures that may explain these gaps. We argue that more work to understand spatial equilibrium is needed and advocate that it should model the economy carefully in each context, use microdata to discipline any model, and use exogenous variation to estimate model parameters accurately.

## 4.2 Are cities larger than optimal?

A related dimension of whether labor is allocated efficiently across space is whether cities are the correct size. Eight of the top ten megacities — cities with more than 10 million inhabitants — are located in low- and middle-income countries. If cities are growing too big, what does this mean for how policymakers should design urbanization in their rapidly-urbanizing countries?

A classic result in this area comes from Henderson (1975) who shows that, in general, each city tends to be too large. As a result, there are too few cities. The result is perhaps easiest to see if there is a single origin location $r$, which is rural, and a single alternative destination $c$, a city. Assume that production in the rural area is subject to constant returns to scale, so that the wage is a constant and equal to $w_r$ and normalize rural amenities and rental rates to 1 so that $V_r = w_r$. In the city assume there is no need for land and that production is subject to an agglomeration externality $\gamma > 1$, so the wage is just $\gamma \bar{A}_d L_d^{\gamma-1}$, which implies that wages rise with the population. We also assume that there is a congestion externality, and in a slight departure from above assume that it implies an optimal city size, so that $V_c$ is an inverted U-shape.[28] Now, equalization of utility across the two locations requires that $V_c = w_r$. This can occur at two points on the concave $V_c$ curve — one with a city smaller than the optimal from the perspective of city residents, and one larger than optimal from the perspective of city residents. A moment's reflection shows that only the latter is stable — if the equilibrium lies to the left of the maximum then a small increase in population will lead to a rise in $V_d$ and hence a further round of migration. The outcome is suboptimal because there is another allocation with a smaller city where some people are made better off and no one is made worse off. The argument is similar to the classic analysis in Harris and Todaro (1970), who instead get cities that are too large because of a binding minimum wage in the city that distorts migration choices. The basic result can also be applied across a set of cities, and implies that there will be too few and too large cities.

We wish to point out a few things about this result, which is derived in a model with no heterogeneity or migration costs. First, the result requires a specific assumption about the shape of congestion externalities. There is very little well-identified work on the shape of congestion in low-income cities, and as we have shown it is not certain that amenities decrease with density. This is an area

[28] This would require a departure from the functional forms above so that $V_d$ is not strictly increasing or decreasing.

that is ripe for research. Second, the result only applies if there are no market failures other than the congestion and agglomeration externalities. For example, suppose that there is a migration cost between the rural area and the city. If that cost exists because of a market failure, for instance the public good nature of road building, then that will cause too little migration to the city, offsetting the result above and leading to unclear welfare implications. This is of course nothing other than a statement of the theory of the second best, but it strikes us as important here. As we have emphasized, development economists tend to be skeptical of perfect markets assumptions in general, and making large policy prescriptions based on them seems hard to justify. Third, the analysis also tends to take the congestion and agglomeration parameters as fixed, but in reality they are likely to be policy choices, reflecting how well governments and others respond to market failures within and across locations. There is a concern that concentrating on a narrative about too many people in cities leads people to ignore the fact that more people could be accommodated if congestion was managed better. Finally, the fundamental cause of the finding that there are too few cities (each of which is too large) comes from a failure to coordinate on opening new cities, and could in principle be fixed by a coordinating institution such as a competitive building sector or a centralized government setting up secondary cities (see, e.g., Duranton and Puga (2004)). The takeaway, as we have highlighted throughout the chapter, is that we must carefully consider the form of the spatial model we use. The model used in Henderson (1975) makes many assumptions, including no migration costs, which are unlikely to be true in many developing countries.

## 4.3 Is urbanization really lower in low-income countries?

The previous two subsections both discuss the optimal allocation of people across space — between rural and urban areas or between different cities. But how to determine what is a city, or what separates a small village, medium-sized town, or dense urban area? There are different methodologies available that each have different implications for how we think about rates of urbanization and the current population distribution between rural and urban locations.

A prevalent narrative suggests that low-income countries are relatively under-urbanized, leading to questions about whether the developing world should increase its urbanization to match levels seen in developed countries. But is urbanization really lower in low-income countries? It turns out that it depends on the methods used to classify urbanization. Until recently researchers had to rely on national definitions for urbanization, decided by individual countries. These definitions vary widely. For example, in the United States, the 2020 census defines census blocks are urban if they have at least 2,000 housing units or 5,000 people, a definition that captures both the high-density urban core and lower-density suburbs. In contrast, Nigeria defines urban areas as towns that have a population of at least 20,000 people. Other countries have urban definitions that depend on the industrial composition of the workforce, level of infrastructure access, or other socioeconomic characteristics.

The population living in urban areas, according to these national definitions, is collected by the United Nations Department of Economic and Social Affairs. The dataset combines each country's census data and surveys with UN estimates for future years, and for years in which administrative data is not available, and spans 1950 to 2050. We present this data as the "country-specific" measure. These country-specific definitions should not be dismissed as they reflect the local perception of urban versus rural and allow for context-dependent designations. However, it is difficult to compare across countries or across time with these definitions, since they vary widely between countries and between years. Country statistical offices also sometimes rely on human judgment to determine urban areas, which makes replication a challenge.

What other options are available? Duranton (2021) presents an overview of methods. For instance, de Bellefon et al. (2021) and Combes et al. (2023) use a methodology that classifies areas as urban if they have considerably more population or building density than other areas of the country, creating replicable but dynamic thresholds for urbanization. Arribas-Bel et al. (2021) use a machine learning approach using data in Spain to classify urban areas based on building density. We focus on one methodology, the Degree of Urbanisation, a classification system created by six agencies led by the European Union (Dijkstra et al., 2021). The definition has the advantage of being standardized across countries, and captures the urban economics assumption that it is population density that matters for agglomeration and congestion. We choose to highlight this approach because it is a simple and transparent methodology that only requires data that is readily available across the globe, but we note that there are many other methods for classifying urbanization that each have their own advantages and disadvantages. The accuracy of the Degree of Urbanisation and other methods also relies on the accuracy of the underlying population or building data they use.

The Degree of Urbanisation determines the urbanization level of a gridded map of the globe according to population level and density from GHSL data.[29] Table 8 summarizes the criteria; in the most expansive definition of urbanization, which includes suburban areas and semi-dense urban areas, an area is considered urban if it has a population of at least 300 people per square km. and at least 5,000 people in a cluster of contiguous grid cells. The data is available every five years from 1975 to 2030 (2025 and 2030 are based on projections). We present this data as the "density-based" or "Degree of Urbanisation" measure.

Fig. 8a (replicated from the introduction) shows the share of population living in an urban area, as defined by each country's definition, for each region between 1950 and 2050. It shows the standard story: Africa and Asia were significantly less urbanized than Europe and the Americas in 195, and have since undergone a period of rapid urbanization. However, in 2025 they are still much

**TABLE 8 Criteria for Different Types of Urban Areas in the Degree of Urbanisation Methodology.**

| Type of Urban Area | Density | Population | Distance to Neighbors |
|---|---|---|---|
| Urban Centre | $\geq 1500/km^2$ | $\geq 50{,}000$ | |
| Dense Urban Cluster | $\geq 1500/km^2$ | 5,000 - 50k | |
| Semi-dense Urban Cluster | $\geq 900/km^2$ | $\geq 2{,}500$ | $\geq 2$ km away |
| Suburban or Peri-urban | $\geq 300/km^2$ | $\geq 5{,}000$ | |

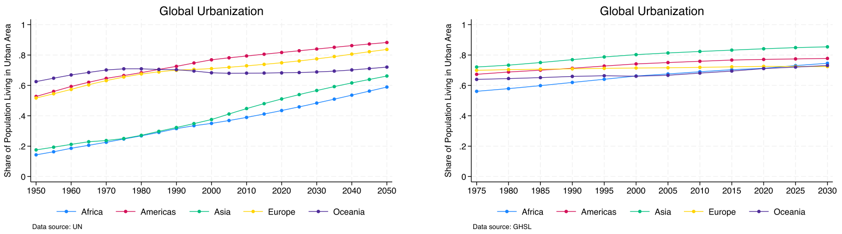


(a) Country-specific measure (b) Density-based measure

**FIGURE 8** Country-specific vs. density measures of urbanization.

less urbanized than the rest of the world and are not predicted to catch up until after 2050, despite high rates of urbanization.

Fig. 8b shows the same plot, but uses the Degree of Urbanisation measure: areas are urban if there are at least 5,000 people living in contiguous cells that have at least 300 people per square km. Using this data, the story of urbanization looks different. At the beginning of the data series in 1975, Asia was the most urbanized of all the regions and Africa was the least urbanized, but there were not major differences. Over the following 50 years, Africa caught up, overtaking Europe, so that all of the regions have similarly high levels of urbanization and low rates of urbanization. This is an important caveat to the traditional narrative that low-income countries are much less urbanized than the rest of the world and shows how important the choice of urbanization definition is. We explore this concept more thoroughly in a related working paper (Bryan et al., 2025).

The differences between definitions are especially prominent when looking at specific countries, as in Fig. 9. The blue line shows the country-specific definition, while the red areas show progressive levels of urbanization from the Degree of Urbanisation categorization. The highest red line shows the most expansive definition, which includes the suburban and semi-dense areas, and which we used in the previous figure for total levels of urbanization. Fig. 9 reveals several facts. First, the national definitions in the wealthier countries, the UK and the US, are most well-aligned with an expansive view of what is urban, while this is

[29] More detail is provided in Appendix A.5.

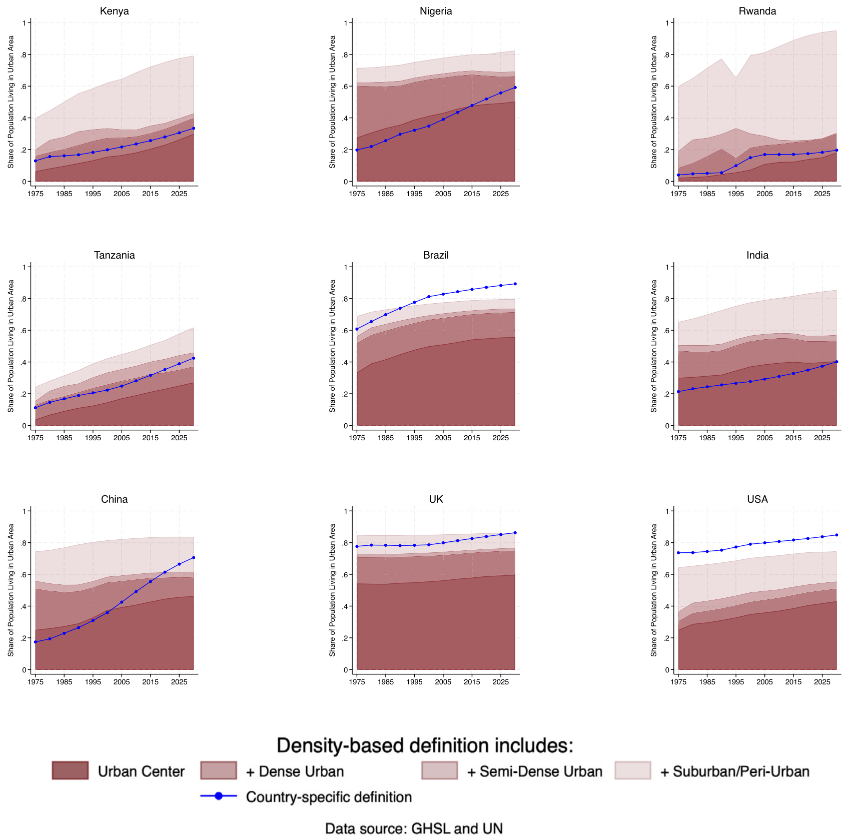


**FIGURE 9** Country Comparisons: Urbanization Patterns.

much less true in the poorer countries. India, Kenya, and Rwanda's national definitions seem much better aligned with the definition of urban that includes only the most dense areas. The Chinese and Nigerian country-based definitions do not seem to accord well with any one of the density-based measures. Second, the rate of urbanization depends crucially on the measure used. In China and Nigeria, for instance, the country-specific definition shows a high rate of urbanization since 1975, but the density-based definition, using any of the thresholds, shows a much slower rate of increase. On the other hand, the country-specific definition in Kenya and in Rwanda after 2000 has a slower rate of urbanization than the density-based definition. In some countries, like Tanzania and the United States, the rates are similar. But globally we see much higher levels of urbanization and lower rates of urbanization when using the density-based definition.

While we don't take a strong stance on *how* urbanization should be determined, we do want to point out that the choice of definition matters. Similarly to thinking about how the canonical model should be adapted to each context, we urge researchers to carefully consider what measure of urbanization is most relevant to their work.

## 5 Avenues for future work

While issues of urbanization were a mainstay of early development economics (e.g., Lewis (1954); Harris and Todaro (1970)), recent work has been dominated by the search for clean identification, which typically means RCTs. The search for clean identification is for good reason. The resources available for anti-poverty and development programming are small, and the costs of wasting these resources on programs that do not work are large. But this focus on RCTs has meant a focus that is often on rural issues, where it is more straightforward to assign villages to treatment and control and get clean estimates that account for any within-community spillover effects. Classic urban issues are much harder to study in this way. Critics of the RCT approach often argue that this move has left important development issues, including urbanization, understudied.

Development economists bring many tools to the table to study urban issues. Standard in the development economist's toolkit is experience with collecting data by surveying households directly; running RCTs, either in partnership with NGOs or government agencies; and experience developing partnerships with firms, NGOs, and government agencies to get access to administrative data. Low-income countries usually have a paucity of formal administrative data, making such "hands-on" research necessary to get data to estimate models. On the other hand, urban economists are fluent in the modeling toolbox needed to interpret and understand endogenous price changes across space, but may not have the experience required to access or generate the data needed to bring the model to the data in a context-specific way or to extend the model in directions relevant for accounting for the pervasive market frictions that exist in many low-income countries. Both the standard development approach of RCTs and the model-based work of urban economists have costs and benefits, but answering our important questions almost surely requires more work that combines the benefits of both.

In this section we aim to chart a path forward for future work. We highlight three areas ripe for future research: incorporating frictions into spatial models, harmonizing data collection to theoretical frameworks, and improving the estimation of context-specific elasticities in data-scarce environments. Our goal in this section is to issue a "call to arms" for collaboration between development and urban economists to combine research methodologies to better estimate the parameters of urban models and add much more realism to the estimates that come from these models. We hope that this type of partnership will push forward the promising research agenda involving urbanization in low-income countries.

## 5.1 Incorporating market frictions into spatial models

The canonical spatial model is an extremely useful workhorse tool. However, simply employing off-the-shelf versions of models that were developed for use in the world's wealthier countries may miss important channels. More research is needed to extend the spatial framework to account for pervasive frictions present in the low-income world.

Gechter and Tsivanidis (2023) make progress on extending the baseline spatial model to account for dual housing markets. The paper combines an interesting historical episode with meticulous data collection and an adaptation of the standard spatial commuting model to account for formal and informal housing sectors. The focus is the redevelopment of a large number of formerly-industrial mills (which account for 15% of the city) in Mumbai. The authors wanted to study the spillover effects of the redevelopment onto the neighboring slum neighborhoods. As the former mills developed, how did redevelopment spill over to informal housing and what is the broader impact of the redevelopment after taking into account the displacement of residents?

As with all spatial papers, a key input is highly geographically disaggregated data. Population and employment data was easily sourced from population and economic censuses, floorspace data digitized from historical valuation books, and slum data meticulously geocoded from slum maps overlaid over satellite images to get historical boundaries of slums. The resultant dataset is a city-level dataset of population, socio-economic status, floorspace price, employment, industry, and slum status at 800 geographical units each containing approximately 15,000 people.

The model is based on the standard commuting model we outlined above. However, the authors make important modifications to capture the duality of housing markets. In the model, developers choose between formal and informal development. If they redevelop a slum, developers need to pay compensation to residents, with the parameterization based on the specific policy. Formal land can have multi-story buildings; informal land has single-story buildings. The model also extends to include dynamic choice, rather than a static one. The authors make good progress in adapting the model to the setting and are successful at estimating almost all the elasticity parameters for their specific context. Empirically, the paper finds that the construction of high-rise buildings had spillovers onto nearby slum neighborhoods consistent with gentrification: the share of slums falls, the cost of housing increases, and population density falls.

Just as this paper extends the model to account for the duality of housing markets, future work could extend the model to account for the duality of labor markets, adding the presence of capital market frictions, and the pervasive risk inherent in agriculture. Political economy considerations are also key and affect many resource allocation decisions, but again are rarely modeled. There is a rich empirical literature documenting the presence of many of these frictions; understanding how they change the conclusions of policy-at-scale or affect the conclusions of programs with high likelihood of spillovers is key.

### 5.2 Cleaner identification of context-specific model elasticities

A second area for future research is cleaner identification of key model elasticities, specific to the context that the model is being estimated. We are particularly excited about combining variation from RCTs that can cleanly identify elasticities that can then be used in spatial models. Work in this space has taken two forms, either utilizing existing RCTs and tying the variation to structural models (an early example is Todd and Wolpin (2006) with the PROGRESA cash transfer experiment in Mexico), or, more innovatively, designing the RCT itself to explicitly identify the structural parameters needed.

An example of a project with a tight link between theory and experimental design is Kreindler (2024). This paper provides a very clear example of how to integrate modern RCT and data collection into a spatial equilibrium model. The RCT seeks to identify commuting responses that are then fed into an equilibrium model of commuting choice to understand optimal design of equilibrium congestion pricing. Congestion pricing is a standard economic prescription to reduce road congestion, a problem in many developing country megacities (Akbar et al., 2023b). The impact of congestion pricing, however, depends crucially on the elasticity of commute time and route to price, and the elasticity of congestion to commute time and route. These elasticities are relatively difficult to estimate given observational data (Small et al., 2005). Kreindler (2024) is able to make excellent use of modern development methods to provide a compelling response. First, he provides a sample of drivers with cell phones that passively track GPS. This allows him to measure travel speeds at different times of day, and to extract estimates of travel time to congestion, making use of within-day variation in congestion as instruments. Second, he gives a stock of money to a set of participants with the GPS app and implements an experimental congestion pricing scheme, giving clear experimental estimates of the relevant commuting elasticities. The results show that, in his setting of Bangalore, the elasticity of commute time to price is very small, and that the congestion externality is close to linear. These two facts combine to mean that there are very small benefits from optimal congestion pricing in his context. The paper is an inspiring example of combining the modern methods of development economics with structural modeling and classic urban questions. Perhaps more important than the results in Bangalore, the paper provides a portable method that can be implemented in any setting to determine where there are likely to be large gains from congestion pricing.

### 5.3 Finding and identifying novel data sources

Finally, a running theme throughout this chapter is the need to estimate elasticities — labor supply, housing, migration, commuting — for the specific context. However, in order to do this, the relevant data needs to exist. In many settings, data is sparse, missing, collected with measurement error, or collected at too-aggregate a scale. For example, many countries do not regularly collect even census data: Ethiopia last undertook a national census in 2007; Nigeria in 2006. The Demographic and Health Survey (DHS) program is one data source widely available across countries, but often only at sporadic intervals. The World Bank collects survey data for many countries through the Living Standards Measurement Survey (LSMS) program, but not on a regular cycle for each country. Standard labor force surveys, especially ones with a panel dimension, are often also not collected. In areas with high informality there are no centralized tax or education records to link individuals over time. This data paucity means that researchers often need to collect their own data. However, such efforts are often local in both space and time. An additional challenge for spatial frameworks is that often a large amount of data — data for every location either directly or indirectly affected — is needed. If people migrate in response to policies, we may also want to track individuals to see how they fare in their new location.

Given the paucity of development, the first step is identifying how to "fill in the gaps" and identify the needed data. Here we again see exciting avenues for future work. Many new data sources — satellite data, cellphone records (although ownership of smart cellphones is not universal amongst the world's poor, nor is cellphone connectivity) — are being used. The chapter Abramitzky et al. (2025) in this handbook includes many such examples of novel data sources.

One example of innovative data collection is Harari and Wong (2024), who study a slum upgrading policy that took place in Jakarta, Indonesia, between 1969 and 1984. Enough time has passed to ask about the long-run effects of the policy. However, the paper quickly runs into a large empirical challenge: how can consistent measures of land development be constructed when a large portion of the city is informal housing, which does not appear on formal sales records? The authors overcome the data challenge by analyzing pictures of neighborhoods drawn from Google Street View imagery. A random sample of the city was drawn and then the pictures are coded to measure building height, which can be done whether the building is a skyscraper or a shack. However, 10% of the selected pixels did not have Google Street View images, either because the streets were too narrow or the community was gated, leading to non-random sample selection. The authors creatively fixed the missing data by collecting it themselves, sending enumerators to the field to take photos and then coding the photos based on building height in a consistent way as the Google Street View images. This bespoke data collection allows the full sample of the city to be analyzed, and gives a measure of development that accounts for both formal and informal structures.

In other settings, theory can be used to fill in the gaps of missing data as illustrated by the exercise in Kreindler and Miyauchi (2023). The authors wish to understand whether cellphone data can predict the distribution of economic activity across two urban areas in Dhaka, Bangladesh and Colombo, Sri Lanka. If the cellphone data does predict economic activity well, then can high-frequency cellphone records be used to understand the economic effects of high-frequency events such as transportation strikes? The exercise uses the gravity relationship derived in Eq. (6) in Section 2. People should tend to commute towards areas with higher wages and so the destination fixed effect in the gravity regression contains the information about how attractive the destination is. The authors therefore estimate gravity equations using the implied commuting flows from cellphone data, and then show that the destination fixed effects from this regression correlate well with survey measures of workplace income. If this relationship holds more generally, then a combination of theory and data — only commuting flows — can help fill in the gaps for unobserved data — wages and incomes. Future work could extend the intuition in this paper to look for other settings where some data is observed and determine whether that allows users to infer unobserved data about poverty and income.

## 6 Conclusion

Huge strides in spatial modeling and data accessibility have enabled rigorous quantitative analysis of critical urban and development economics questions that would have been unthinkable just a decade ago. Research at this intersection is not only of academic interest but is increasingly sought after by policymakers worldwide and backed by major research initiatives. For example, one of the International Growth Centre's (IGC's) four thematic priorities for the developing world is cities. The IGC funds and promotes research aimed at "harnessing the positive aspects of density ... while reducing its downsides of pollution, traffic, and disease." Achieving this goal requires precisely the type of innovative methodologies outlined in this chapter. Reflecting this demand, the IGC's Cities that Work initiative has even developed an R package (https://cran.r-project.org/web/packages/IGCities/index.html) — a simplified yet accessible version of a within-city spatial model — to help policymakers in low- and middle-income countries assess the urban economic implications of proposed policies. This underscores the growing interest in leveraging urban economic research to tackle pressing policy challenges.

The spatial models outlined in this chapter are designed to be user-friendly, relatively easy to apply to data, and capable of addressing key issues such as spillovers and scale-up effects. However, the models do not yet fully capture the market frictions that characterize developing economies. This limitation points to a crucial avenue for future research: integrating spatial models with the bespoke data collection, randomized evaluations, and attention to market failures that are central to development economics. Pursuing this work will significantly enhance the ability to design effective, evidence-based policies that improve urban living conditions, particularly for the poor. By refining models to reflect the complexities of real-world developing economies, researchers can provide policymakers with even more powerful tools to navigate challenges such as infrastructure investment, growing population, and climate change adaptation. The result is sure to enhance the impact of our work, the quality of our predictions, and, hopefully, the lives of the poor.

# Appendix A Data

## A.1 IPUMS data

The IPUMS data (Ruggles et al., 2024) used in Section 3.6 is available from https://international.ipums.org/international-action/samples and was downloaded in March–May, 2024, except for the datasets for Burkina Faso, Egypt, Ghana, Mali, and Sudan, which were downloaded in 2015. It provides harmonized micro census data compiled from national statistics offices. For each of our sample of countries, we download the most recent dataset for the country and restrict our analysis to heads of households. The geographic unit of analysis is the subnational geographic level available for the recorded birth place — usually the 1st level but sometimes the 2nd level. We count the number of people in each origin-destination pair (i.e. the number of people who were born in place $o$ and migrated to place $d$), we well as the number of people who stayed in place.

We use this data to calculate the modified Head-Ries index of migration cost for each origin-destination pair, setting $\theta = 3$:

$$c_{od} = \left( \frac{\pi_{od}}{\pi_{oo}} \times \frac{\pi_{do}}{\pi_{dd}} \right)^{\frac{-1}{2\theta}}$$

We also calculate the distance between the centroid of the origin and the centroid of the destination using the GIS shapefiles provided by IPUMS.

Table A.1 summarizes the year of data and the level of geographical region for each country. We also wish to acknowledge the statistical offices that provided the underlying data making this research possible. We include the relevant statistical office in the table.

## A.2 World urbanization prospects 2018

The World Urbanization Prospects is a project of the Union Nations Department of Economic and Social Affairs. It collects information on urbanization provided by national statistics offices. The 2018 dataset includes both the information on urbanization provided by the individual countries as well as projections from the UN to fill in the data at five-year intervals. Its website has the link to download the data and also includes information on the sources and methodology: https://population-un-org.stanford.idm.oclc.org/wup/. We downloaded the data on January 31, 2024.

We take "share urban," given in percentage form, from the datasheet "WUP2018-F21-Proportion_Urban_Annual" and multiply it with total population from "WUP2018-F05-Total_Population." This gives the number of people living in an urban area at every five year period in each country. For the region-level analysis, we sum up the total number of people and the number of people living in urban areas for all areas in the region, using the classification from the UN Statistics Division: https://unstats-un-org.stanford.idm.oclc.org/unsd/methodology/m49/overview.

**TABLE A.1 Summary of IPUMS Data.**

| Country | Year | Geographical Level | Statistical Office |
|---|---|---|---|
| Burkina Faso | 2006 | 2nd (Province) | National Institute of Statistics and Demography |
| Cambodia | 2019 | 1st (Province) | National Institute of Statistics |
| Canada | 2011 | 1st (Province) | Statistics Canada |
| Chile | 2017 | 1st (Province) | National Institute of Statistics |
| Egypt | 2006 | 1st (Governate) | Central Agency for Public Mobilisation and Statistics |
| Ghana | 2010 | 1st (Region) | Ghana Statistical Service |
| Indonesia | 2010 | 1st (Province) | BPS Statistics Indonesia |
| Laos | 2015 | 1st (Province) | Statistics Bureau |
| Mali | 2009 | 1st (Region) | Central Bureau of the Census |
| Mexico | 2020 | 1st (State) | National Institute of Statistics, Geography, and Informatics |
| Senegal | 2013 | 2nd (Department) | National Agency of Statistics and Demography |
| Sierra Leone | 2015 | 2nd (Chiefdom/ward) | Statistics Sierra Leone |
| Sudan | 2008 | 1st (State) | Central Bureau of Statistics |
| Tanzania | 2012 | 1st (Region) | Bureau of Statistics |
| United States | 2020 | 1st (State) | Bureau of the Census |
| Uruguay | 2011 | 1st (Department) | National Institute of Statistics |

## A.3 World Bank GDP data

GDP data comes from the World Bank: NY.GDP.PCAP.PP.KD. It contains the GDP per capita (PPP) in 2021 dollars. It was downloaded on October 26, 2023. The methodology is available on the World Bank website. https://data.worldbank.org/indicator/NY.GDP.PCAP.PP.KD.

## A.4 Share of employment in agriculture

The data on the share of employment in agriculture for each country come from the modeled International Labour Organization (ILO) estimates accessed through the World Bank Development Indicators: SL.AGR.EMPL.ZS. It was downloaded on January 16, 2025. Employment is defined as persons of working age who were engaged in any activity to produce goods or provide services for pay or profit, whether at work during the reference period or not at work due to temporary absence from a job, or to working-time arrangement. The agriculture sector consists of activities in agriculture, hunting, forestry and fishing, in accordance with division 1 (ISIC 2) or categories A-B (ISIC 3) or category A (ISIC 4). More information is available on the World Bank website. https://databank.worldbank.org/id/25b2e99e.

## A.5 GHSL data

In this handbook chapter, we use the Degree of Urbanisation data that are part of the Global Human Settlement Layer (GHSL). Information about the GHSL can be accessed here: https://human-settlement.emergency.copernicus.eu/download.php?ds.

We use data on the share of population living in urban areas from the GHS-DUC-2023A V2_0, downloaded on December 18, 2024. It includes the country-level "level 0" total number of people, number of people living in urban areas, and number of people living in rural areas, according to the Degree of Urbanisation classification system. As with the country-specific data, we then add up the total number of people and the number of people living in urban areas for all countries in each region.

To obtain the number of people living in urban areas we add together the different types of "urban" locations in the main figures, and also break them up into the component parts: [UCentre_pop, DUC_Pop, SDUC_Pop, SUrb_Pop], which are respectively [urban centre population, urban cluster population, semi-dense urban cluster population, suburban or peri-urban population].

### *A.5.1 Details on GHSL methodology and verification exercises*

Here, we briefly outline the methodology of the GHSL dataset as of Version 2 2023 (R2023A), but for specifics we refer to the documentation on the GHSL website: https://human-settlement.emergency.copernicus.eu/download.php?ds.

In order to classify cells as rural or urban, the GHSL uses several input layers. The first is the GHS-BUILT-S dataset, which contains a gridded map of buildings in the world. The dataset also distinguishes between residential buildings and non-residential buildings. The data come from satellite images from the Landsat and Sentinel-2 (S2) collections in 1975, 1990, 2000, 2014, and 2018. The model also uses data from Facebook settlement delineation, Microsoft, and Open Street Map building delineation as a learning set. The data is interpolated to get the output dataset at five year intervals from 1975 to 2030, and also interpolated when data is not available. Although the input satellite imagery is at the 10-meter resolution, the end product of GHS-BUILT-S is at the 100 meter or 1 km. resolution: the fraction of land that is built-up in each grid cell.

The previous release was the R2019 dataset, which is in the process of being validated. The data is validated by comparing the built-up classification to the input satellite images by human eye and by comparing the classification to other datasets. In the human inspection of 250,000 datapoints there was an overall accuracy of 82.6%, and the majority of the discrepancies were driven by disagreement about whether a building was residential or non-residential. The validation showed that the dataset is most accurate for regions in Oceania, South Eastern Asia, and Europe, and least accurate for Southern Africa, Central Asia, and Western Asia. Relative to the 2019 version, the 2023 version is more accurate. The validation exercise also shows that the data is more accurate in the most recent epochs, a point which Henderson et al. (2021a) notes.

The GHS-BUILT-H dataset is a corollary dataset which gives global building height at the 100-meter resolution level for 2018, using data from the ALOS Global Digital Surface Model ALOS World 3D-30m (AWD3D30), the NASA Shuttle Radar Topographic Mission data - 30m (SRTM30), and the Sentinel-2 global pixel based image composite from L1C data 2017-2018. GHS-BUILT-H and GHS-BUILT-S are combined to form GHS-BUILT-V, which gives gridded (100 meter and 1 km) global built-up volume in cubic meters: the built-up area of the grid multiplied by the building height. GHS-BUILT-V is available for each five year period 1975–2030.

Next, the GHS-BUILT-V is combined with census data to create GHS-POP: a dataset of the number of people per cell from 1975-2030 at the 100 meter or 1 km. level. Global census data comes from the Gridded Population of the World, version 4.11 (GPWv4.11), harmonized by CIESIN. The GHS layer disaggregates the population from the census unit level to the grid cells using the information on buildings from GHS-BUILT-V and only assigning people to residential buildings. The methodology for this is described in the GHSL documentation and not discussed here.

The GHS-POP layer is used to do the classification of urbananization, resulting in the output dataset GHS-SMOD, which classifies each pixel (1 km.) as urban or rural, further breaking down urban into its component clusters noted above. Here, we copy the definitions for each type of classification directly from the GHSL documentation:

**“Urban Centre” (also “High Density Cluster” - HDC)**: An urban centre consists of contiguous grid cells (4-connectivity cluster) with a density of at least 1,500 inhabitants per km$^2$ of permanent land, and has at least 50,000 inhabitants in the cluster with smoothed boundaries (3-by-3 conditional majority filtering) and <15 km$^2$ holes filled

**“Dense Urban Cluster” (also “Dense, Medium Cluster”)**: A Dense Urban Cluster consists of contiguous cells (4-connectivity cluster) with a density of at least 1,500 inhabitants per km$^2$ of permanent land and has at least 5,000 and less than 50,000 inhabitants in the cluster

**“Semi-dense Urban Cluster” (also “Semi-dense, Medium Cluster”)**: A Semi-dense Urban Cluster consists of contiguous grid cells (using four-point contiguity) with a density of at least 900 inhabitants per km$^2$ of permanent land, has at least 2,500 inhabitants in the cluster and is at least 2-km away from other Urban Clusters

**“Suburban or peri-urban grid cells” (also “Semi-dense grid cells”)**: Suburban cells are all other cells that belong to an urban cluster — contiguous grid cells (8-connectivity cluster) with a density of at least 300 inhabitans per km$^2$ of permanent land and a population of at least 5,000 in the cluster — but are not part of a Urban Centre, Dense Urban Cluster or a Semi-dense Urban Cluster

All other cells are classified as rural cells.

In this chapter we use the GHS-DUC dataset, which simply aggregates the information from GHS-SMOD at the level of administrative units. This gives us country-level information on population living in urban areas.

**TABLE A.2** Definitions of Labor Force and Household Indicators.

| Indicator | Definition |
|---|---|
| **Labor Force Including Subsistence Agriculture** | Somebody who, in the last seven days, worked for a wage, worked for themselves in agriculture or a non-farm enterprise, worked as a trainee or apprentice, did not work but has a job to return to, or sought work in the past four weeks. |
| **Labor Force Not Including Subsistence Agriculture** | Same as above, except subsistence agriculture is excluded if their only participation in the labor force is through working for themselves in agriculture. |
| **Formal Employment** | Somebody who worked for a wage or as a trainee/apprentice in the last seven days. |
| **Informal Employment Including Subsistence Agriculture** | Somebody who worked for themselves in agriculture or non-agriculture in the last seven days. |
| **Informal Employment Not Including Subsistence Agriculture** | Same as above, except subsistence agriculture is excluded if their only employment was through working for themselves in agriculture. |
| **Worked 7 Days Including Subsistence Agriculture** | Somebody who is either formally employed or informally employed, including subsistence agriculture. |
| **Worked 7 Days Not Including Subsistence Agriculture** | Somebody who is either formally employed or informally employed, excluding subsistence agriculture. |
| **Monthly Wage (Naira)** | Amount paid in the last pay period in Naira for those who reported wage work in the past seven days. |
| **Number in Household** | Number of people who normally live and eat their meals together in a household. |
| **Slum** | A binary indicator for whether the household is considered a slum, based on four UN-Habitat Urban Indicators: inadequate access to safe water, inadequate access to sanitation, poor structural quality of housing, or overcrowding. |
| **Own Residence** | Whether the household owns the dwelling they occupy. |
| **Imputed Monthly Rent (Naira)** | Amount being paid in rent, or the amount the household could charge if renting out an owned dwelling. |
| **Share of Consumption on Food** | The proportion of food consumption out of the total consumption (includes food, non-food, education, health, and rent). |
| **Share of Consumption on Rent** | The proportion of rent consumption out of the total consumption (includes food, non-food, education, health, and rent), where rent is hedonic rent if the household owns the dwelling. |

## A.6 Nigeria LSS 2019

The Nigeria Living Standards Survey 2018–2019 is available on the World Bank website: https://microdata.worldbank.org/index.php/catalog/3827/get-microdata. We downloaded the dataset on October 12, 2023. The survey is a collaboration from Nigeria's National Bureau of Statistics and the World Bank. The urban/rural demarcation is from the survey itself and therefore the country's definition of urban. Table A.2 shows how we define each variable.

NB: We define "subsistence agriculture" as doing agricultural work for the household where the products are "only for family use" or "mainly for own/family use but some for sale/barter." We do not consider household agricultural work "only for sale/barter" or "mainly for sale/barter but some for own/family use" as subsistence agriculture.

The UN-Habitat Indicators for a slum are precisely defined as follows:

- Inadequate access to safe water: main source of water during rainy or dry season one of unprotected dug well, unprotected spring, or surface water
- Inadequate access to sanitation and other infrastructure: toilet household usually uses is one of latrine without slab, bucket, hanging toilet, or there are no facilities
- Poor structural quality of housing: at least one of the following holds:
  - Dwelling is tent, improvised home (kiosk, container), or uncompleted building
  - Wall materials are mud, wood or bamboo, iron sheets, or cardboard
  - Roof is made of thatch, plastic sheet, or mud
  - Floor is sand/dirt/straw
- Insufficient area: more than 3 people share the same habitable room

Table A.3 shows summary statistics for the components of a slum among households designated as a slum.

**TABLE A.3 Components of Slum Households.**

| | Urban | Rural |
|---|---|---|
| Unclean Water | 0.09 | 0.46 |
| Bad Sanitation | 0.48 | 0.73 |
| Insufficient Area | 0.49 | 0.22 |
| Nondurable Housing | 0.27 | 0.67 |
| N | 2862 | 12418 |

*Notes:* Data source: Nigeria LSS Survey (2018-2019). Weighted at the household level. Sample only includes households designated as slums.

# Appendix B Derivations

## B.1 Share of people from origin $o$ migrating to location $d$ ($\pi_{od}$)

Each individual chooses to live in the destination $d$ that maximizes utility:

$$\max_d \frac{V_d}{c_{od}} \epsilon_d^i.$$

Therefore the probability of individual $i$ migrating to location $d$ is:

$$P\left(\frac{V_d}{c_{od}}\epsilon_d^i > \frac{V_d}{c_{od}}\epsilon_{d'}^i\right) \text{ for all } d' \neq d$$

We can use the properties of the Fréchet distribution to simplify this:

$$\begin{aligned}
&P\left(\frac{V_d}{c_{od}}\epsilon_d^i > \frac{V_{d'}}{c_{od'}}\epsilon_{d'}^i \quad \forall d' \neq d\right)\\
&= \int_0^\infty P\left(\frac{V_d}{c_{od}}\epsilon_d^i > \frac{V_{d'}}{c_{od'}}\epsilon_{d'}^i \quad \forall d' \neq d \mid \frac{V_d}{c_{od}}\epsilon_d^i = t\right)\theta\left(\frac{V_d}{c_{od}}\right)^\theta\\
&\quad \times t^{-(\theta+1)} \exp\left(-t^{-\theta}\left(\frac{V_d}{c_{od}}\right)^\theta\right) dt\\
&= \int_0^\infty P\left(\frac{V_{d'}}{c_{od'}}\epsilon_{d'}^i < t \quad \forall d' \neq d\right)\theta\left(\frac{V_d}{c_{od}}\right)^\theta t^{-(\theta+1)} \exp\left(-t^{-\theta}\left(\frac{V_d}{c_{od}}\right)^\theta\right) dt\\
&= \int_0^\infty P\left(\epsilon_{d'}^i < \frac{t c_{od'}}{V_{d'}} \quad \forall d' \neq d\right)\theta\left(\frac{V_d}{c_{od}}\right)^\theta t^{-(\theta+1)} \exp\left(-t^{-\theta}\left(\frac{V_d}{c_{od}}\right)^\theta\right) dt\\
&= \int_0^\infty \left(\frac{V_d}{c_{od}}\right)^\theta \theta t^{-(\theta+1)} \prod_{d' \neq d} \exp\left(-t^{-\theta}\left(\frac{V_{d'}}{c_{od'}}\right)^\theta\right) \exp\left(-t^{-\theta}\left(\frac{V_d}{c_{od}}\right)^\theta\right) dt\\
&= \int_0^\infty \left(\frac{V_d}{c_{od}}\right)^\theta \theta t^{-(\theta+1)} \exp\left(-t^{-\theta}\sum_{d'}\left(\frac{V_{d'}}{c_{od'}}\right)^\theta\right) dt\\
&= \frac{\left(\frac{V_d}{c_{od}}\right)^\theta}{\sum_{d'}\left(\frac{V_{d'}}{c_{od'}}\right)^\theta}\left[-\exp\left(-t^{-\theta}\sum_{d'}\left(\frac{V_{d'}}{c_{od'}}\right)^\theta\right)\right]_0^\infty\\
&= \frac{\left(\frac{V_d}{c_{od}}\right)^\theta}{\sum_{d'}\left(\frac{V_{d'}}{c_{od'}}\right)^\theta}[-0 - -1]\\
&= \frac{\left(\frac{V_d}{c_{od}}\right)^\theta}{\sum_{d'}\left(\frac{V_{d'}}{c_{od'}}\right)^\theta}
\end{aligned}$$

By the law of large numbers, this is also $\pi_{od}$, the share of people from origin $o$ migrating to location $d$.

## B.2 Head-ries index for migration cost

Recall that we defined the share of people from origin $o$ migrating to location $d$, $\pi_{od}$ as:

$$\pi_{od} = \frac{\left(\frac{V_d}{c_{od}}\right)^{\theta}}{\Phi_o}$$

By symmetry,

$$\pi_{do} = \frac{\left(\frac{V_o}{c_{do}}\right)^{\theta}}{\Phi_d}$$

We can also define the probability of staying in the same location, setting the migration cost to be 1 (i.e., costless):

$$\pi_{oo} = \frac{\left(\frac{V_o}{1}\right)^{\theta}}{\Phi_o}$$

$$\pi_{dd} = \frac{\left(\frac{V_d}{1}\right)^{\theta}}{\Phi_d}$$

We can write out the ratio of people choosing to migrate:

$$\frac{\pi_{od}}{\pi_{oo}} = \left(\frac{V_d}{V_o} \times \frac{1}{c_{od}}\right)^{\theta}$$

$$\frac{\pi_{do}}{\pi_{dd}} = \left(\frac{V_o}{V_d} \times \frac{1}{c_{do}}\right)^{\theta}$$

If we assume $c_{od} = c_{do}$ we can multiply the two ratios and solve for $c_{od}$:

$$\frac{\pi_{od}}{\pi_{oo}} \times \frac{\pi_{do}}{\pi_{dd}} = \left(\frac{1}{c_{od}^2}\right)^{\theta}$$

$$\left(\frac{\pi_{od}}{\pi_{oo}} \times \frac{\pi_{do}}{\pi_{dd}}\right)^{\frac{1}{\theta}} = \frac{1}{c_{od}^2}$$

$$c_{od} = \left(\frac{\pi_{od}}{\pi_{oo}} \times \frac{\pi_{do}}{\pi_{dd}}\right)^{\frac{-1}{2\theta}}$$

This final result, $c_{od} = \left(\frac{\pi_{od}}{\pi_{oo}} \times \frac{\pi_{do}}{\pi_{dd}}\right)^{\frac{-1}{2\theta}}$, gives us the Head-Ries index for migration costs.

## References

Abramitzky, R., Boustan, L., Storeygard, A., 2025. New data and insights in regional and urban economics. In: Donaldson, D., Redding, S.J. (Eds.), Handbook of Regional and Urban Economics, vol. 6. Elsevier North Holland, Amsterdam.

Adamopoulos, T., Brandt, L., Chen, C., Restuccia, D., Wei, X., 2024. Land security and mobility frictions. The Quarterly Journal of Economics 139 (3), 1941–1987.

Adamopoulos, T., Brandt, L., Leight, J., Restuccia, D., 2022. Misallocation, selection, and productivity: a quantitative analysis with panel data from China. Econometrica 90 (3), 1261–1282.

Akbar, P., Couture, V., Duranton, G., Storeygard, A., 2023a. Mobility and congestion in urban India. American Economic Review 113 (4), 1083–1111.

Akbar, P.A., Couture, V., Duranton, G., Storeygard, A., 2023b. The Fast, the Slow, and the Congested: Urban Transportation in Rich and Poor Countries.

Allen, T., Arkolakis, C., 2014. Trade and the topography of the spatial economy. The Quarterly Journal of Economics 129 (3), 1085–1140.

Alves, G., 2021. Slum growth in Brazilian cities. Journal of Urban Economics 122, 103327.

Arribas-Bel, D., Garcia-López, M.À., Viladecans-Marsal, E., 2021. Building(s and) cities: delineating urban areas with a machine learning algorithm. Journal of Urban Economics 125, 103217.

Asher, S., Jha, K., Novosad, P., Adukia, A., Tan, B., 2024. Residential Segregation and Unequal Access to Local Public Services in India: Evidence from 1.5m Neighborhoods.

Asher, S., Novosad, P., 2020. Rural roads and local economic development. American Economic Review 110 (3), 797–823.

Balboni, C., Bryan, G., Morten, M., Siddiqi, B., 2020. Transportation, Gentrification, and Urban Mobility: the Inequality Effects of Place-Based Policies. Working Paper.

Bandiera, O., Elsayed, A., Heil, A., Smurra, A., 2022. Economic Development and the Organization of Labour: Evidence from the Jobs of the World Project.

Banerjee, A., Duflo, E., 2005. Growth theory through the lens of development economics. In: Handbook of Economic Growth, pp. 473–552.

Banerjee, A., Duflo, E., Sharma, G., 2021. Long-term effects of the targeting the ultra poor program. American Economic Review: Insights 3 (4), 471–486.

Bank, W., 2009. Reshaping Economic Geography. Number 2009 in World Development Report. World Bank, Washington, D.C.

Baseler, T., 2023. Hidden income and the perceived returns to migration: experimental evidence from Kenya. American Economic Journal: Applied Economics 15 (4), 321–352.

Baseler, T., Narayan, A., Ng, O., Sinha Roy, S., 2024. Social Welfare Portability and Migration: Evidence from India's Public Distribution System. Working Paper.

Bazzi, S., 2017. Wealth heterogeneity and the income elasticity of migration. American Economic Journal: Applied Economics 9 (2), 219–255.

Beaman, L.A., 2012. Social networks and the dynamics of labour market outcomes: evidence from refugees resettled in the U.S. The Review of Economic Studies 79 (1), 128–161.

Bird, J., Venables, T., 2020. The efficiency of land-use in a developing city: Traditional vs modern tenure systems in Kampala, Uganda. Journal of Urban Economics 119, 103268.

Bobba, M., Flabbi, L., Levy, S., Bobba, M., Flabbi, L., Levy, S., 2022. Labor market search, informality and schooling investments. International Economic Review, 1–67.

Boserup, E., 1975. The impact of population growth on agricultural output. The Quarterly Journal of Economics 89 (2), 257–270.

Breza, E., Kaur, S., Shamdasani, Y., 2021. Labor rationing. American Economic Review.

Bryan, G., Chowdhury, S., Mobarak, A.M., 2014. Under-investment in a profitable technology: the case of seasonal migration in Bangladesh. Econometrica 82 (5), 1671–1748.

Bryan, G., Frye, K., Morten, M., 2025. Urbanization in the Developing World: Measurement and Decomposition. Working Paper.

Bryan, G., Morten, M., 2019. The aggregate productivity effects of internal migration: evidence from Indonesia. Journal of Political Economy 127 (5), 2229–2268.

Combes, P.-P., Gorin, C., Nakamura, S., Roberts, M., Stewart, B., 2023. An Anatomy of Urbanization in Sub-Saharan Africa.
Comin, D., Lashkari, D., Mestieri, M., 2021. Structural change with long-run income and price effects. Econometrica 89 (1), 311–374.
Conte, B., 2022. Climate change and migration: the case of Africa. SSRN Electronic Journal.
Cruz, J.-L., Rossi-Hansberg, E., 2024. The economic geography of global warming. The Review of Economic Studies 91 (2), 899–939.
Cutler, D.M., Glaeser, E.L., 1997. Are ghettos good or bad? The Quarterly Journal of Economics 112 (3), 827–872.
de Bellefon, M.-P., Combes, P.-P., Duranton, G., Gobillon, L., Gorin, C., 2021. Delineating urban areas using building density. Journal of Urban Economics 125, 103226.
de Janvry, A., Emerick, K., Gonzalez-Navarro, M., Sadoulet, E., 2015. Delinking land rights from land use: certification and migration in Mexico. American Economic Review 105 (10), 3125–3149.
De Soto, H., 2000. The Mystery of Capital: Why Capitalism Triumphs in the West and Fails Everywhere Else. Basic Books.
Diamond, R., 2016. The determinants and welfare implications of US workers' diverging location choices by skill: 1980-2000. American Economic Review 106 (3), 479–524.
Dijkstra, L., Florczyk, A.J., Freire, S., Kemper, T., Melchiorri, M., Pesaresi, M., Schiavina, M., 2021. Applying the degree of urbanisation to the globe: a new harmonised definition reveals a different picture of global urbanisation. Journal of Urban Economics 125, 103312.
Donaldson, D., Hornbeck, R., 2016. Railroads and American economic growth: a "market access" approach. The Quarterly Journal of Economics.
Duranton, G., 2021. Classifying locations and delineating space: an introduction. Journal of Urban Economics 125, 103353.
Duranton, G., Puga, D., 2004. Micro-foundations of urban agglomeration economies. In: Henderson, J.V., Thisse, J.-F. (Eds.), Handbook of Regional and Urban Economics. In: Cities and Geography, vol. 4. Elsevier, pp. 2063–2117.
Eaton, J., Kortum, S., 2002. Technology, geography, and trade. Econometrica 70 (5), 1741–1779.
Eckert, F., Peters, M., 2024. Spatial Structural Change. Working Paper.
Fajgelbaum, P., Redding, S.J., 2022. Trade, structural transformation and development: evidence from Argentina 1869-1914. Journal of Political Economy 130.
Feler, L., Henderson, J.V., 2011. Exclusionary policies in urban development: under-servicing migrant households in Brazilian cities. Journal of Urban Economics 69 (3), 253–272.
Field, E., 2007. Entitled to work: urban property rights and labor supply in Peru. The Quarterly Journal of Economics 4 (122), 1561–1602.
Foster, A.D., Rosenzweig, M.R., 2022. Are there too many farms in the world? Labor-market transaction costs, machine capacities and optimal farm size. Journal of Political Economy.
Frankenhoff, C.A., 1967. Elements of an economic model for slums in a developing economy. Economic Development and Cultural Change 16 (1), 27–36.
Franklin, S., Imbert, C., Abebe, G., Mejia-Mantilla, C., 2024. Urban public works in spatial equilibrium: experimental evidence from Ethiopia. American Economic Review 114 (5).
Galiani, S., Schargrodsky, E., 2010. Property rights for the poor: effects of land titling. Journal of Public Economics 94 (9–10), 700–729.
Gechter, M., Tsivanidis, N., 2023. Spatial Spillovers from High-Rise Developments: Evidence from the Mumbai Mills.
Glaeser, E., Henderson, J.V., 2017. Urban economics for the developing world: an introduction. Journal of Urban Economics 98, 1–5.
Glaeser, E.L., Gyourko, J., Saks, R., 2005. Why is Manhattan so expensive? Regulation and the rise in housing prices. The Journal of Law & Economics 48 (2), 331–369.
Glaeser, E.L., Maré, D.C., 2001. Cities and skills. Journal of Labor Economics 19 (2), 316–342.
Gollin, D., Kaboski, J.P., 2023. New views of structural transformation: insights from recent literature. Oxford Development Studies 51 (4), 339–361.

Gollin, D., Kirchberger, M., Lagakos, D., 2021. Do urban wage premia reflect lower amenities? Evidence from Africa. Journal of Urban Economics 121, 103301.
Gollin, D., Lagakos, D., Waugh, M., 2014. The agricultural productivity gap in developing countries. The Quarterly Journal of Economics 129 (2), 939–993.
Guedes, R., Iachan, F.S., Sant'Anna, M., 2023. Housing supply in the presence of informality. Regional Science and Urban Economics 99, 103875.
Harari, M., Wong, M., 2024. Slum upgrading and long-run urban development: evidence from Indonesia. The Review of Economic Studies.
Harris, J., Todaro, M., 1970. Migration, unemployment and development: a two-sector analysis. American Economic Review 60 (1), 126–142.
Head, K., Mayer, T., 2013. Gravity equations: workhorse, toolkit, and cookbook. In: Handbook of International Economics, p. 63.
Henderson, J.V., 1975. Congestion and optimum city size. Journal of Urban Economics 2 (1), 48–62.
Henderson, J.V., Nigmatulina, D., Kriticos, S., 2021a. Measuring urban economic density. Journal of Urban Economics 125, 103188.
Henderson, J.V., Regan, T., Venables, A.J., 2021b. Building the city: from slums to a modern Metropolis. The Review of Economic Studies 88 (3), 1157–1192.
Henderson, J.V., Turner, M.A., 2020. Urbanization in the developing world: too early or too slow? The Journal of Economic Perspectives 34 (Summer), 150–173.
Hicks, J.H., Kleemans, M., Li, N., Miguel, E., 2021. Reevaluating agricultural productivity gaps with longitudinal microdata. Journal of the European Economic Association 19 (3), 1522–1555.
Hsiao, A., 2023. Educational Investment in Spatial Equilibrium.
Hsieh, C.-T., Hurst, E., Jones, C., Klenow, P.J., 2019. The allocation of talent and US economic growth. Econometrica 87 (5), 1439–1474.
Imbert, C., Papp, J., 2020. Short-term migration, rural public works, and urban labor markets: evidence from India. Journal of the European Economic Association 18 (2), 927–963.
Imbert, C., Seror, M., Zhang, Y., Zylberberg, Yanos, Zylberberg, Y., 2022. Migrants and firms: evidence from China. American Economic Review 112 (June 2022), 1885–1914.
Jensen, R., 2012. Do labor market opportunities affect young women's work and family decisions? Experimental evidence from India. The Quarterly Journal of Economics 127 (2), 753–792.
Khanna, G., 2023. Large-scale education reform in general equilibrium: regression discontinuity evidence from India. Journal of Political Economy 131 (2), 549–591.
Kreindler, G.E., 2024. Peak-hour road congestion pricing: experimental evidence and equilibrium implications. Econometrica 92 (4), 1233–1268.
Kreindler, G.E., Miyauchi, Y., 2023. Measuring commuting and economic activity inside cities with cell phone records. Review of Economics and Statistics 105 (4), 899–909.
Kumar, A., Barrett, F., 2008. Stuck in traffic: Urban transport in Africa.
Lagakos, D., Marshall, S., Mobarak, A.M., Vernot, C., Waugh, M.E., 2020. Migration policy and observational returns to rural-urban migration in the developing world. Journal of Monetary Economics 113 (C), 138–154.
Lagakos, D., Mobarak, M., Waugh, M., 2023. The welfare effects of encouraging rural-urban migration. Econometrica 91 (3), 803–837.
Lall, S.V., Henderson, J.V., Venables, A.J., Avner, P., 2017. Africa's Cities: Opening Doors to the World. World Bank, Washington, DC.
Lee, K., Miguel, E., Wolfram, C., 2020. Experimental evidence on the economics of rural electrification. Journal of Political Economy 128 (4), 1523–1565.
Lewis, W.A., 1954. Economic development with unlimited supplies of labour. Manchester School 22 (2), 139–191.
Lipton, M., 1977. Why Poor People Stay Poor: Urban Bias in World Development, first edition. Harvard University Press, London.
Marx, B., Stoker, T., Suri, T., 2013. The economics of slums in the developing world. The Journal of Economic Perspectives 27 (4), 187–210.

Marx, B., Stoker, T., Suri, T., 2019. There is no free house: ethnic patronage and diversity in a Kenyan slum. American Economic Journal: Applied Economics 11 (4), 36–70.

McFadden, D., 1974. Conditional logit analysis of qualitative choice behavior. In: Frontiers in Econometrics. Academic Press, New York.

Meghir, C., Mobarak, A.M., Mommaerts, C., Morten, M., 2020. Migration and informal insurance. The Review of Economic Studies.

Meghir, C., Narita, R., Robin, J.-M., 2015. Wages and informality in developing countries. American Economic Review 105 (4), 1509–1546.

Morten, M., 2019. Temporary migration and endogenous risk sharing in village India. Journal of Political Economy 127 (1), 1–46.

Morten, M., Oliveira, J., 2024. The effects of roads on trade and migration: evidence from a planned capital city. American Economic Journal: Applied Economics 16 (2), 389–421.

Munshi, K., Rosenzweig, M., 2016. Networks and misallocation: insurance, migration, and the rural-urban wage gap. American Economic Review 106 (1), 46–98.

National Bureau of Statistics (NBS), 2019. Living Standards Survey 2018-2019.

Panman, A., Lozano Gracia, N., 2022. Titling and beyond: evidence from Dar es Salaam, Tanzania. Land Use Policy 117, 105905.

Redding, S., Rossi-Hansberg, E., 2017. Quantitative spatial economics. Annual Review of Economics 9, 21–58.

Redding, S.J., 2025. Quantitative urban models. In: Donaldson, D., Redding, S.J. (Eds.), Handbook of Regional and Urban Economics, vol. 6. Elsevier North Holland, Amsterdam.

Rubin, D., 1980. Randomization analysis of experimental data: the Fisher randomization test comment. Journal of the American Statistical Association 75 (371).

Ruggles, S., Cleveland, L., Lovaton, R., Sarkar, S., Sobek, M., Burk, D., Ehrlich, D., Heimann, Q., 2024. Integrated Public Use Microdata Series. International.

Small, K.A., Winston, C., Yan, J., 2005. Uncovering the distribution of motorists' preferences for travel time and reliability. Econometrica 73 (4), 1367–1382.

Tang, Y., 2024. To Follow the Crowd? Benefits and Costs of Migrant Networks.

Todd, P., Wolpin, K., 2006. Assessing the impact of a school subsidy program in Mexico: using a social experiment to validate a dynam fertility. American Economic Review 96 (5), 1384–1417.

Udry, C., 2024. The 33rd Kuznets Memorial Lecture: Structural Change and Declining Agricultural Productivity: Evidence from Sub-Saharan Africa.

Ulyssea, G., 2020. Informality: causes and consequences for development. Annual Review of Economics 12 (1), 525–546.

UN-Habitat, 2002. Expert Group Meeting on Urban Indicators. Technical report.

United Nations Human Settlements Programme (Ed.), 2003. The Challenge of Slums: Global Report on Human Settlements, 2003. Earthscan Publications, London; Sterling, VA.

Young, A., 2013. Inequality, the urban-rural gap, and migration. The Quarterly Journal of Economics 128 (4), 1727–1785.

Zárate, R.D., 2023. Spatial Misallocation, Informality, and Transit Improvements: Evidence from Mexico City.