Migration Policy and Observational Returns to Rural-Urban Migration in the Developing World

David Lagakos (UCSD and NBER)

Samuel Marshall (University of Warwick)

Ahmed Mushfiq Mobarak (Yale University)

Corey Vernot (Yale University)

Michael E. Waugh (New York University and NBER)

Nov 11, 2019

ABSTRACT

Recent studies find that observational returns to rural-urban migration are near zero in several developing countries. Do this imply that the gains from policies that reduce migration costs will also be low? To answer this question we build a simple Roy model in which workers are heterogeneous in both their return to urban-rural migration and their cost of migrating. In the model, observational returns to migration are driven by individuals that have high returns to migration net of costs, and many rural individuals with high potential returns to migration do not actually migrate in equilibrium due to even higher costs of migrating. The model predicts that rural regions with higher migration costs send fewer migrants and have higher observational returns to migration. We provide new evidence supporting this predictions from panel tracking surveys from six developing countries. In the case of Bangladesh, we show further that observational returns to migration are substantially lower than the returns for those induced to migrate through reduced migration costs. We conclude that migration policy should not look to observational returns to migration for guidance about the gains from increasing internal migration by reducing migration costs.

Email: lagakos@ucsd.edu, ahmed.mobarak@yale.edu, mwaugh@stern.nyu.edu. For outstanding research assistance we thank Tomer Mangoubi, Hideto John Mori, Sebastian Quaade, Pranav Bhandarkar, Michael Mao, Lixing Liang, Min Byung Chae, Liana Wang, Hannah Moreno, Jorge Colmenares, Xiaouya Gao, Lukas Fesser, Huang Songjie, and Jaclyn Schess. All potential errors are our own.

1. Introduction

A growing body of evidence suggests that encouraging rural-urban migration may be of the most promising policy options available to developing countries in order to increase aggregate productivity and facilitate economic growth. Average levels of income and consumption tend to be much higher in urban areas than in rural areas in most developing countries, even for observationally similar workers (Young, 2013; Gollin, Lagakos, and Waugh, 2014; Herrendorf and Schoellman, 2018). Direct experimental evidence points to large returns to encouraging rural workers in Bangladesh to migrate seasonally to urban areas (Bryan, Chowdhury, and Mobarak, 2014). At least partly in light of evidence of this type, many developing countries have begun to pay more attention to the implicit borders that divide rural and urban areas within their countries, and to strategies for facilitating movements of rural workers into more productive urban centers.

Yet recent findings cast doubt on rural-urban migration as a viable growth strategy for developing nations. Using panel tracking studies from Kenya and Indonesia, Hicks, Kleemans, Li, and Miguel (2017) find that observational returns to migration – meaning the returns for those observed to migrate in the data – are near zero on average. Alvarez (Forthcoming) finds a similar result using Brazilian panel tracking data. Relative to the cross-sectional gaps, the consumption gains from rural-urban migration are an order of magnitude smaller. The authors of these studies join a chorus of other recent papers positing that rural-urban gaps in average living standards are in fact due to efficient sorting of workers across space within countries (Lagakos and Waugh, 2013; Young, 2013; Herrendorf and Schoellman, 2018). In short, these studies posit that those living in urban areas are simply those with better skills for urban jobs, and that rural workers, who lack these same skills, will not be benefit much from migrating to cities. As such, policy makers should not spend time trying to reduce migration costs for rural households.

In this paper we propose an alternative interpretation for the low observational returns to migration relative to the large cross-sectional wage gaps. Our theory builds upon the premise that migration costs – broadly defined – are highly heterogeneous across individuals. Individuals who are geographically or linguistically isolated from urban areas face would higher costs of migration (see e.g. Morten and Oliveira, 2019). Conversely, rural individuals with strong networks in urban destinations, for example, would face a lower cost to finding work and housing in cities after migrating there. The same is true for those whose family responsibilities include caregiving for other older or younger family members. As such, those who are observed to migrate from rural to urban areas in the data are likely to be those whose benefits of migration are higher than the costs, not just those who have a high benefit to migration overall. Left behind may be those with high potential benefits from migration, but higher potential cost. The basic implication of our theory is that lowering migration costs could in fact have large aggregate productivity effects, by helping those with high benefits – but also high costs – of moving to urban areas.

We formalize this theory in a simple model of rural-urban migration in which workers are heterogeneous in comparative advantage as well as in their cost of migration. The model economy consists of three regions: an urban region and two rural regions. One of the rural regions has a high cost of migration to the urban area and the other has a low cost. At birth, individuals observe a vector of labor productivities in each region as in Roy (1951) and decide whether to stay where they are or migrate. Sorting otherwise works as in other recent Roy models in this literature, such as the ones studied by Lagakos and Waugh (2013); Young (2013), Hicks, Kleemans, Li, and Miguel (2017) and Nakamura, Sugrudsson, and Steinsson (2019). In our model, as in these, sorting is efficient conditional on migration costs, and a benevolent social planner would not want to re-allocate workers across space.

In the model, the workers that migrate in equilibrium are those for whom the return to migration exceeds the costs. These include workers with high overall comparative advantage in urban areas, as well as those with modest benefits but even lower costs. The non-migrants may have strong comparative advantage in urban areas, but sufficiently high moving costs that they choose not to migrate. Lowering migration costs in the model can lead to substantial increases in income for those in the high-cost region, and the returns to migration for those induced to move by the lower cost exceed the average observational return before the migration. More generally, the model highlights how the observational returns to migration are not directly informative about the aggregate gains from lowering migration costs.

The model makes a clear testable implication about how migration costs, out-migration rates and average returns to migration covary across rural regions. Specifically, it predicts that rural regions with higher migration costs should be those with lower migration rates and higher average returns to migration. The intuition is that the only individuals to migrate from regions with higher migration costs in equilibrium are those with the highest potential returns to migration. This prediction highlights the importance of cost in determining both which individuals decide to migrate in equilibrium and in what their observed return to migration will be.

To test the model's assumptions and predictions, we turn to recent panel tracking surveys from six developing countries: China, Indonesia, Malawi, South Africa, Tanzania, and Ghana). Together, our data cover nearly 200,000 individuals across 1,600 communities, and successfully track over 10,000 rural-urban migrants. The surveys have detailed measures of consumption at the household level that is measured in the same way for wage earners and the self employed, unlike measures of income, which have been the focus of previous studies. The surveys differ in important ways, but allow us to measure how the observations returns to migration in a consistent way using panel regressions with individual fixed effects as in the work of Hicks, Kleemans, Li, and Miguel (2017). The data also allow to measure how rural-urban migration rates covary with observational returns to migration across a large set of different rural regions.

Our data largely confirm the finding of Hicks, Kleemans, Li, and Miguel (2017) and Alvarez (Forthcoming) that observational returns to rural-urban migration are substantially smaller than cross-sectional gaps. In Malawi and Tanzania, for example, we estimate that consumption per adult equivalent is 52 and 60 log points higher in urban areas than in rural areas in the cross-section. Though once individual fixed effects are included, the gaps fall to 4 and 5 log points, with neither estimate being statistically significantly different from zero.

On the other hand, our data are not consistent with the findings of Hicks, Kleemans, Li, and Miguel (2017) that individual fixed effects render the estimated urban-rural gaps near zero on average. Even with individual fixed effects, the estimated urban dummies in the regression are 19 log points in China, 18 log points in Ghana, 44 log points in South Africa and 12 log points in Indonesia. The simple average across our six countries is 17 log points. This average increases to 26 log points when we include only rural to urban migrants in the sample, excluding urban to rural migrants, which we show theoretically is preferable. The observational returns we estimate compare consumption per adult, rather than consumption per hour, which Pulido and Swiecki (2018) argue more accurately reflects the margin relevant for a potential migrant. Our focus on consumption, rather than income, also allows us to sidestep serious measurement challenges that are present for measures of income taken from household surveys in developing parts of the world.

Across five of our six countries, we find find support for our model's prediction that rural regions with higher migration rates tend to have the highest observational returns to migration. As one way to show this, we run two sets of panel fixed-effect regressions like those described above, one with rural-urban migrants from low-migration regions and the other with migrants from high-migration regions. In China, for example, the observational return to migration is 32 log points from the low-migration rural regions and 7 log points from the high-migration regions, and the difference is highly significant. The other countries, other than Malawi, with an insignificant difference, exhibit similar patterns. These empirical findings largely corroborate the model's predictions and its emphasis on migration costs as an important barrier to migration for an important portion of rural households.

We lastly turn to experimental evidence on migration costs in Bangladesh, re-visiting the experimental evidence of Bryan, Chowdhury, and Mobarak (2014). These data (but few other data sets) allow us to compare experimental estimates of the return to migration to observational estimates using panel regressions with individual fixed effects as in the recent literature. We find that the observational returns to migration in those panel data (using only from the control group) are 9 log points and statistically significant at the 10-percent level. The experimental returns to migration are 36 log points and significant at the 5-percent level. The difference is large and statistically significant at the 1-percent level. The sharply larger experimental returns to migration than the observational returns supports our paper's conclusion that observational return may not paint an accurate picture of the gains from lowering migration costs.

Our paper's conclusions about the positive effects of lowering migration barriers echoes that of Bryan and Morten (2019). Their paper also documents a pattern of higher wages for migrants in Indonesia that come from further distance, which is an important type of migration cost. Their paper does not confront how low average observational returns to migration relate or not to gains from reducing migration barriers, as we do. Empirically, our observed returns to migration are also closer to the higher numbers of Beegle, De Weerdt, and Dercon (2011) than the lower ones of Hicks, Kleemans, Li, and Miguel (2017) and Alvarez (Forthcoming), though our work builds on these important studies and is motivated by their findings.

2. Model of Migration with Heterogeneous Migration Costs

In this section, we present a simple model to help relate returns to migration: specifically, the cross-sectional gap, the observed return to migration among migrants, and the gains from encouraging migration more broadly.

2.1. Model Environment

The economy is composed of a unit measure of individuals who live for two periods. There are three closed competitive labor markets—an urban market indexed by u and two rural markets (low and high) indexed by r_L and r_H. The two rural labor markets differ in how costly it is to migrate out of that market. Individuals who choose to migrate out of a rural labor market lose a portion of their income (which we will explain below), m, in the next period, where $m \in \{m_H, m_L\}$ and $0 < m_H < m_L < 1$.

To motivate the two types of rural areas, we can think of an economy that consists of a city and two villages. One village (the low-type) is close to the city. Individuals from that village can easily take a bus to the city. The other village is geographically isolated from the city by a mountain range which makes it more difficult to migrate to the city.

We restrict our model so that migration is only possible at the end of the first period. This allows us to observe the income of migrants both before and after migration.

A share π_u of individuals is born in the urban market and the share λ of those are born in the rural area are born into high-type rural markets. Thus, $(1 - \pi_u)$ start out in the high-type rural market and $(1 - \lambda)(1 - \pi_u)$ start out in the low-type rural market.

Endowments and Preferences. Individual are endowed with one unit of time which they supply inelastically to labor in each period. Individuals are also endowed with a bivariate skill vector $\{z_r,\, z_u\}$ that determines how many efficiency units of labor she supplies in rural and urban labor. Skills are distributed Frechet. For the rural area, rural skills are where $Z_r \sim F(1, \theta)$ with the mean value normalized to one. Urban skills are likewise distributed $Z_u \sim F(T_u, \theta)$ where T_u is mean urban ability and θ indexes the dispersion parameter. Skills across space are independent of each other. Let $G(Z_r, Z_u)$ be the joint distribution of skills and there is no inherent differences in the distribution of skills across space (i.e. the rural and urban talent distribution are the same).

Individuals are endowed with a factor $\psi > 1$ which makes them more productive in their home labor market. One can think of ψ as location-specific knowledge that makes native workers more productive than non-native workers of the same ability. Bazzi, Gaduh, Rothenberg, and Wong (2016) argue that exactly this kind of specific knowledge prevents skills from being transferable across regions.

Migration Choices and Consumption Taking their skill vector, migration cost, and wages as given, individuals make a migration decision at time one to maximize their income in the second period. For now, we will assume that those in the urban location stay in the urban location. And then that rural migrants only flow towards the urban area. The assumption of home-specific knowledge supports this pattern of flows. The consumption of a work in period two is then

$$w = \begin{cases} \max\{\psi\, z_r, \quad m_L\, A\, z_u\} & \text{if } l = r_L \\ \max\{\psi\, z_r, \quad m_H\, A\, z_u\} & \text{if } l = r_H \\ \psi\, A\, z_u & \text{if } l = u. \end{cases} \tag{1}$$

Technology. A representative firm in each market takes labor as its sole input to produce a single consumption good, Y. Production in both rural labor markets is the same and production in each market is linear in labor:

$$Y_u = A\,L_u \qquad \text{and} \qquad Y_r = L_r, \tag{2}$$

where

$$L_l = \int_{\Omega^l} z_l\, dG \quad \text{and} \quad N_l = \int_{\Omega^l} dG$$

and Ω^l denotes the set of workers employed in labor market l and N_l is the total number of workers employed in that labor market.

Market Clearing. In each period, the labor and consumption goods markets clear:

$$N_{r_L t} + N_{r_H t} + N_{ut} = 1 \tag{3}$$

$$Y_{r_L t} + Y_{r_H t} + Y_{ut} = C_t, \tag{4}$$

where $t \in \{1, 2\}$ and C_t is aggregate consumption in the economy.

Equilibrium. We normalize the price of the consumption good to one so that the competitive equilibrium price of one efficiency unit of labor in both rural labor markets is one. Specifically, the urban and rural wages per efficiency unit are

$$w_u = A_u \qquad \text{and} \qquad w_{r_L} = 1 \qquad \text{and} \qquad w_{r_H} = 1. \tag{5}$$

2.2. Equilibrium Results

The model above is deliberately simple to allow us to think through the how the observational returns to migration relate to the cross-sectional gaps and returns for those induced to migrate by lowering migration costs.

Cross-Sectional Gap. The cross-sectional gap is simply average consumption in the urban area compared to the average consumption in rural area. It is analogous to the cross-sectional gaps studied in Gollin, Lagakos, and Waugh (2014), Young (2013) and Herrendorf and Schoellman (2018). The cross-sectional gap takes a simple representation in our model in period one, when no one has moved:[1]

$$\text{CSG} := \frac{\psi A_u E z_u}{\psi E z_r} = T_u A_u. \tag{6}$$

[1]While the higher average urban productivity in period one is assumed exogenously here, it is meant to capture the steady state of a more dynamic environment. For example, Lagakos, Mobarak, and Waugh (2019) endogenously generate higher average worker productivity in urban areas in the stationary distribution of their dynamic two-sector Roy model.

Equation (6) tells us that the cross-sectional gap reflects two forces: (i) differences in ability between urban and rural areas – governed by T_u – and (ii) the effect of being in the urban area – captured by A_u. This formulation exactly picks up the idea that the cross-sectional gap is a mixture of both inherent talent differences and the treatment effect of the urban area. The latter is causal and could be exploited by policy. The former is not.

Observational Returns to Migration. We will define the observational returns to migration as the average return to migration among those that choose to migrate in equilibrium. This is the return measured by Hicks, Kleemans, Li, and Miguel (2017). Note that in their data, and in our model, we measure the gains inclusive of migration costs. That is, what we see $A_u z_u$ in the urban area, not consumption net of the moving cost, $mA_u z_u$. For a particular individual that migrates, her observed return is given by:

$$R \equiv \log A_u z_u - \log \psi z_r, \tag{7}$$

which has to be positive if she decides to move in equilibrium. It is useful to rearrange this return and put some labels on each term:

$$\underbrace{\log A_u}_{\text{urban effect}} \quad + \quad \underbrace{\log z_u - \log z_r}_{\text{selection on c.a.}} \quad - \quad \underbrace{\log \psi}_{\text{local knowledge loss}} \quad > 0 \tag{8}$$

so the observed return to migration depends on several things. The first is the benefit of being in the urban area, A_u; the second term is about individual talents, or what we will call selection on comparative advantage; the third term reflects losses in the "local knowledge" that rural-urban migrants give up. When we integrate over workers from across rural locations, we get the observational return to migration:

$$\text{ORM} = \log A_u - \log \psi + \sum_{l \in H,L} \pi_{u,r_l} \mathrm{E}\left\{\log z_u - \log z_r | \text{migrate}_l\right\}, \tag{9}$$

$$\text{migrate}_l \text{ if } \quad \log z_u - \log z_r > \log m_l - \log A_u + \log \psi,$$

where the individual selection term is replaced with a sum over conditional expectations and π_{u,r_l} is the fraction of urban migrants from each locations. The location-specific conditional expectation term reflects the individual-specific gains from migration. It is conditional because these gains are only observed if it is beneficial to migrate. In the "standard case" of Roy models where comparative advantage aligns with absolute advantage (which is true with the Frechet distribution and many others), we should expect this term to be positive, i.e. on average the individual-specific gains to migration should be positive.

Can A_u be large while the observed returns to migration are low? Yes. One could observe low returns to migration for several reasons. First, all migrants give up benefits from being in the rural area which are represented by the local knowledge loss and are not recouped in the urban area. The ψ term is that it allows our model to generate low observed returns—even though selection may be at work and even though A_u, the causal effect of going to the urban area may be positive.

To see this take the difference between the cross-sectional gap and the observational returns in (12):

$$\text{CSG} - \text{ORM} = \log T_u + \log \psi - \sum_{l \in H,L} \pi_{u,r_l} \mathrm{E}\left\{\log z_u - \log z_r | \text{migrate}_l\right\} \tag{10}$$

which embeds three terms: intrinsic talent differences, knowledge loss, and migrant selection. The cross-sectional gap could be larger than observed returns because talent is intrinsically different in the urban area, this is the T_u term. This is the Hicks et al interpretation of the data: the cross-sectional gap is much larger than the observational return, and so the difference must be all due to efficient sorting of heterogeneous individuals. These need not be the case, however, it could also be because of the ψ term. That is, the opportunity cost of rural migration is high which lowers the returns to migration.

There is a second reason why observed returns to migration may seem low and that is differential selection upon moving costs. The selection term in the observed return to migration is weighted towards low-migration cost regions. In equilibrium, most rural migrants come from the low-migration region, and this skews the migration returns toward the average return among migrants from there. In particular, the mass of individuals migrating to the urban area from each rural location l is:

$$\pi_{u,r_l} = \frac{A_u^\theta}{A_u^\theta + (\psi m_l)^\theta}. \tag{11}$$

which shows that locations with larger migration costs will have lower migration probabilities. In other words, the individual specific gains to migration are larger for high-moving-cost locations than low-moving-cost locations. But because (12) puts lower weight on the high-return locations, these high-returns may be masked.

Gains from Lowering Migration Costs. Suppose that a policymaker were interested in the gains from lowering internal migration costs. Could she turn to the observational returns to migration for guidance? The answer in this model is no. Equation (12) is informative about all those migrating in equilibrium already before any reductions in cost. To compute the gains from lowering migration costs by some amount Δm, one would have to integrate all returns R, defined as in (7), for all individuals that do *not* migrate in equilibrium, but do migrate when costs are lowered by Δm:

$$\text{GLC} = \log A_u - \log \psi + \sum_{l \in H,L} \pi^*_{u,r_l} \mathrm{E}\left\{\log z_u - \log z_r | \text{not migrate}_l \ \& \ \text{migrate}_{\Delta m}\right\} > 0, \quad (12)$$

$$\text{not migrate}_l \text{ if } \log z_u - \log z_r \leq \log m_l - \log A_u + \log \psi,$$
$$\text{migrate}_{\Delta m} \text{ if } \log z_u - \log z_r > \log(m_l - \Delta m) - \log A_u + \log \psi,$$

and π^*_{u,r_l} represents the probability of moving from location l after costs are lowered.

To illustrate, Figure 1 illustrates how the model works for some arbitrary parameter values ($A = 1.5$, $\psi = 1.5$, $\theta = 5$, $T_u = 1$.) We illustrate the observational returns and gains from lowering migration costs over the range $m_h \in (0.2)$ and under 4 values for m_l: (.01, .05, .1, .25). The policy experiment is lowering the costs by 75 percent in each region. The blue circles lower the cost only in the high cost region, while the red circles lower the cost in the low and high cost regions. The x-axis is m_h, and the y-axis is the average return, defined as in (7), for all migrants. Each of the four graphs reflects one particular value for m_l, as specified in the title.

As Figure 1 makes clear, the returns to migrants induced to move by lowering migration costs can be higher or lower than the observational returns. The difference depends on amount of heterogeneity in cost between the low and high cost regions. It also depends on whether the cost reductions are targeted toward the high-cost region, with larger gains from targeting. Finally, one can see that the gains from lowering migration costs are higher than the observational returns when the costs are higher, and lower for low costs in the high region.

Figure 1: Observational Returns to Migration and Gains From Lowering Migration Costs

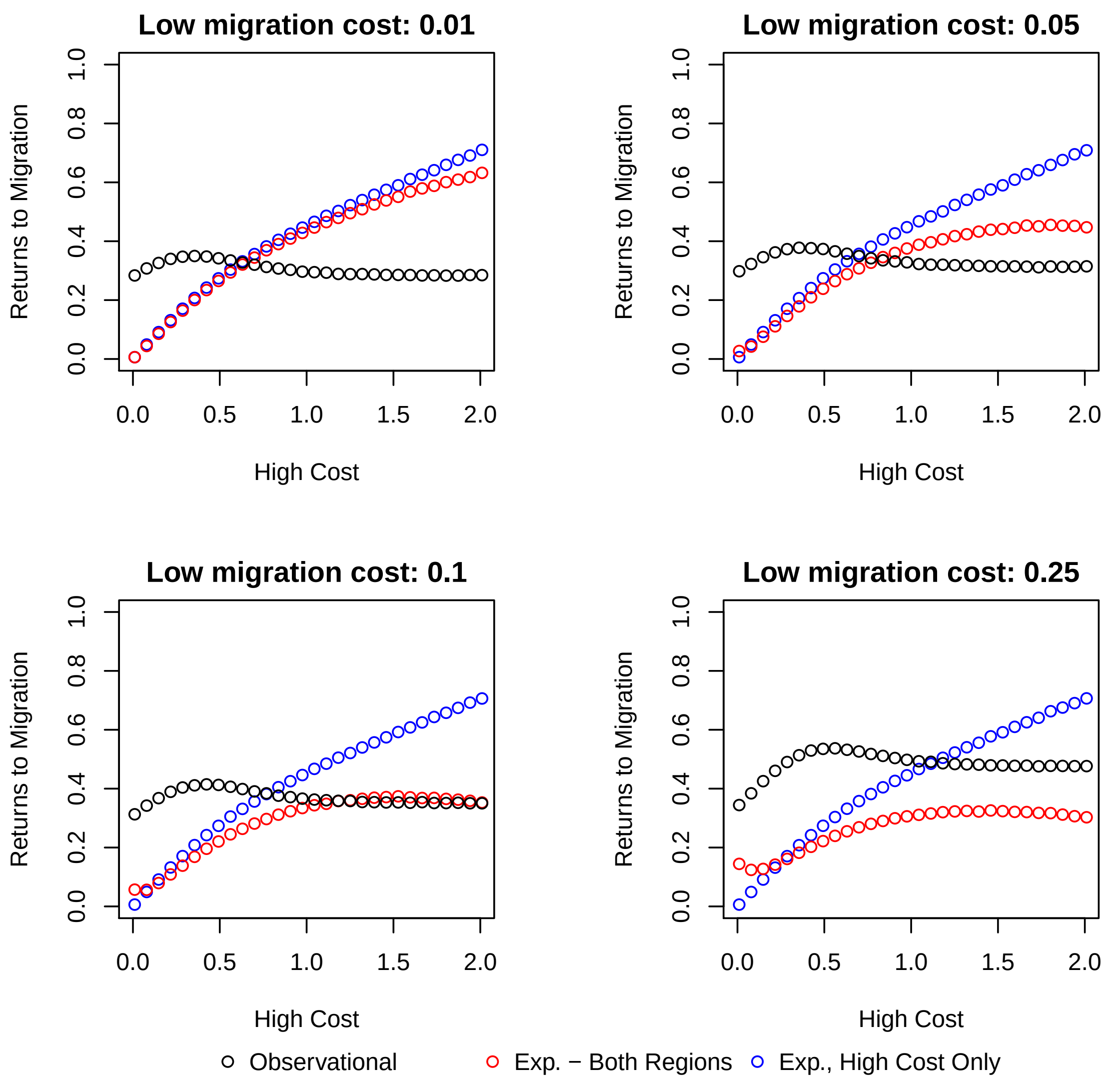

3. Observational Returns to Migration: New Evidence

In this section we describe the panel tracking surveys that covering six developing countries. We use them to estimate the observational returns to migration overall and by region of origin. We confirm the finding of Hicks, Kleemans, Li, and Miguel (2017) that cross-sectional wage gaps are larger than observational returns to migration, which we estimate using panel regressions with individual fixed effects. However, we find larger returns in the five countries not studied by Hicks, Kleemans, Li, and Miguel (2017) than the one they studied that is in our sample (Indonesia). When looking by region of origin, we find that rural regions with high migration rates have significantly lower observational returns to migration than do regions with lower migration rates. This is consistent with what our model predicts.

3.1. Data Overview

Our primary data comes from household panel surveys conducted in six developing countries: China, Indonesia, Ghana, Malawi, South Africa, and Tanzania. We chose panel surveys that shared three characteristics necessary for our analysis. First, the survey attempts to track and re-survey individuals that migrate between survey waves waves. Tracking survey respondents across space is very costly, but it allows us to estimate the consumption gains from changing locations while controlling for unobserved time-invariant characteristics as in Beegle, De Weerdt, and Dercon (2011), Hicks, Kleemans, Li, and Miguel (2017) and Alvarez (Forthcoming). Second, the survey collects detailed geographic information on each individual in each wave, allowing us to classify all individuals into urban-rural status and then further by more disaggregated areas. This allows us to estimate gains from rural-urban migration overall and by source destination. Third, the survey collects detailed evidence on consumption expenditures and self-produced items (particularly food) and creates a single consumption aggregate for each household. We describe below the advantages of consumption as a measure of living standards relative to income.

3.2. Consumption Measures and Urban-Rural Status

Our main outcome of interest is household consumption per adult equivalent member (or "consumption per adult" for simplicity). Consumption has a number of advantages relative to measures of income which have been used most often in this literature. One clear advantage is that consumption is measured in the same way for self-employed households as households that primarily earn wage income. To see why this may matter, suppose that a survey asks self-employed respondents to report their monthly income directly, and that households systematically under-report their actual income. Suppose further that wage workers report their monthly wage earnings with no bias (or less bias than the self-employed). Then a migrant moving from rural self employment to urban wage work will show up with a large spurious measured return from migration which actually largely captures an increase in the accuracy of income reporting.

Another related advantage of consumption relative to income as an outcome variable is that income may go toward higher expenses in urban. In principle income measures can simply be deflated by an urban-rural spatial price index. In practice, however, price indices are hard (or at least costly) to produce by surveyors, and often not included with surveys. In contrast the consumption surveys record many consumption for many goods in physical quantities, for instance kilograms of rice or maize. In other cases the consumption expenditures covered in our surveys are adjusted explicitly for rural-urban price differences using spatial price indices.

More generally, while consumption surveys are still face measurement challenges of various sorts, they are generally more widely believed than income as measures of living standards in the developing world (see e.g. Ravallion, 2003). There are also large gaps in consumption per worker between rural and urban areas (Young, 2013) and between the agricultural and non-agricultural sectors (Gollin, Lagakos, and Waugh, 2014), so there basic puzzle of large gaps is still clearly present in measures of consumption.

Consumption is measured similarly across all six or our countries, though with varying granularity and some clear differences in methodology for certain consumption items (as we describe below). Survey respondents report the monetary value of all goods and services purchased or consumed in the last two weeks, month, or (for a few select items) year. These values are then annualized, summed and divided by the number of adult equivalents. Importantly, we consider only contemporaneous questions about consumption expenditures. Pulido and Swiecki (2018) argue that using retrospective questions about income, as in the study of Hicks, Kleemans, Li, and Miguel (2017), leads to a downward bias in the estimated return to migration because of measurement error.

Our main independent variable of interest is a binary indicator for rural or urban status. Urban status comes from the geographic identifier variables provided by the surveys. In addition, following Hicks, Kleemans, Li, and Miguel (2017) we include controls for age and years of education in our regression specifications. Additional details on the measurement, timeline, and scope of each of our panel surveys are given below.

3.3. Panel Tracking Surveys and Characteristics

Table 1 presents some basic characteristics of our surveys. There are at least two surveys waves in each country, and as many as five in Indonesia. The surveys all contain a rich set of geographic detail. The Chinese survey is divided up into 176 different geographic communities, which is the lowest, and Tanzania (the highest) is divided up into 409 communities. The surveys collectively cover a very large number of individuals and households. A total of 183,169 individuals are surveyed across the countries and waves. Each individual is surveyed on average 2.6 times across all the surveys, leading to a total of 472,741 observations.

The countries in our surveys are still largely rural. This is almost a defining characteristic of being a developing country. Malawi has the highest percent of the population in rural areas, at 72 percent. Tanzania and Ghana are next at 62 percent each, followed by China at 54 percent, Indonesia at 53 percent and South Africa at 49 percent. There is no puzzle in these countries as to why people are not moving out of rural areas given the large urban-rural gaps there. People are moving rapidly to cities, and rural-urban migration rates are positive in all countries. We compute, but do not report, urban-rural migration rates, and these are everywhere smaller than rural-urban migration rates.

Table 1: Characteristics of the Panel Tracking Studies

	Waves	Communities	Individuals	Observations	P(Rural)	P(Migrate RU)
China	4	176	50,965	143,923	.54	.027
Ghana	2	334	15,964	26,786	.62	.016
Indonesia	5	296	48,184	131,803	.63	.011
Malawi	3	204	13,969	38,165	.72	.008
South Africa	5	400	38,430	104,090	.49	.025
Tanzania	3	409	15,887	35,103	.62	.029

Note: Columns 1-4 give the number of Waves communities (enumeration areas in uprban locations) adult individuals and individual-wave observations in each dataset. Column 5 gives the fraction of adults living in a rural location in wave 1. Column 6 gives the per-year rural-urban migration rate of adults in wave 1, annualized by the average number of years the sample was in the panel.

In what follows we summarize the six surveys in more detail, outlining their key similarities and differences.

China. The data come from the China Family Panel Study (CFPS). The CFPS is a longitudinal survey with waves in 2010, 2012, 2014, and 2016. The baseline sampling frame included 16,000 households across 25 provinces and aimed to represent 95% of the Chinese population. Enumerators attempted to track all migrant households and split-off families, and re-response rates from the previous year were 85.3 percent and 89.7 percent in 2012 and 2014 respectively. Household consumption was measured by asking about the value of 31 categories of items consumed. Food consumption accounted for only two of these 31 items (purchased and own-produced food).

Ghana. The data come from the EGC-ISSER Ghana Panel Survey. Data includes a total of 5,009 households representative of all 10 regions of Ghana collected in two waves in 2009 and 2013. Migrant and split-off households were tracked in between subsequent waves, and 87 percent of wave 1 households were re-interviewed in wave 2. The consumption module includes 85 food items and 53 non-food items. Households reported both the quantity and the value of each item consumed. We deflate our measure based on differences in prices in rural and urban locations. Urban-rural status is reported as a single binary variable.

Indonesia. The Indonesian data come from the Indonesia Family Life Survey (IFLS), and our variables used mirror those used by Hicks, Kleemans, Li, and Miguel (2017). Tracking was particularly successful in the IFLS re-interview rates across adjacent waves rounds are above 90% for all waves. Our cleaned panel differs slightly from the main dataset used in Hicks, Kleemans, Li, and Miguel (2017) in that we do not use any recall data. Our measure of rural-urban status is based on the rural-urban status of the households location in each wave according to the Indonesian Bureau of Statistics, rather than relying on self-reported prior location.

Malawi. The dataset comes from the Malawi Integrated Household Panel Survey (IHPS) which includes an initial sample of 3,104 household surveyed over three waves in 2010, 2013, and 2016. Enumerators attempted to track all individuals over 12 within mainland Malawi between waves 1 and 2. Between waves 2 and 3, enumerators attempted to track migrant households from a randomly selected half of the original enumeration areas. Tracking rates were high: attrition from 2013-2016 was only 4 percent, and attrition from 2010 to 2016 (among tracking enumeration areas) was 10 percent. The consumption module includes 136 food items and 120 non-food items, and our consumption aggregate is deflated for spatial price differences.

South Africa. Our South African data come from the National Income Dynamics Study (NIDS), which included an initial sample of 7,296 households surveyed over 5 waves from 2008 to 2017. Enumerators attempted to track all internal migrants between waves, and in total 73 percent of the wave sample remained in wave 5. The consumption module includes a total of 32 food items and 53 non-food items.

Tanzania. Our Tanzania data comes from the Tanzania National Panel Survey (NPS) which included an initial nationally representative sample of 3,280 households interviewed over three waves in 2008, 2010, and 2012. Enumerators attempted to track all internal migrants, and were highly successful; attrition from wave 2 to wave 3 was only 3.5 percent, and attrition from wave 1 to 3 was 4.8 percent. Consumption modules included a total of 72 food items and 46 non-food items, and our consumption aggregate is deflated by a spatial price deflator.

3.4. Results: Cross-Sectional Gaps vs Observational Returns to Migration

We first examine our estimates of the rural-urban consumption gaps across our six countries, which are presented in Table 2. Column (1) shows the cross-sectional gap, defined as the coefficient on an urban dummy in a regression of log consumption on an urban dummy and no other controls. Just as found in prior studies, the cross-sectional consumption gaps are large and statistically significant across all countries. The gap ranges from 42 log points in Ghana to 73 log points in South Africa, and the gap of 65 log points estimated for Indonesia is similar in magnitude to the 58 point estimate found by Hicks, Kleemans, Li, and Miguel (2017). The cross-sectional consumption gaps are generally smaller than those reported in Gollin, Lagakos, and Waugh (2014) and Young (2013).

Column (2) reports the results of estimated urban coefficient in regressions of log consumption per adult that also include year fixed effects and individual fixed effects are included. In these regressions, the urban premium is identified from consumption changes for individuals that are observed in both the urban areas and rural areas in different years. The urban coefficients here thus represent one estimate of the observational return to migration that are at the heart of the Hicks, Kleemans, Li, and Miguel (2017) and Alvarez (Forthcoming) studies. Column (2) shows that the observational returns to migration are substantially lower than the cross-sectional urban gaps, which consistent with the above studies. The urban coefficient falls from 55 to 19 log points in China, from 42 to 18 log points in Ghana, from 65 to 4 log points in Malawi, from 73 to 44 log points in South Africa and from 60 to 5 log points in Tanzania. Indonesia, which is also studied by Hicks et al, falls from 65 log points to 12 log points.

However, even with the addition of individual fixed effects, meaningful urban gaps remain in many countries. The simple average urban coefficient estimate across our countries gap is 17 log points. In South Africa, a 44 point gap remains even after the addition of individual fixed effects. Four of the six have statistically significant gaps (at the 5 percent level or lower) in Column (2): China, Ghana, Indonesia and South Africa. Malawi and Tanzania have positive remaining gaps but are statistically insignificant from zero.

Columns (3) and (4) repeat the panel fixed effect regressions of column (2) but with sample restrictions. Column (3) reports the urban coefficients with only individuals starting in urban areas. The urban coefficients in this column are therefore identified only through workers that move urban to rural areas. Column (4) presents the urban coefficients estimated with only workers that start in rural areas. These regressions identify the urban coefficient only through rural-urban migrants. As we show in the model section above, the return in (2) is a combination of the estimate in (3) and (4). In the event that workers moving from urban to rural areas have a consumption gain, as in the case of Malawi, with a negative urban coefficient of 37 log points, the coefficient in (2) will be biased toward zero. If one is interested in the gains to rural workers from moving to urban areas, the estimates in column (4) is more informative than the estimates in column (2) or (3).

As column (4) shows, none of the urban coefficients are approximately zero. The lowest is the coefficient from Indonesia, which was studied by Hicks et al., at 13 percent. Tanzania and Malawi are next at 16 and 20 percent, with the former statistically significant at the one-percent level and the latter insignificant. Ghana has a significant return of 22 percent and China a significant return of 25 percent. South Africa is the largest by far at 59 log points (which is closer to 80 percent). The unweighted average is 26 log points, and five of the six returns are statistically significant. The evidence here overall points to substantial observational returns to migration, not returns that are roughly zero. This is a substantively different conclusion from Hicks, Kleemans, Li, and Miguel (2017), and more similar to the observational returns found for Tanzania in Beegle, De Weerdt, and Dercon (2011) and the experimental returns of Bryan, Chowdhury, and Mobarak (2014).

A word is in order about hours controls, which Hicks, Kleemans, Li, and Miguel (2017) include in their preferred specification. We do not. One reason is that the majority of our surveys do not have reliable data on hours worked. But Pulido and Swiecki (2018) make a compelling case not to include hours control at all. They argue that potential migrants are not simply choosing between the higher of two fixed hourly wage rates as in simple two-sector models like the one in Gollin, Lagakos, and Waugh (2014). The choice is in fact the higher income and higher hours of urban areas versus the lower income and lower hours of the rural area. For welfare what matters of course is both consumption and leisure. Though one may wonder whether the value of leisure is very high at the margin for rural workers near subsistence consumption levels. In models of non-homothetic preferences, such as Stone-Geary, the marginal utility of consumption is so high for poor households that they are glad to forgo leisure in order to raise consumption. Empirically, the substantially higher average work hours in poor countries (Bick, Fuchs-Schündeln, and Lagakos, 2018) suggests that households in poor countries may not put much weight on leisure relative to consumption at the margin. If so, looking only at the level of consumption offered by the city, rather than consumption per hour worked, may better capture the outcome of interest for a poor rural individual considering migrating to the city.[2]

[2]Why don't rural workers just work a lot more hours then if the marginal value of leisure is so low and the marginal utility of consumption so high? One potential answer is that rural work is often subsistence self employment, which has a sharply decreasing marginal product of hours worked each day. The evidence of Bandiera, Burgess, Das, Gulesci, Rasul, and Sulaiman (2017) provide support for this story. They document

Table 2: Observational Returns to Migration in Six Developing Countries

	(1)	(2)	(3)	(4)
China	0.548^{***}	0.186^{***}	0.066	0.251^{***}
	(0.005)	(0.027)	(0.063)	(0.030)
Ghana	0.423^{***}	0.182^{**}	0.344^{***}	0.217^{**}
	(0.011)	(0.077)	(0.130)	(0.102)
Indonesia	0.646^{***}	0.121^{***}	0.114^{**}	0.134^{***}
	(0.006)	(0.024)	(0.051)	(0.028)
Malawi	0.520^{***}	0.037	-0.372^{***}	0.199
	(0.012)	(0.090)	(0.120)	(0.138)
South Africa	0.734^{***}	0.441^{***}	-0.002	0.588^{***}
	(0.006)	(0.023)	(0.048)	(0.027)
Tanzania	0.595^{***}	0.046	0.019	0.160^{***}
	(0.032)	(0.033)	(0.048)	(0.047)
Individual FE	No	Yes	Yes	Yes
Year FE	No	Yes	Yes	Yes
Sample	Full	Full	Start Urban	Start Rural

Note: The table presents the coefficients of an urban dummy variable for the six countries in our data. The dependent variable is log consumption per adult. Specifications in columns 2-4 include individual and wave fixed effects as well as controls for age and years of education. Standard errors are clustered at the level of the wave 1 household. Robust standard errors are in parenthesis. $^{*}p < .1, ^{**}p < .05, ^{***}p < .01$

3.5. Observational Returns to Migration by Region

We turn next to observational returns by rural region of origin. Our model predicts that migrants from regions with higher costs of migration will have higher returns to migration than those from regions with lower costs. Unfortunately, we cannot directly observe migration costs. We can however look at out-migration rates across all rural areas in each country. Our theory predicts that areas with higher out-migration rates are the ones with lower migration costs. We use this to test our model by examining whether average returns to migration among migrants from high-migration rural areas are lower than returns from low-migration areas.The number of rural migrant-sending communities in each of our countries ranges from 72 in Malawi to 198 in South Africa. Our approach is to estimate the returns to migration for migrants from the half of rural communities with the highest migration rates

that poor landless families in rural Bangladesh work very few hours on average due to lack of productive opportunities. Once they are given a transfer of livestock, hours worked rise substantially.

to the half with the lowest migration rates.

We present these results in Table 3. As our model predicts, across five of our six countries migrants from low-migration communities have higher returns to migration than those from high-migration communities. The difference is statistically significantly different from zero in China, Indonesia and Tanzania at the 10-percent level or lower. South Africa has a p-value of 0.108. The difference in Ghana is large but insignificant. In Malawi the model's prediction is exactly backwards, with higher returns from the high migration region, though with an insignificant difference. The simple mean of the returns across countries increases from 18 log points in high-migration communities to 30 points in low-migration communities. In China, the returns to migration for migrants from high and low-migration communities respectively are 7 and 32 log points. In Ghana, the returns are 12 and 29 log points (p = .39). In Indonesia, the returns are zero and 16 log points. In South Africa, the returns are 54 and 63 log points. In Tanzania, the returns are 8 and 24 log points.

It bears noting that these results are estimated only using communities where we observe rural-urban migration. In our median country, a third of our rural sample lives in a community that sent no urban migrants. Our theory predicts that the returns to migration are even higher for the 'marginal migrant' in these communities.

The results of Table 3 are largely consistent with the predictions of our model. In addition, they suggest that there are a meaningful number of non-migrants with both high cost and high potential returns to migration. Estimates of the returns to migration for status-quo migrants understate the gains from policies that facilitate migration for these high-cost individuals.

Table 3: Observational Returns to Migration by Region

	(1) High Migration	(2) Low Migration	(3) p-value: (1) = (2)
China	0.069	0.317***	0.000
	(0.044)	(0.043)	
Ghana	0.120	0.292***	0.391
	(0.173)	(0.112)	
Indonesia	0.002	0.155***	0.019
	(0.047)	(0.049)	
Malawi	0.265	0.144	0.659
	(0.207)	(0.183)	
South Africa	0.543***	0.633***	0.108
	(0.044)	(0.039)	
Tanzania	0.076	0.239***	0.071
	(0.058)	(0.070)	
Individual FE	Yes	Yes	
Year FE	Yes	Yes	
Sample	High Migration	Low Migration	

Note: The table presents the coefficients of an urban dummy variable for the six countries in our data. The dependent variable is log consumption per adult. Specifications are as in Table 2, column 4. The sample is split by the rural-urban migration rate in the origin community throughout the panel, based on the median origin migration rate among migrants. Standard errors are clustered at the level of wave 1 household. Robust standard errors are in parenthesis. $^{*}p < .1, ^{**}p < .05, ^{***}p < .01$

4. Observational versus Experimental Returns to Migration

While our results from Table 3 demonstrate that the observational returns to migration understate the returns to migration some non-migrants, it is ambiguous whether the observational gains to migration will overstate or understate the gains from any given migration-reducing policy. This is impossible to test in our panel datasets because we lack any policy experiment which exogenously lowered migration costs. Instead, we examine the difference between the observational and experimental returns to migration using data from Bryan, Chowdhury, and Mobarak (2014), an RCT that experimentally varied the costs of seasonal rural-urban migration at the household level in rural villages in northern Bangladesh.

The experiment in Bryan, Chowdhury, and Mobarak (2014) involved offering travel subsidies valued at roughly \$8.50 USD to rural households conditional on the household sending at least one migrant to an urban area during the agricultural lean season, a period between planting and harvest where rural labor demand falls. This subsidy effectively reduced the cost of migration and increased temporary rural-urban migration from 36 percent to 58 percent. The authors measure household migration and per-capita consumption during the lean season for all households in both 2008 (the year of the subsidy) and 2009.

This data gives us the rare opportunity to compare the observational returns to migration to the returns to migration to those induced to move by reduced costs. We can estimate the consumption gains for induced migrants by instrumenting for migration using the treatment. In addition, due to the panel nature of the data, we can also estimate the returns to migration for a subset of migrants that were migrating in the status quo, using the same individual-year fixed-effects regression we used in our main analysis.

Table 4 shows two estimates of the returns to migration in this context. The first, in column 1, estimates the returns to migration for status-quo migrants in the control group who switched between migrant and non-migrant status from 2008 and 2009. This non-experimental estimate includes household and year fixed-effects as in tables 2 and 3. In column 2, we show the Local Average Treatment Effect (LATE) of migration on consumption using treatment status as an instrument for migration. This estimates the returns to migration for migrants who are induced to migrate due to a reduction in migration costs. Column 3 shows the difference between these estimates, where we bootstrap the standard error of this difference by re-sampling households stratified by treatment status 1000 times. The difference between the estimates is an economically meaningful and statistically significant 27 log points; the returns to migration for policy-induced migrants in this context are some 300 percent greater than the returns for those that change from migrant to non-migrant in the status quo.

Admittedly, this is a very specific context which may not generalize well. Our experimental

Table 4: Observational vs. Experimental Returns to Seasonal Migration

	(1) Observational	(2) Experimental	(3) Difference
Seasonally Migrated	0.092*	0.357**	
	(0.053)	(0.156)	
year=2008	0.000		
	(.)		
year=2009	0.210***		
	(0.027)		
Difference in Returns			0.265***
			(0.095)
Constant	6.739***	6.652***	
	(0.022)	(0.084)	
Individual FE	Yes	Yes	
Year FE	Yes	Yes	
Observations	1194	1867	1867

Note: The dependent variable in the regressions is the log of consumption per adult. The data come from Bryan, Chaudhary, and Mobarak (2014). Column 1 is estimated using households in the control group with household and year fixed effects. Column 2 presents the Local Average Treatment Effect (LATE) of migration on consumption using treatment as an instrument for migration. Column 3 presents the difference between columns 1 and 2. Standard errors in columns 1 and 2 are clustered at the village level. The standard error in column 3 is computed from 1000 bootstrap replications. Robust standard errors in parenthesis. $^{*}p < .1,^{**} p < .05,^{***} p < .01$

sample involves only temporary rather than permanent migration, and includes households from a single district in northern Bangladesh. However, it is informative that in this one setting where we can observe both the observational and experimental returns, the observational returns significantly understate the gains to migration for those induced to migrate by reduced migration costs.

5. Conclusion

Recent evidence from several developing countries show that observational returns to rural-urban migration – i.e. returns for those that migrate even before any intervention – are not far above zero. One may wonder then if policymakers should give up on trying to induce even more rural households to move to cities. We argue that they should not. Some households may not already be migrating in equilibrium even if their potential returns would be high for the simple reason that migration costs may be higher. To the extent that many households face high moving costs, and policymakers can take steps to reduce them, this may still be a viable strategy to raise incomes for poor rural families.

So what are these migration costs, and what can policy makers do about them? This is the million dollar question. Alas, it not the question this paper has tried to answer. Fortunately, other researchers are working on this important question. Baseler (2019) argues that information frictions may serve as important barrier to migration. In his experiment in Kenya, an information treatment given to a random group of rural villagers increased the fraction of rural households sending a migrant to Nairobi from 20 percent to 28 percent. Morten and Oliveira (2019) analyze the effects of road building in Brazil and find that robust evidence that roads promote internal migration by lowering moving costs. Through the lens of a general equilibrium model of regional trade and migration, they estimate substantial welfare gains from reduced migration costs, with the largest gains for the most remote regions. Akram, Chowdhury, and Mobarak (2017) show that network effects could be another important migration barrier that can interact with other reductions in migration costs. Future work should continue to explore ways to lower migration costs in the developing world.

References

Akram, A. A., S. Chowdhury, and A. M. Mobarak (2017): "General Equilibrium Effects of Emigration on Rural Labor Markets," Unpublished Working Paper, Yale University.

Alvarez, J. A. (Forthcoming): "The Agricultural Wage Gap: Evidence from Brazilian Micro-data," *American Economic Journal: Macroeconomics*.

Bandiera, O., R. Burgess, N. Das, S. Gulesci, I. Rasul, and M. Sulaiman (2017): "Labor Markets and Poverty in Village Economies," *Quarterly Journal of Economics*, 132(2), 811–870.

Baseler, T. (2019): "Hidden Income and the Perceived Returns to Migration: Experimental Evidence from Kenya," Unpublished Working Paper, University of Rochester.

Bazzi, S., A. Gaduh, A. Rothenberg, and M. Wong (2016): "Skill Transferability, Migration, and Development: Evidence from Population Resettlement in Indonesia," *American Economic Review*, 106(9), 2658–98.

Beegle, K., J. De Weerdt, and S. Dercon (2011): "Migration and Economic Mobility in Tanzania," *Review of Economics and Statistics*, 93, 1010–1033.

Bick, A., N. Fuchs-Schündeln, and D. Lagakos (2018): "How do Hours Worked Vary with Income? Cross-Country Evidence and Implications," *American Economic Review*, 108(1), 170–199.

Bryan, G., S. Chowdhury, and A. M. Mobarak (2014): "Under-investment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh," *Econometrica*, 82(5), 1671–1748.

Bryan, G., and M. Morten (2019): "The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia," *Journal of Political Economy*, 127(5), 2229–2268.

Gollin, D., D. Lagakos, and M. E. Waugh (2014): "The Agricultural Productivity Gap," *Quarterly Journal of Economics*, 129(2), 939–993.

Herrendorf, B., and T. Schoellman (2018): "Wages, Human Capital, and Barriers to Structural Transformation," *American Economic Journal: Macroeconomics*, 10(2), 1–23.

Hicks, J. H., M. Kleemans, N. Y. Li, and E. Miguel (2017): "Reevaluating Agricultural Productivity Gaps with Longitudinal Microdata," NBER Working Paper 23253.

LAGAKOS, D., M. MOBARAK, AND M. E. WAUGH (2019): "Welfare Implications of Encouraging Rural-Urban Migration," Unpublished Working Paper, University of California San Diego.

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

MORTEN, M., AND J. OLIVEIRA (2019): "The Effects of Roads on Trade and Migration: Evidence from a Planned Capital City," Unpublished Working Paper, Stanford University.

NAKAMURA, E., J. SUGRUDSSON, AND J. STEINSSON (2019): "The Gift of Moving: Intergenerational Consequences of a Mobility Shock," Unpublished Working Paper, University of California Berkeley.

PULIDO, J., AND T. SWIECKI (2018): "Barriers to Mobility or Sorting? Sources and Aggregate Implications of Income Gaps across Sectors and Locations in Indonesia," Unpublished Working Paper, Vancouver School of Economics.

RAVALLION, M. (2003): "Measuring Aggregate Welfare in Developing Countries: How Well Do National Accounts and Surveys Agree?," *Review of Economics and Statistics*, 85(3), 645–652.

ROY, A. D. (1951): "Some Thoughts on the Distribution of Earnings," *Oxford Economic Papers, New Series*, 3(2), 135–146.

YOUNG, A. (2013): "Inequality, the Urban-Rural Gap, and Migration," *The Quarterly Journal of Economics*, 128(4), 1727–1785.