

Urban-Rural Gaps in the Developing World: Does Internal Migration Offer Opportunities?

David Lagakos

One prominent feature of virtually every developing country is an enormous divide between rural and urban living standards, measured by income, consumption, or various nonmonetary aspects of life. As a result, much of the inequality within the developing world—home to about half of the planet’s nearly 8 billion people—is accounted for by the urban-rural gap.

To illustrate, Table 1 presents a number of comparisons of rural and urban living standards for those residing in Nigeria and India—the most populous countries in Africa and South Asia, respectively—drawing on real outcomes measured from the Demographic and Health Surveys (DHS). In Nigerian and Indian villages, the floor in your home would most likely be made of dirt; in urban areas, floors are most commonly made of wood or stone. About one-half of rural Indians and one-third of rural Nigerians have no toilet facility of any kind—not even a pit latrine or composting toilet—while virtually all urban residents have one, however rudimentary. Fewer than four in ten rural Nigerians can point to a power outlet inside their home, compared with eight out of ten urbanites. Rural Indians similarly lag behind their urban counterparts in electricity connections. In both countries, television ownership rates in cities are about twice as high as in rural areas.

Similar patterns emerge when looking at mortality rates and other health metrics. In both Nigeria and India, you would be just over half as likely to perish before your fifth birthday in a city than in a village. Among adults, a body-mass index

■ *David Lagakos is Associate Professor of Economics, University of California San Diego, La Jolla, California. He is also a Research Associate, National Bureau of Economic Research, Cambridge, Massachusetts. His email address is dlagakos@ucsd.edu.*

For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at <https://doi.org/10.1257/jep.34.3.174>.

Table 1
Real Urban and Rural Living Standards in India and Nigeria

	<i>Urban</i>	<i>Rural</i>
Percent with finished floors		
India:	70.4	40.3
Nigeria:	88.1	60.8
Percent with toilet facility		
India:	89.5	45.9
Nigeria:	84.6	67.5
Percent with electricity		
India:	97.5	83.2
Nigeria:	82.7	38.9
Percent owning a television		
India:	87.0	53.5
Nigeria:	70.7	30.0
Under-five mortality (per 1,000 births)		
India:	36	59
Nigeria:	86	155
Percent with BMI below 18.5		
India:	15.5	26.8
Nigeria:	9.6	14.4

Note: Compiled from the Demographic and Health Surveys, funded by the US Association for International Development and publicly available at <https://dhsprogram.com/>. The statistics are calculated in the most recent year available, which is most commonly 2018.

of less than 18.5 is a commonly used indicator of serious malnutrition. For example, someone who is 5' 8" tall (172 centimeters) would have to weigh just 122 pounds (55 kilograms) to have a BMI of 18.5. In both countries your chance of having such a low BMI would be about 50 percent higher in rural areas than in cities.

Of course, well-being depends on a host of other factors other than the few presented here, and some of these are harder to measure. A full accounting of urban-rural differences might also take into consideration the value (positive or negative) placed on the hustle and bustle of urban centers or the security of traditional kinship ties in rural villages. However, the observable differences are so large that it is hard to believe that rural-urban gaps are simply artifacts of inaccurate measurement. The real issue is how to interpret these gaps and whether policy should try to do anything about them.

In this essay, I first set the stage by offering some more systematic evidence on the size and prevalence of the urban-rural gap from a variety of recent data sources. I then discuss whether the urban-rural gap can be explained by sorting; after all, it is clearly the case that those in urban areas tend to have more education, and they are probably selected on other less observable abilities as well. In addition, I review an array of evidence on outcomes of worker migration and find that rural-to-urban migrants do typically obtain higher incomes, which suggests that the pre-migration

situation cannot entirely be due to efficient sorting by skill and education. In addition, while rural-to-urban migrants do experience gains, migration itself does not close the urban-rural gap. Thus, I turn to other potential explanations. I emphasize that a variety of frictions—information, financial, and in land markets—may help to explain the persistence of urban-rural gaps, though much more work is needed here.¹

Throughout time and space, economic development has been associated with massive population movements from rural to urban areas, and from agriculture to non-agricultural activities. A main reason urban-rural gaps are worthy of study is that they may be informative about this process of structural transformation. Some writers have pointed to the large urban-rural gaps as suggesting the possibility of massive productivity gains from helping move workers in developing countries out of rural agriculture (Caselli 2005; Restuccia, Yang, and Zhu 2008; Vollrath 2009; McMillan, Rodrik, and Verduzco-Gallo 2014). I offer a more restrained conclusion. Although some rural-to-urban migrants already make gains, many choose not to migrate for a variety of reasons. Some of these reasons may be efficient in the sense that a benevolent social planner wouldn't want them to do any differently. However, others may be held back inefficiently from rural-to-urban migration, and policy may play an important role in reducing the frictions that keep them from the higher living standards that cities would offer them.

Urban-Rural Gaps: Recent Systematic Evidence

Until recently, most of the evidence on urban-rural gaps has come from nominal income or consumption expenditure data by region that has been deflated using spatial price indices. However, price deflation is never quite as straightforward as the breezy descriptions that we (or at least I) offer in undergraduate macroeconomics lectures. The main challenge is that many goods and services are not that easy to compare between cities and rural villages. For example, an urban apartment will carry a higher monthly rent than a straw hut of comparable size in a remote area, but will involve a number of quality differences including access to electricity and running water. This makes apples-to-apples comparisons difficult. Prices for food items are also tricky to compare, because such a large fraction of food in rural areas is home-produced (Deaton and Dupriez 2011).

Ravallion, Chen, and Sangraula (2007) construct arguably the best rural-urban price deflators available for a large set of low-income countries, drawing on spatial

¹For engaging overviews of urban economics in the developing world, the interested reader might begin with the review articles by Bryan, Glaeser, and Tsivanidis (2019) and Brueckner and Lall (2015) and the book by Glaeser (2017), which helps put today's developing-world cities in a broader historical context. While rural-urban migration will be discussed below, I will focus mostly on how it informs us about the sources of the urban-rural gaps. For a more general review of the literature on rural-urban migration, a useful starting point is the article by Lucas (2015). Similarly, for more on the role of migration and rural-urban gaps in dual economy models in economics, the essay by Gollin (2014) is essential reading.

price data collected by the World Bank. On average, they find that the “same basket,” to the extent that such a thing can be measured, costs around 30 percent more in cities than in rural areas. They argue that these price differentials are not nearly enough to offset the higher poverty rates of rural areas that are implied by nominal consumption expenditures. This finding of higher real poverty rates in rural areas is echoed in the work of Ferré, Ferreira, and Lanjouw (2012), among many others, all of whom grapple with similar challenges in making spatial price adjustments.

The pioneering paper by Young (2013) offers an alternative approach to measuring real rural-urban consumption gaps in the developing world that sidesteps the need for spatial price indices. The methodology he develops, while less transparent than approaches based on regional price deflation, allows him to estimate a single proxy for “consumption” per head in rural and urban areas of 67 developing countries using the data from the Demographic and Health Surveys discussed above. The idea is to infer household consumption levels using cross-sectional correlations between educational attainment and the consumption of each “good” measured in the data (Young 2012, 2013). The “goods” are actually 23 real outcomes, including ownership of durable goods (like televisions), housing conditions (like access to electricity), and children’s health outcomes. Thus, Young’s measure of consumption is broader in some ways than typical national accounts measures of consumption, but it covers a narrower set of goods. Across Young’s (2013) set of developing countries, average consumption per household is 4.5 times as high in urban areas as in rural areas. In addition, Young calculates that 40 percent of total inequality is accounted for by the urban-rural gaps themselves. The share of inequality explained by the urban-rural gap is higher in countries with more inequality. For example, in Zimbabwe, one of most unequal countries in Young’s data, the urban-rural gap explains close to 70 percent of all inequality.

The urban-rural gaps in consumption and income are related—though clearly not identical—to the gaps in income per head between agricultural and non-agricultural workers. Researchers who focus on agricultural and non-agricultural incomes can make use of national accounts data on value-added by sector, which is widely available. Gollin, Lagakos, and Waugh (2014) draw on sectoral value added data for 151 counties (of all income levels) and construct employment by sector using nationally representative surveys. On average, they find that value added per worker is 3.5 times as high in the non-agricultural sector as in agriculture. This ratio is well above one in all but a handful of countries in their data.

One key concern with sector-level national accounts data is that while the guidelines behind the UN Systems of National Accounts clearly include non-market production of goods (including agricultural goods) as part of value added, in practice such output may be underestimated, as Gollin, Parente, and Rogerson (2004) hypothesize. Indeed, Herrendorf and Schoellman (2015) document this pattern in US agricultural output data. Gollin, Lagakos, and Waugh (2014) seek to address this issue by constructing value added using household income and consumption data for ten countries from the World Bank’s Living Standards Measurement Surveys (LSMS), which contain detailed questions about agricultural production both

for own consumption or to be sold in a market. Such measures continue to show large gaps in value added per worker between agricultural and non-agricultural workers. For example, the urban-rural gap in Ghana is 2.2 according to the national income accounts and 2.3 using the household surveys. Cote d'Ivoire has a gap of 4.7 according to the national income accounts and 4 using the household surveys. These multiples are not identical, to be sure, but recalculating based on household survey data continues to leave a large gap.

To circumvent the national accounts data altogether, Herrendorf and Schoellman (2018) focus on wage data by sector, where wages are defined as labor income divided by hours worked. Their data come from nationally representative population censuses compiled by the Integrated Public Use Microdata Series (IPUMS). On average, the simple ratio of non-agricultural average wages to agricultural average is around 1.8 in the 13 countries they study, with a range of 1.5 to 2.7, and gaps of a similar magnitude when splitting non-agriculture into industry and services. These gaps are not as large as some calculated by other methods, but in none of their countries does a simple comparison of average wages yield anything close to parity between the agriculture and non-agricultural sectors.

How do the urban-rural gaps in developing countries compare to those of high-income countries? The limited available evidence points in the direction of larger gaps in developing countries, but this evidence is far from definitive. Drawing on rich regional-level data, Chauvin et al. (2017) document that in China and India, a doubling of population density is associated with an increase in average real wages of around 6 percent. In other words, more urbanized regions offer substantial wage premia relative to more rural areas. In the United States, a doubling of density leads real wages to increase by only about 2 percent, consistent with a smaller urban-rural divide. On the other hand, in Brazil, the wage-density gradient is even smaller at around 1 percent. Looking at agricultural productivity gaps, Gollin, Lagakos, and Waugh (2014) find that the median gap is 4.3 among countries in the bottom quartile of the world income distribution, compared to 1.7 in the top quartile, while the mean gap is 5.6 in the bottom quartile and 2.0 in the top quartile. Thus, these gaps are substantially larger on average in poorer countries. Still, there is a lot of variation within each income quartile, and many high-income countries exhibit very large gaps as well.

A related question is whether urban-rural gaps tend to decrease over time as economies develop and populations urbanize. While many economists have the prior that they do, the historical evidence is thin and some studies have pointed to quite different patterns. For example, Hatton and Williamson (1992) show that the rural-urban gap actually widened in the United States from the late nineteenth century up until right before World War II, and Lundh and Prado (2015) find that the gap in Sweden changed very little over the twentieth century. Hnatkovska and Lahiri (2018) document that, since the 1980s, urban-rural wage gaps have widened in China and decreased substantially in India. A valuable task for future research would be to look systematically at how these gaps have evolved over long time horizons in a broad set of countries.

Sorting and Evidence from Internal Migrants

Urban households differ from rural households in many ways: they perform different jobs, possess different skills, and as one would expect, receive different compensation levels. To what extent can the urban-rural gaps in average incomes be explained by a situation with more productive individuals sorting into cities and less productive types ending up mostly in rural areas? Suppose, for argument's sake, that the answer is 100 percent. Then there would be a sense in which the urban-rural gap would not be policy relevant. In the same way that dentists earn more on average than Uber drivers, urbanites may just be different people from villagers. We should not then be trying to turn villagers into urbanites any more than we should be encouraging Uber drivers to earn their living filling cavities.

In this section, I first discuss some models and quantitative analyses of sorting and then turn to evidence from long-term panel data and experimental studies of internal migration. The theory of sorting makes a strong case that the gains from expanding rural-to-urban migration will be much smaller than a naïve interpretation of the urban-rural gap might suggest because average levels of education and skill are lower in rural areas. However, I will also argue that the reduced-form evidence on sorting can be hard to interpret because these data mix together migrants with many varying motivations for migration—for example, the fact that some migrants are choosing a move toward opportunity while others are making a forced move of necessity—and do not take into account the many varying motivations for non-migration. Overall, sorting alone does not seem to explain the ongoing urban-rural gaps in income and standard of living, though it certainly explains part of it.

A Sorting Approach to Urban-Rural Gaps

At least in cross-sectional data, the signs that there is urban-rural sorting based on observable variables like education are overwhelming. As one example, Gollin, Lagakos, and Waugh (2014) measure average years of schooling among non-agricultural and agricultural workers in 124 countries using census data from IPUMS and household survey data from various other sources. In literally every country, non-agricultural workers have higher average schooling levels, and on average, they have almost twice as many years of schooling than those working in agriculture. The same basic finding shows up in every other cross-sectional comparison of which I am aware, whether looking at non-agriculture versus agriculture or urban versus rural. Taking a historical perspective, Porzio and Santangelo (2019) argue that the rise of schooling globally over the twentieth century was one of the main factors behind the large movement of workers into the non-agricultural sector observed in most countries over this period.

Taking this a step further, Young (2013) looks not only at where workers are currently located, but where they were raised as children. He finds that rural-to-urban migrants have substantially higher education levels than those who were raised in rural areas and stayed there. Urban-to-rural migrants, in contrast, have

much lower education levels than those born in urban areas who remained there. His data cover 170 surveys from a diverse set of developing countries, and thus emphasize that sorting on education is a basic fact of life in the developing world. Theoretical models have been built on this insight; for example, the Lucas (2004) model of rural-urban migration and long-run economic growth begins with the premise that education is valuable only in urban areas.

However, the obvious question is the importance of this sorting in accounting for the urban-rural gaps, and here the literature remains divided to some extent. For example, Gollin, Lagakos, and Waugh (2014) take a basic approach and convert years of schooling into “units of human capital” using an off-the-shelf estimate of 10 percent (Mincerian) return to a year of schooling. Their approach follows the literature on development accounting that does more or less the same thing but in a cross-country setting (for example, see Hsieh and Klenow 2010). The resulting human capital stocks estimated by Gollin, Lagakos, and Waugh (2014) are around 40 percent higher for non-agricultural workers, which explains only a modest fraction of the overall gaps. Country-specific returns and attempts to adjust for schooling quality explain a bit more, but still leave large residual gaps.

Herrendorf and Schoellman (2018) draw on their individual-level wage data to estimate the Mincerian returns to education and experience by sector for their set of 13 countries. They find smaller returns to education among agricultural workers than for workers in other sectors, which squares with the common perception that education is not that useful in farming. Once they add controls, including for education, they find that the raw wage gaps of 1.8 in the median country fall to 1.3. Vollrath (2014) takes a similar approach but by using data from the Living Standards Measurement Surveys for 14 developing countries. His approach includes controls for age and age-squared (to capture experience effects), occupation, and, in one specification, occupation-specific returns to education. His extra controls get the gaps to fall further, though his estimated wages are still lower in agriculture even with the most stringent set of controls.

Of course, economists are fully aware that the measurable variable of education is likely to be highly correlated with a number of unobservable variables, like cognitive abilities or whether a family offers support for education. Furthermore, the importance of sorting is likely to be more important among educated workers than for those performing unskilled manual tasks, among which skill heterogeneity is probably not that substantial. When developing a model for putting education in the context of a broader group of observed and unobservable variables, one sooner or later arrives at the Roy (1951) model of occupational choice. In brief, the Roy model posits that each worker has a vector of “skills”—one for each occupation—rather than a single skill level. Workers then sort into the occupation that yields them the highest income level.

The most difficult challenge with Roy-style models of sorting is that one rarely observes the paths not chosen—that is, the life that rural-urban migrants would have led had they stayed in the village. One approach, first taken in this literature by Lagakos and Waugh (2013), is to draw on more structure to make inferences,

using parameterized versions of the Roy model calibrated to data on wage distributions by sector. The papers of Young (2013) and Herrendorf and Schoellman (2018) introduce richer Roy models that combine sorting on observables, in particular, education as well as unobservable ability. The key theme in these models is that the urban sector is more skill-intensive. Workers are endowed with education, but education is not a skill in and of itself and only increases the probability of becoming skilled. Those becoming skilled stay and work in the skill-intensive urban sector. Those born in the rural area but who become skilled—even with their lower education level—migrate to the city to work in the urban sector. The reverse is true for those not becoming skilled, who find their way to the rural area to work in the unskilled-intensive sector.

The theory is consistent with the clear sorting on education by region in the data, with the more educated primarily locating in cities. Because education and ability are positively correlated, those with higher ability are also primarily located in cities. In the models presented in these papers, the higher average education and ability of workers in cities combine quantitatively to account for nearly all of the urban-rural gap, leaving little or no role for any other explanatory factor.

Compelling as these models may be, there is only so far one can go with sorting stories by looking at cross-sectional data, with one wage outcome per individual (Heckman and Honore 1990). What is more informative, though still no panacea, is to draw on detailed panel data to observe what actually happens to those that move between rural and urban areas.

Panel Data on Internal Migration

Hicks et al. (2017) carry out such an exercise using two long and large household panel surveys: the Indonesian Family Life Survey and the Kenyan Life Panel Survey. What makes these surveys so attractive, besides their length and large sample sizes, is that they make a serious effort to track every respondent that moves between survey waves. Such tracking is not easy: many migrants leave behind little trace of their whereabouts, requiring survey enumerators to do some real detective work.

The punchline of their study is that the gaps between urban and rural areas are far larger than the changes in income and consumption experienced by those moving from rural to urban areas (or from agriculture into a non-agricultural job). In Indonesia, for example, the urban-rural earnings gap is about 1.7 when calculated without any other controls—roughly in line with the cross-sectional estimates described above. But once individual fixed effects are included, the urban coefficient falls nearly to zero. In Kenya, the cross-sectional urban-rural gap in earnings is around 2.4 without any controls, but this falls by two-thirds with individual fixed effects. The results for non-agricultural worker status yield different numbers but the same conclusion: cross-sectional gaps are cut down dramatically once individual fixed effects are added to the regressions.²

²Hicks et al. (2017) find that cognitive ability scores, measured using Raven's Progressive Matrices, are around 0.3 standard deviations higher in both countries among migrants than for those who remain

More recent evidence from other studies and from countries other than Indonesia and Kenya has tended to find similar patterns; that is, the observed gains from rural-to-urban migration are much smaller than urban-rural cross-sectional gaps. Using detailed wage data from a large survey in Brazil, Alvarez (2020) finds that sectoral movers gain a lot less than one would naively expect given the large cross-sectional wage gap in Brazil. Non-agricultural workers in the cross-section earn a premium of around 62 percent relative to agricultural workers. Once individual fixed effects are included, the premium is a just 9 percent for manufacturing and a paltry 4 percent for services. He concludes that sorting on observables explains close to the entire Brazilian cross-sectional gap.

Lagakos et al. (forthcoming) follow suit by looking at the returns to migration for rural-urban migrants in China, Ghana, Indonesia, Malawi, South Africa, and Tanzania. The surveys they draw on are nationally representative panels and also make substantial efforts to track migrants across space. Like Hicks et al. (2017) and Alvarez (2020), these panel data confirm the substantially smaller returns to rural-urban migration for those choosing to migrate than the cross-sectional gaps. However, their estimated returns are not near-zero after controls for individual fixed effects, and instead, average a substantial 25 percent across their countries. Lagakos et al. (forthcoming) show that their larger average estimated returns come from their different set of countries, not differences in methodology. Their estimated return for Indonesia, also studied by Hicks et al., is the smallest of their six countries.

There are other earlier studies from developing countries that estimate the gains from migrating from urban to rural areas, though these tend to have smaller sample sizes and not to be nationally representative. These earlier studies have generally found substantial returns to migration for those observed to migrate. For example, in a study of 772 rural Indian individuals that were surveyed in 1975 and again in 2005, Dercon, Krishnan, and Krutikova (2013) find consumption per capita was 42 percent higher for those that migrated since the first survey than for those that stayed put. Using panel tracking data from northern Tanzania, Beegle, De Weerd, and Dercon (2011) find that among 912 households surveyed, those moving out of the community had 36 percent higher consumption levels than those that remained behind after controlling for education, age, and other co-variates. Individuals that moved further away tended to have larger consumption gains.

in rural areas. They cite this, convincingly, as direct evidence that rural-urban migrants are positively selected on ability, in addition to education. While the topic of international migration is beyond the scope of the current essay, it is worth pointing out that the literature on international migration has also found strong evidence of positive selection into migration, both on observables and unobservables. For example, McKenzie, Stillman, and Gibson (2010) provide experimental evidence that the large wage gains for Tongan migrants to New Zealand are largely due to selection on who chooses to apply for a migration lottery. For the much larger set of Mexican migrants to the United States, Chiquiar and Hanson (2005) document strong positive selection on education relative to Mexican non-migrants.

Difficulties in Interpreting Observational Returns to Migration

One might be tempted to conclude that because the gains from migration are substantially smaller than the cross-sectional gaps, there is little scope for policy aimed at encouraging rural-urban migration. Yet this conclusion should be proffered with some skepticism. The “observational returns to migration” estimated from non-experimental panel data require some care to interpret and do not translate as easily as one might think into lessons about the effects of incentivizing others to migrate internally. In general, the concerns relate to the non-random nature of who migrates and to the fact that many people do not actually migrate.

First, worker heterogeneity may extend to migration costs, not just to the migration benefits (as posited in most Roy-style models). Individuals who migrate in equilibrium may be those with relatively low costs and low benefits of migrating, as might be the case for one whose village is close to a major urban center and connected to it by a high quality road. Conversely, those who do not migrate because of high costs—even if they might also experience large potential gains—will never help identify the urban coefficient in a regression that relies only on migrants for identification (Lagakos et al. forthcoming). There is clearly more work to be done to improve our understanding of how migration costs differ across individuals, as opposed to just the returns to migration.

Second, migrating workers who switch sectors may do so for reasons other than choosing the best sector for themselves in a permanent sense. This possibility is consistent with the findings of Pulido and Świącki (2018) who, using the same Indonesian panel data as Hicks et al. (2017), find that around one in five of the movers from the non-agriculture to agriculture sectors describe the shift as “forced” rather than voluntary—for example, when the employer was closed or relocated, or the worker’s job was relocated. Those forced to move sectors due to job loss on average witness substantial wage loss.

The underlying problem is that in an observational study, one never really knows what motivates a worker to migrate or not. Once one is “assigned” to migrate using some controlled (or at least well understood) external incentive to migrate, some of the inferences may become clearer. For this reason, many researchers have turned to experimental and quasi-experimental approaches to measuring the returns to internal migration.

Experimental and Quasi-Experimental Returns to Migration

The ideal experiment would be to induce some rural farmers in a developing country to permanently move to urban areas and to observe them and their non-migrant counterparts (plus all of their offspring, while we are at it) for the rest of their lives. But most people aren’t likely to move permanently away from their homes in exchange for a modest payment from some experimenting economist. Temporary moves may prove somewhat more feasible to induce, particularly during times when opportunities at home are poor.

In a first-of-its-kind experiment, Bryan, Chowdhury, and Mobarak (2014) tried—quite successfully as it turns out—to induce rural Bangladeshi households to

send migrants to more productive places during the so-called “lean season” between the rice planting and harvest. In the Rangpur region of northern Bangladesh, the lean season brings on a large fall of perhaps one-half in average income, rendering many households so poor as to skip meals. Around one-third of households were already sending out a migrant during the lean season at the time of the experiment, with many going to the urban centers of Chittagong or Dhaka to work as rickshaw drivers, day laborers on construction sites, or some other low-skilled job. Bryan, Chowdhury, and Mobarak (2014) offered households in a randomly selected set of villages an incentive of \$11.50 (equal to a few weeks wages) conditional on sending a migrant in the lean season. This modest sum induced a 22 percentage point increase in migration, raising the fraction of households with a migrant from one-third to above one-half. The households with an additional migrant saw consumption rise by a surprising 30 percent per household member. Those in treatment villages were more likely to migrate even three years after the experiment, though only somewhat more than households in the control villages.

A follow-up experiment by Akram, Chowdhury, and Mobarak (2018), carried out in the same region of Bangladesh, offered a richer set of migration incentives and more comprehensive household surveys. The simplest migration incentives, which were about the same size as those in the original experiment, induced a similarly large number of households to send a migrant. The migrants and their families were contacted (pestered, one might say) every week during the lean season with specific questions about the migrants’ employment, earnings, and the remittances sent back to their relatives in the villages. In a second treatment arm, a larger fraction of households was offered the incentive in some villages, and this randomly selected second group of villages sent even more migrants. The two treatment arms allow Akram, Chowdhury, and Mobarak (2018) to estimate that rural wages rise by 2 percent for every 10 percent increase in the rural out-migration rate as rural workers become scarcer. In both treatment arms, income and consumption were substantially higher in the treatment villages than in the controls.

Since permanent migration is harder to induce, it is useful to study episodes of forced migration in which individuals in a given area were induced to move out of rural agricultural areas by some strong external force. Perhaps the most relevant study for our discussion is by Sarvimäki, Uusitalo, and Jäntti (2019) who analyze the long-term consequences of when Finland ceded a large portion of its eastern region to the Soviet Union after World War II and had to resettle 430,000 people (11 percent of its population). While Finland is a rich country by any metric today, its GDP per capita was under \$5,000 in 1950 (in purchasing power parity terms), and the workforce was mostly agricultural like most developing nations today. A quarter century later in 1971, the groups that were resettled consistently had higher income than comparison groups (like those just on the other side of the border) who had not resettled. The main reason was that being forced to move increased the changes of leaving farming and joining the non-agricultural sector—with its higher wages—by about 50 percent. Interestingly, the children of resettled farmers

also have higher incomes and education than the children of farmers not resettled, pointing to important intergenerational effects of migration.³

While all of the evidence presented in this section is specialized in some way, it reinforces the message that the urban-rural gap cannot be solely about sorting. After all, if the rural-urban gap were all about efficient sorting of better workers into urban areas, then an external force inducing people to migrate out of rural agricultural areas should not lead their wages to rise. However, the reasons why these workers were not migrating more often to begin with—and what accounts for the rest of the gap—are not settled. In the next sections, I offer my perspective on two broad classes of possible explanations: compensating differentials of rural life and migration frictions of various sorts. I argue that the latter is the more promising explanation of the two.

Non-Monetary Amenities of Rural Areas

One can easily imagine that the higher wages of developing world cities reflect a premium for lower non-monetary amenities than in rural areas. In fact, this was almost certainly the case in the “killer cities” of the past, which had much higher death rates than the rural hinterlands (Costa and Kahn 2006; Hanlon and Tian 2015; Jedwab and Vollrath 2016). More generally, other non-monetary amenities of rural life may represent the compensating differentials that underly the “spatial equilibrium” assumption common in urban economics in which households are indifferent on average between locations with high wages and fewer amenities and those with lower wages but more amenities (Glaeser and Gottlieb 2009).

In an attempt to shed light on this hypothesis, Gollin, Kirchberger, and Lagakos (2019) analyze spatial data for 20 African countries covering a select number of non-monetary “amenities” related to public goods, pollution, and crime. Theirs is hardly an exhaustive list of all possible amenities but rather, some of the candidates most commonly discussed. They find that public goods are generally much less common in rural areas, including electricity, piped water, and sewage systems (as highlighted in the Nigeria and India comparisons at the start of this article). Indoor air pollution is clearly worse in rural areas because rural households burn solid fuels such as wood for their cooking, which creates a lot of smoke. The World Health Organization (2014) estimates that around four million people die prematurely each year due to burning solid fuels indoors. Perhaps surprisingly, outdoor air pollution is somewhat worse on average in rural areas in Africa. The rural areas in this study tend to be closer to the Sahara Desert where fine particulate matter

³In a related study, Nakamura, Sigurdsson, and Steinsson (2019) study the after-effects of a volcanic eruption on a rural fishing community in Iceland which destroyed a random selection of houses. Those under 25 at the time of the eruption who migrated had earnings 83 percent higher than those that stayed behind, some of which happened because the movers completed 3.5 more years of schooling. Bazzi et al. (2016) study an episode of forced migration across islands in Indonesia, in which rural farmers were moved from denser to less dense islands. The overall wage gains were not that large, suggesting that moving workers within the agricultural sector may not be a fruitful way to raise overall productivity.

in the air is highest, while African cities have low levels of manufacturing activity given their level of GDP per capita (Gollin, Jedwab, and Vollrath 2016).

Crime is the one area where African cities appear worse on some metrics than their hinterlands, but the differences are not dramatic. In rural areas, 10 percent of respondents reported that they or a household member were physically attacked in the last year; in urban areas, the rate was 12 percent. Rates of theft are modestly higher in urban areas. When asked about whether they ever felt unsafe in their homes, 37 percent of those in rural areas answered in the affirmative compared to 45 percent in urban areas.

Does evidence on nonmonetary amenities from African cities jibe with the situation in South Asia or other parts of the developing world? More systematic evidence is needed here. In the dimension of air pollution, for example, cities in India have some of the worst air quality in the world and appear much worse than rural areas. In terms of crime patterns, some cities like Bangkok and Manila are thought to have much higher crime rates than their rural hinterlands. Yet in Madagascar, Fafchamps and Moser (2003) find that *rural areas* have higher rates of homicide, burglary, and insecurity than do urban households. Other non-monetary amenities that have not been systematically explored in a developing world context include commuting times and sanitation—these could certainly play some role in explaining some of the urban wage premium. Yet I am skeptical that, taken as a whole, they will explain all that much.

Arguably a more promising version of the amenities story is one in which individuals have idiosyncratic tastes for rural and urban amenities, as in the recent work in the urban economics literature (for example, Kline and Moretti 2014). The idea is that some rural individuals (“the country mice”) may optimally choose to remain there even if moving to the city would result in substantial income gains. For example, some rural residents may particularly value living in sparsely populated areas or the bucolic way of life. Such a story can help reconcile the persistence of urban-rural gaps despite income gains for those induced to migrate. The task of separating the idiosyncratic taste shocks from the frictions that hold back migration is a worthwhile—and challenging—job for future research.

Frictions: Information, Financial, and Land Markets

Few economists would dispute the notion that markets in developing countries are full of frictions. A growing body of evidence suggests that some of these frictions may be important factors holding back rural-urban migration and that migration frictions more generally lead to substantially lower aggregate productivity (for example, Bryan and Morten 2018; Tombe and Zhu 2019).⁴

⁴Interestingly, one seemingly obvious type of migration friction—poor road networks—seems to have rather modest effects on internal migration rates and aggregate productivity. At least, this is the conclusion reached by Asher and Novosad (2020); Banerjee, Duflo and Qian (2020); and Morten and Oliveira (2018) in their studies of road building projects in India, China, and Brazil, respectively.

For starters, one wonders how much people in rural parts of the developing world even know about wages and living conditions in distant cities. If most rural residents are not even aware of how much higher wages are in cities, it might help to explain why internal migration rates are not higher. Bryan, Chowdhury, and Mobarak (2014) put this information-frictions theory to the test in one arm of their experimental setting in 16 poor rural villages in northern Bangladesh. Certain households were presented with information on the most common jobs available for seasonal migrants like rickshaw driver, construction worker, or day laborer in four common destinations along with likelihood of finding one of these jobs and average wages. As it turns out, this information treatment had a precise zero effect on migration. In this setting—and remember, around one-third of households were already sending seasonal migrations from this region—it seems natural to assume that households already had sound information about job prospects for seasonal migrants.

However, in a study about migration expectations among rural Kenyans, Baseler (2019) reaches virtually the opposite conclusion. Even though uneducated Kenyans earn twice as much as their rural counterparts, rural workers substantially underestimate the magnitude of this wage gap. To understand why, Baseler (2019) runs an information experiment on a set of villagers where he informs them of the average wages and prices of food in Nairobi and other urban centers, plus the most common jobs for migrants in each potential destination. This simple intervention raises expectations about average urban wages and increases migration to Nairobi, from 20 percent of rural households to 28 percent. Two years later, migration rates were still higher among those getting the information treatment, and migrants reported higher subjective well-being on average.

So why are these Kenyans so poorly informed to begin with? Baseler (2019) theorizes that migrants tend to underreport their earnings in cities to their rural brethren so as not to have to send back too much of their income as remittances. To test his theory, Baseler (2019) runs a second experiment where he spills the beans about all the hidden savings by families with migrants to a random set of others in the villages. This information treatment group responds by updating their expectations about urban wages and their plans to send more migrants in the future. The lesson is that even if out-migration is common in rural areas, the locals still may have imperfect information about wages in the cities. The extent to which information frictions hold back migration in other settings is certainly a topic worth more exploration, especially given how cheap it is to provide information about urban wages and job prospects relative to, say, the cost of providing an additional year of formal education.

Financial frictions of various sorts have long been thought to be important barriers to migration. In particular, financial frictions might bite if migration to cities is a risky enterprise, as emphasized in the classic paper by Harris and Todaro (1970), and potential migrants face borrowing constraints. Bryan et al. (2014) interpret the outcomes of their migration experiments (discussed earlier)

as being about migration risk and borrowing constraints, which keep productive rural people from moving while they do not have a sufficiently large buffer stock of savings to self-insure.

But on further inspection, the explanation that migration risk and credit constraints are what held back migration in this setting doesn't seem quite right.⁵ Lagakos, Mobarak, and Waugh (2019) reinterpret the experimental evidence of Bryan et al. (2014) and argue that the experimental data are more consistent with a model in which rural households generally prefer to be in rural areas and only migrate to cities when they are desperate enough to make it worth their while. In the data, it is the households with *lower* assets and consumption levels that are more likely to seasonally migrate, rather than the other way around. Also, when offered cash that is not tied to migration, few households actually decide to migrate in response. Kleemans (2015) finds a similar result in Indonesia, where rural households—particularly those with low asset levels—send more temporary migrants in response to negative rainfall shocks. Lagakos, Mobarak, and Waugh (2019) use their model to simulate the effects of permanently offering migration subsidies and show that the welfare effects from such a policy come largely through offering better insurance to vulnerable rural households.

Munshi and Rosenzweig (2016) and Morten (2019) make the case that a related type of financial friction—the lack of insurance markets—is a constraint on migration that holds down average productivity levels. The idea is that many people are stuck in rural villages because they lack formal insurance, including public safety-net programs, and thus rely on the informal risk-sharing arrangements prevalent in rural communities (Townsend 1994; Udry 1994). In support of this theory, Munshi and Rosenzweig (2016) show that in a set of Indian villages, those that have riskier incomes are less likely to have migrant members. Counterfactual simulations from their model predict that improving insurance markets would lead to substantial reductions in the misallocation of rural workers, many of whom are currently stuck inefficiently in rural areas.

Frictions in land markets in poor countries also may be an important factor holding back rural-urban migration and income levels for many rural households. Unlike in the United States, where real estate is bought and sold readily just about anywhere, few Africans or South Asians hold a title to their land that could allow them to sell it, even in principle. To be sure, land-titling programs are growing, but traditional systems of land rights are still the norm, and markets for land are still nonexistent in most rural areas. In the context of internal migration, the lack of land markets, along with financial frictions discussed above, suggest that it will

⁵Though financial constraints certainly may play an important role in holding back internal migration in other settings, as Cai (2020) shows for China. He randomizes the rollout of a microfinance program in rural Chinese villages and shows that those getting micro-finance loans are much more likely to send migrants seasonally to nearby cities. Migrants experienced increases in earnings of around 36 percent relative to the mean of the control group. Angelucci (2015) and Bazzi (2017) provide evidence that financial constraints hold back *international* migration among very poor households in Mexico and Indonesia, respectively.

be hard for villagers to save up to fund a migration episode, which requires liquid wealth. In addition, those without formal title may be unlikely to migrate for fear of losing their land.

Land market frictions can be potent in holding back migration. A study by de Janvry, Emerick, Gonzalez-Navarro, and Sadoulet (2015) analyzes the rollout of a new land-titling system in Mexico that began in the 1990s, which gave official ownership of land to rural households that had previously farmed the land informally. The authors find that those receiving a title over their land were 28 percent more likely to send migrants. Their interpretation is that the insecurity over land ownership, rather than the lack of liquidity, was the main reason migration was not higher before the titling program. In related studies, Chen (2017) and Gottlieb and Grobovsek (2019) argue that land frictions like these, which keep too many workers in agriculture, lead to substantial misallocation of talent. Using a model of structural change calibrated to Ethiopia, Gottlieb and Grobovsek (2018) simulate the effects of improving markets for land and find that it would raise aggregate productivity by 9 percent, as previously misallocated agriculture workers move into more productive non-agricultural activities. Chen (2017) finds an even larger effect in Malawi, which had virtually none of its land titled as recently as 2007.

There are certainly other frictions that hold back migration within developing countries, and research on this important topic is still in its infancy in many ways. Future work should thus continue to help identify and measure the frictions holding back migration of rural workers into more productive opportunities in cities. More analyses that help quantify their importance of specific barriers to migration for development outcomes, like aggregate productivity, would also be most valuable.

■ *For helpful comments on this essay, I thank the JEP editors: Gordon Hanson, Enrico Moretti, Heidi Williams, and especially Timothy Taylor, plus Doug Gollin, Jim Rauch, and Todd Schoellman.*

References

- Akram, Agha Ali, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak.** 2018. "Effects of Emigration on Rural Labor Markets." Unpublished.
- Alvarez, Jorge A.** 2020. "The Agricultural Wage Gap: Evidence from Brazilian Micro-data." *American Economic Journal: Macroeconomics* 12 (1): 153–73.
- Angelucci, Manuela.** 2015. "Migration and Financial Constraints: Evidence from Mexico." *Review of Economics and Statistics* 97 (1): 224–28.
- Asher, Sam, and Paul Novosad.** 2020. "Rural Roads and Local Economic Development." *American Economic Review* 110 (3): 797–823.

- Banerjee, Abhijit, Esther Duflo, and Nancy Qian.** 2020. "On the Road: Access to Transportation Infrastructure and Economic Growth in China." *Journal of Development Economics* 145: Article 102442.
- Baseler, T.** 2019. "Hidden Income and the Perceived Returns to Migration: Experimental Evidence from Kenya." Working Paper, University of Rochester.
- Bazzi, Samuel.** 2017. "Wealth Heterogeneity and the Income Elasticity of Migration." *American Economic Journal: Applied Economics* 9 (2): 219–55.
- Bazzi, Samuel, Arya Gaduh, Alexander D. Rothenberg, and Maisy Wong.** 2016. "Skill Transferability, Migration, and Development: Evidence from Population Resettlement in Indonesia." *American Economic Review* 106 (9): 2658–98.
- Beegle, Kathleen, Joachim De Weerd, and Stefan Dercon.** 2011. "Migration and Economic Mobility in Tanzania: Evidence from a Tracking Survey." *Review of Economics and Statistics* 93 (3): 1010–33.
- Brueckner, Jan K., and Somik V. Lall.** 2015. "Cities in Developing Countries: Fueled by Rural-Urban Migration, Lacking in Tenure Security, and Short of Affordable Housing." In *Handbook of Regional and Urban Economics*, edited by Gilles Duranton, J. Vernon Henderson, and William Strange, vol. 5A, 1399–1456. Amsterdam: Elsevier.
- Bryan, Gharad, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak.** 2014. "Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh." *Econometrica* 82 (5): 1671–1748.
- Bryan, Gharad, Edward Glaeser, and Nick Tsivanidis.** 2019. "Cities in the Developing World." NBER Working Paper 26390.
- Bryan, G. and M. Morten.** 2014. "The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia." *Journal of Political Economy* 127 (5): 2229–68.
- Cai, Shu.** 2020. "Migration under Liquidity Constraints: Evidence from Randomized Credit Access in China." *Journal of Development Economics* 142: 1–17.
- Caselli, Francesco.** 2005. "Accounting for Cross-Country Income Differences." In *Handbook of Economic Growth*, edited by Philippe Aghion and Steven N. Durlauf, 679–741. Amsterdam: Elsevier.
- Chauvin, Juan Pablo, Edward Glaeser, Yueran Ma, and Kristina Tobio.** 2017. "What is Different about Urbanization in Rich and Poor Countries? Cities in Brazil, China, India and the United States." *Journal of Urban Economics* 98: 17–49.
- Chen, Chaoran.** 2017. "Untitled Land, Occupational Choice, and Agricultural Productivity." *American Economic Journal: Macroeconomics* 9 (4): 91–121.
- Chiquiar, Daniel, and Gordon H. Hanson.** 2005. "International Migration, Self-Selection, and the Distribution of Wages: Evidence from Mexico and the United States." *Journal of Political Economy* 113 (2): 239–81.
- Costa, Dora L. and Matthew E. Kahn.** 2006. "Public Health and Mortality: What Can We Learn from the Past?" In *Public Policy and the Income Distribution*, edited by Alan J. Auerbach, David Card, and John M. Quigley, 359–98. New York: Russell Sage.
- De Janvry, Alain, Kyle Emerick, Marco Gonzalez-Nava Rro, and Elisabeth Sadoulet.** 2015. "Delinking Land Rights from Land Use: Certification and Migration in Mexico." *American Economic Review* 105 (10): 3125–49.
- Deaton, Angus, and Olivier Dupriez.** 2011. "Spatial Price Differences within Large Countries." Unpublished.
- Dercon, Stefan, Pramila Krishnan, and Sofya Krutikova.** 2013. "Changing Living Standard in Southern Indian Villages 1975–2006: Revisiting the ICRISAT Village Level Studies." *Journal of Development Studies* 49 (12): 1676–93.
- Fafchamps, Marcel, and Christine Moser.** 2003. "Crime, Isolation, and Law Enforcement." *Journal of African Economies* 12 (4): 625–71.
- Ferré, Céline, Francisco H. G. Ferreira, and Peter Lanjouw.** 2012. "Is There a Metropolitan Bias? The Relationship between Poverty and City Size in a Selection of Developing Countries." *The World Bank Economic Review* 26 (3): 351–82.
- Glaeser, Edward.** 2011. *Triumph of the City*. New York, New York: Penguin Group.
- Glaeser, Edward L., and Joshua D. Gottlieb.** 2009. "The Wealth of Cities: Agglomeration Economies and Spatial Equilibrium in the United States." *Journal of Economic Literature* 47 (4): 983–1028.
- Gollin, Douglas.** 2014. "The Lewis Model: A 60-Year Retrospective." *Journal of Economic Perspectives* 28 (3): 71–88.
- Gollin, Douglas, Remi Jedwab, and Dietrich Vollrath.** 2016. "Urbanization with and without Industrialization." *Journal of Economic Growth* 21: 35–70.
- Gollin, Douglas, Martina Kirchberger, and David Lagakos.** 2019. "Do Urban Wage Premia Reflect Lower

- Amenities? Evidence from Africa." Unpublished.
- Gollin, Douglas, David Lagakos, and Michael E. Waugh.** 2014. "The Agricultural Productivity Gap." *Quarterly Journal of Economics* 129 (2): 939–93.
- Gollin, Douglas, Stephen L. Parente, and Richard Rogerson.** 2004. "Farm Work, Home Work and International Productivity Differences." *Review of Economic Dynamics* 7 (4): 827–50.
- Gottlieb, Charles, and Jan Grobovšek.** 2019. "Communal Land and Agricultural Productivity." *Journal of Development Economics* 138: 135–52.
- Hanlon, W. Walker, and Yuan Tian.** 2015. "Killer Cities: Past and Present." *American Economic Review: Papers & Proceedings* 105 (5): 570–75.
- Harris, John R., and Michael P. Todaro.** 1970. "Migration, Unemployment and Development: A Two-Sector Analysis." *American Economic Review* 60 (1): 126–42.
- Hatton, Timothy J., and Jeffrey G. Williamson.** 1992. "What Explains Wage Gaps between Farm and City? Exploring the Todaro Model with American Evidence, 1980–1941." *Economic Development and Cultural Change* 40 (2): 267–94.
- Herrendorf, Berthold, and Todd Schoellman.** 2015. "Why Is Measured Productivity So Low in Agriculture?" *Review of Economic Dynamics* 18 (4): 1003–22.
- Herrendorf, Berthold, and Todd Schoellman.** 2018. "Wages, Human Capital, and Barriers to Structural Transformation." *American Economic Journal: Macroeconomics* 10 (2): 1–23.
- Hicks, Joan Hamory, Marieke Kleemans, Nicholas Y. Li, and Edward Miguel.** 2017. "Reevaluating Agricultural Productivity Gaps with Longitudinal Microdata." NBER Working Paper 23253.
- Hnatkovska, Viktoria, and Amartya Lahiri.** 2018. "Urbanization, Structural Transformation and Rural-Urban Disparities in China and India." Unpublished.
- Hsieh, Chang-Tai, and Peter J. Klenow.** 2010. "Development Accounting." *American Economic Journal: Macroeconomics* 2 (1): 207–23.
- Jedwab, Remi, and Vollrath, Dietrich.** 2016. "The Urban Mortality Transition and the Rise of Poor Mega-Cities." Unpublished.
- Kleemans, Marieke.** 2015. "Migration Choice under Risk and Liquidity Constraints." Unpublished.
- Kline, Patrick, and Enrico Moretti.** 2014. "People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs." *Annual Review of Economics* 6: 629–62.
- Lagakos, David, Samuel Marshall, Ahmed Mushfiq Mobarak, Michael E. Waugh, and Corey Vernot.** Forthcoming. "Migration Costs and Observational Returns to Rural-Urban Migration." *Journal of Monetary Economics*.
- Lagakos, David, Ahmed Mushfiq Mobarak, and Michael E. Waugh.** 2019. "Welfare Effects of Encouraging Rural-Urban Migration." Unpublished.
- Lagakos, David, and Michael E. Waugh.** 2013. "Selection, Agriculture, and Cross-Country Productivity Differences." *American Economic Review* 103 (2): 948–80.
- Lucas, Robert E. B.** 2015. "Internal Migration in Developing Countries: An Overview." KNOWMAD Working Paper 6.
- Lucas, Robert E., Jr.** 2004. "Life Earnings and Rural-Urban Migration." *Journal of Political Economy* 112 (S1): 29–59.
- Lundh, Christer, and Svante Prado.** 2015. "Markets and Politics: the Swedish Urban-Rural Wage Gap, 1865–1985." *European Review of Economic History* 19 (1): 67–87.
- McKenzie, David, Steven Stillman, and John Gibson.** 2010. "How Important Is Selection? Experimental vs. Non-Experimental Measures of the Income Gains from Migration." *Journal of the European Economic Association* 8 (4): 913–45.
- McMillan, Margaret, Dani Rodrik, and Íñigo Verdugo-Gallo.** 2014. "Globalization, Structural Change, and Productivity Growth, with an Update on Africa." *World Development* 63: 11–32.
- Morten, Melanie.** 2019. "Temporary Migration and Endogenous Risk Sharing in Village India." *Journal of Political Economy* 127 (1): 1–46.
- Morten, Melanie, and Jaqueline Oliveira.** 2018. "The Effects of Roads on Trade and Migration: Evidence from a Planned Capital City." NBER Working Paper 22158.
- Munshi, Kaivan, and Mark Rosenzweig.** 2016. "Networks and Misallocation: Insurance, Migration, and the Rural-Urban Wage Gap." *American Economic Review* 106 (1): 46–98.
- Nakamura, Emi, Jósef Sigurdsson, and Jón Steinsson.** 2019. "The Gift of Moving: Intergenerational Consequences of a Mobility Shock." Unpublished.
- Porzio, Tommaso, and Gabriella Santangelo.** 2019. "Does Schooling Cause Structural Transformation?" Unpublished.

- Pulido, José, and Tomasz Świącki.** 2018. "Barriers to Mobility or Sorting? Sources and Aggregate Implications of Income Gaps across Sectors and Locations in Indonesia." Unpublished.
- Ravallion, Martin, Shaohua Chen, and Prem Sangraula.** 2007. "New Evidence on the Urbanization of Global Poverty." *Population and Development Review* 33(4): 667–701.
- Restuccia, Diego, Dennis Tao Yang, and Xiaodong Zhu.** 2001. "Agriculture and Aggregate Productivity: A Quantitative Cross-Country Analysis." *Journal of Monetary Economics* 55 (2): 234–50.
- Restuccia, Diego, Dennis Tao Yang, and Xiaodong Zhu.** 2008. "Agriculture and Aggregate Productivity: A Quantitative Cross-Country Analysis." *Journal of Monetary Economics* 55 (2): 234–50.
- Roy, A. D.** 1951. "Some Thoughts on the Distribution of Earnings." *Oxford Economic Papers* 3 (2): 135–46.
- Sarvimäki, Matti, Roope Uusitalo, and Markus Jäntti.** 2019. "Habit Formation and the Misallocation of Labor: Evidence from Forced Migrations." Unpublished.
- Tombe, Trevor, and Xiaodong Zhu.** 2009. "Trade, Migration, and Productivity: A Quantitative Analysis of China." *American Economic Review* 109(5): 1843–72.
- Tombe, Trevor, and Xiaodong Zhu.** 2019. "Trade, Migration and Productivity: A Quantitative Analysis of China." *American Economic Review* 109 (5): 1843–72.
- Townsend, Robert M.** 1994. "Risk and Insurance in Village India." *Econometrica* 62(3): 539–91.
- Udry, Christopher.** 1994. "Risk and Insurance in a Rural Credit Market: An Empirical Investigation in Northern Nigeria." *The Review of Economic Studies* 61 (3): 495–526.
- Vollrath, Dietrich.** 2009. "How Important are Dual Economy Effects for Aggregate Productivity?" *Journal of Development Economics* 88 (2): 325–34.
- Vollrath, Dietrich.** 2014. "The Efficiency of Human Capital Allocations in Developing Countries." *Journal of Development Economics* 108: 106–18.
- Young, Alwyn.** 2012. "The African Growth Miracle." *Journal of Political Economy* 120 (4): 696–739.
- Young, Alwyn.** 2013. "Inequality, the Urban-Rural Gap, and Migration." *The Quarterly Journal of Economics* 128 (4): 1727–85.