Migration Costs and Observational Returns to Migration in the Developing World

David Lagakos (UCSD and NBER)

Samuel Marshall (University of Warwick)

Ahmed Mushfiq Mobarak (Yale University, NBER, Deakin University and CEPR)

Corey Vernot (Yale University)

Michael E. Waugh (New York University and NBER)

March 2020

ABSTRACT —————————————————————————————————————

Recent studies find that observational returns to rural-urban migration are near zero in three developing countries. We revisit this result using panel tracking surveys from six countries, finding higher returns on average. We then interpret these returns in a multi-region Roy model with heterogeneity in migration costs. In the model, the observational return to migration confounds the urban premium and the individual benefits of migrants, and is not directly informative about the welfare gain from lowering migration costs. Patterns of regional heterogeneity in returns, and a comparison of experimental to observational returns, are consistent with the model's predictions.

———

1. Introduction

A growing body of evidence suggests that encouraging rural-urban migration may be a promising policy option for developing countries looking to increase aggregate productivity and reduce poverty rates. 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 lowering migration costs for rural workers to migrate seasonally to urban areas (Bryan, Chowdhury, and Mobarak, 2014; Akram, Chowdhury, and Mobarak, 2017). At least partly in light of evidence of this type, developing countries have begun to pay more attention to the implicit borders that divide rural and urban areas, and to strategies for facilitating movements of rural workers into more productive urban centers.

Yet other recent findings cast doubt on rural-urban migration as a viable growth strategy for developing nations. Using panel tracking studies from Kenya and Indonesia and an individual fixed effects specification, 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 (2020) finds a similar result in Brazil using a large panel survey. Relative to the large cross-sectional gaps, the gains for rural-urban migrants are an order of magnitude smaller. These studies join a chorus of other recent papers positing that rural-urban gaps are due to efficient sorting of workers across space within countries (e.g. Lagakos and Waugh, 2013; Young, 2013; Herrendorf and Schoellman, 2018). In short, these studies postulate that the people who work in urban areas are simply those with better skills for urban jobs, and that rural workers, who lack these same skills, will not benefit much from migrating to cities. The implication is that policy makers should not spend time trying to reduce migration costs for rural households to help them migrate.

In this paper we draw on new evidence and theory to reassess these conclusions and policy implications. On the empirical front, we draw on panel tracking surveys from six developing countries and the same fixed effects specification to expand the set of places for which observational returns to migration can be estimated. Using consumption measures as our main outcome of interest, we find observational returns to migration that are larger on average than those found in previous studies. In our preferred specification, returns to migration in Indonesia are similar to those of Hicks et al. (2017), but larger on average in our other five countries – China, Ghana, Malawi, South Africa and Tanzania – with an overall average return of 23 percent in our meta-analysis. Returns measured based on income are similar, for the subset of our countries for which income data are available. In all countries, observational returns to migration are substantially lower than the cross-sectional gaps, just as in the work of Hicks et al. (2017) and Alvarez (2020).

We then build an equilibrium model of migration to understand how observational returns to migration and cross-sectional gaps relate to the precise policy question that governments grapple with: *Would there be large benefits to lowering internal migration costs?* Feasible policy options that could affect the internal spatial distribution of people in large scale are likely related to lowering mobility costs. Our model builds upon the premise that migration costs – broadly defined – are quite heterogeneous across individuals. Examples abound. People who are geographically or linguistically isolated from urban areas face higher costs of migrating there than those living closer by (see e.g. Bryan and Morten, 2019; Bazzi, Gaduh, Rothenberg, and Wong, 2016). Rural individuals with strong networks in cities face a relatively lower cost of finding work and housing after migrating (Bryan et al., 2014; Akram et al., 2017), while those with greater ties to risk-sharing networks in villages have relatively higher costs of leaving those behind (Munshi and Rosenzweig, 2016; Morten, 2019; Meghir, Mobarak, Mommaerts, and Morten, 2019). Those who are poorly informed about the relative wages of urban centers have higher migration costs than those who are better informed (Baseler, 2019). Most strikingly, China's *hukou* policy explicitly erects greater barriers to internal migration for the low-skilled compared to the high-skilled (Khanna, Liang, Mobarak, and Song, 2019).

Our model of migration is a multi-region Roy (1951) model in which workers are heterogeneous in their comparative advantage in the urban location, and face differing migration costs depending, for simplicity, on the region from which they migrate. The urban region offers a productivity premium for all workers, and rural-urban migrants face skill loss from migration making them less productive than urban incumbents. We use the model to relate the cross-sectional consumption gap, the observational returns to migration by region, and the welfare gains from lowering migration costs.

The model predicts that observational returns to migration may not be as informative about migration policy as previously thought. In our general formulation, the observational return is just the cross-sectional gap plus an adjustment for the skill composition of the migrants. Migrants are not on the margin in the model, and migrate whenever their individual return from migrating exceeds their cost. The observational return to migration thus reflects not just the urban premium but the average productivity of those for whom the benefit from migration exceeds the costs. One cannot, therefore, draw inferences about the urban productivity premium without knowing the joint distribution of costs and benefits of migration across workers.

The model does, however, predict a different use of the observational returns *by region*. In the model, comparing the observational returns in two different regions reveals differences in the extent to which migrants are selected or not across regions. For example, workers from both high and low migration cost regions share a common benefit of migration, but only the most productive workers from high-cost region will sort into the urban area, whereas many relatively less productive workers from the low-cost region migrate. Under standard—and ad-

2

mittedly restrictive—parametric assumptions, these differential observed returns exactly equal the difference in migration costs between regions. The upshot is that variation in observational returns by region may be informative about the underlying differences in migration costs.

We use the model and a utilitarian social welfare function to derive simple expressions for the welfare gains to lowering migration costs in any given region. In the most general formulation, the welfare gains take a simple form: the mass of migrants multiplied by the average productivity of those migrants. This result is surprisingly intuitive: if a region sends more migrants to begin with, then lowering migration costs there will help more people. If the migrants from a particular region are particularly productive, then lower migration costs for that region will do more to raise aggregate productivity. How do observational returns relate to these welfare expressions? Unfortunately, it is not obvious. Per se, the observational returns to migration may not be that informative about the gains from lowering migration costs. Under parametric assumptions, we do illustrate how relative differences in observed returns (as we emphasized above) help identify regions in which lower migration costs may lead to relatively high returns from a social welfare perspective.

Both our analysis of the observational returns and gains to lower migration costs suggest studying how observational returns vary across regions in the data. We do so by examining heterogeneity in the observational returns to migration across regions within each country in our data. Consistent with the model, we see that observational returns are generally higher in regions where fewer people are currently migrating. When aggregating across countries, the difference in average returns between high- and low-migration regions is a statistically significant 12 log points. Viewed through the lens of the model, this empirical result suggests that lying behind the modest observational returns calculated across all regions are situations in which the returns for some individuals are large. Combined with the model's welfare expressions, the implication is that the gains from lowering migration costs may be high in some areas, even if observational returns are modest on average.

We conclude by turning to experimental evidence on the effects of lowering migration costs, drawing on data from the travel subsidy experiment of Bryan et al. (2014) that randomly lowered migration costs for some rural Bangladeshis. These data are unique in that they allow us to directly compare experimental estimates of the effect of a policy lowering migration costs to observational estimates for the exact same population derived using the same type of panel regression with individual fixed effects used in the recent literature. The observational return to migration – estimated using year-to-year variation in migration decisions in the control group – is only 9 log points. In contrast, the experimental return (to lowering migration costs) is 36 log points. The difference between the two estimates is large and statistically significant at the one-percent level. The sharply larger experimental returns to migration compared to the observational returns supports our paper's conclusion that estimated observational returns may not

paint an accurate picture of the gains from lowering migration costs.

This paper builds on a large literature about rural-urban migration in the developing world, and necessarily abstracts from many important issues. More comprehensive treatments can be found in Lucas (1997) and Brueckner and Lall (2015). Our focus on cross-sectional gaps versus panel returns to migration follows studies in the United States conducted by Glaeser and Maré (2001) and Herrendorf and Schoellman (2018). In the developing-county setting, our estimates build on those of Beegle, De Weerdt, and Dercon (2011), Dercon, Krishnan, and Krutikova (2013) and Garlick, Leibbrandt, and Levinsohn (2016) in addition to the studies mentioned above. Our work also follows the earlier literature focused on identifying Roy models of selection, in particular the work of Heckman and Honoré (1990) and Heckman and Vytlacil (2005); French and Taber (2011) provide an overview of the various approaches this literature has taken. Our empirical finding of larger experimental returns than observational returns to migration is broadly consistent with the evidence from the natural experiments studied by Nakamura, Sigurdsson, and Steinsson (2019) and Sarvimäki, Uusitalo, and Jäntti (2019), in which individuals forced to migrate experienced substantial returns on average relative to those staying behind.

2. Observational Returns to Migration: New Evidence

In this section, we draw on data from panel tracking surveys covering six developing countries. We use these data to estimate the observational returns to migration, by which we mean the urban consumption premium estimated in regressions with individual fixed effects. We confirm the finding of Hicks et al. (2017) that cross-sectional gaps are larger than observational returns to migration in all of our countries. However, we find larger returns in the five countries not studied by Hicks et al. (2017) than the one they studied that is also in our sample (Indonesia).

2.1. Data Overview

Our data come from household panel surveys conducted in six developing countries: China, Indonesia, Ghana, Malawi, South Africa, and Tanzania. The surveys all share three common characteristics necessary for our analysis. First, they attempt to track and re-survey all individuals that migrate between survey waves. This tracking of survey respondents across space, which is very costly, is necessary to reliably estimate the returns to moving across locations while controlling for unobserved time-invariant characteristics as in Beegle et al. (2011), Hicks et al. (2017) and Alvarez (2020). Second, the surveys collect detailed geographic information on each individual in each survey wave, allowing us to classify all individuals by urban-rural status and then further by more disaggregated region. This allows us to estimate gains from rural-urban migration overall and also by source-destination pair. Third, the surveys collect detailed evidence on consumption expenditures and self-produced items, particularly food,

which allows for the creation of a comprehensive consumption aggregate for each household. We conducted a broader search for viable panel datasets, and settled on these six countries because the available data met these criteria.

2.2. Consumption Measures

Our main outcome of interest is household consumption per adult equivalent member (or "consumption per adult" for simplicity). Consumption is, for the most part, measured similarly across our set of 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 et al. (2017), leads to a downward bias in the estimated return to migration because of measurement error.

Consumption has a number of advantages relative to the more popular income measures. Importantly, consumption is measured in the same way for both the self-employed and wage earners. To see why this may matter in our application, suppose that the self -employed are more likely to under-report income in surveys relative to wage workers, who may receive consistent paychecks each week that make their exact earnings more salient. Then a migrant moving from rural self employment to urban wage work will display a large spurious measured return from migration, largely driven by improvements in the accuracy of income reporting.

Another important advantage of consumption as an outcome variable is that urban income may go toward higher expenses in cities. In principle, income measures can simply be deflated by an urban-rural spatial price index. In practice, however, price indices are difficult and costly for surveyors to produce, and often not included with surveys. In contrast, the consumption surveys record consumption for many goods in physical quantities, such as 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.

There are several challenges to creating fully comparable measures of consumption across urban and rural locations. First, a larger fraction of rural food consumption comes from crops produced by the household. All of our consumption measures include questions about the value of food consumed from the household's own farm. Second, some durables are purchased infrequently, and may not get captured in annualized expenditure surveys. Our consumption aggregates for five countries excludes expenditures on durables, while in Malawi we include the annualized use value of durables. A third complication is correctly inferring the value of

housing. For rental housing, this can be estimated with actual expenses on rent. In our consumption aggregates, housing expenditure from both rental and owner-occupied housing are included for Indonesia, Malawi, and South Africa; housing value is excluded entirely for South Africa and Ghana. The consumption aggregate for China includes rental housing but not the value of owner-occupied housing. Table 1 summarizes the characteristics of our consumption aggregates.

While consumption surveys face various measurement challenges, they are widely perceived as a more accurate measure of living standards in the developing world compared to income (see, e.g., Ravallion, 2003). The literature has documented large gaps in consumption per worker between rural and urban areas (Young, 2013) and between the agricultural and non-agricultural sectors (Gollin et al., 2014), so the basic puzzle of internal spatial gaps is clearly present in measures of consumption as well.

2.3. Urban-Rural Status and Other Controls

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. We follow Hicks et al. (2017) and use the self-reported coding of locations as either villages, small towns, or cities in Indonesia. We choose to treat cities as urban areas and the rest as rural, though our results are not substantively different if we include small towns as urban. In addition, following Hicks et al. (2017) we include controls for age and years of education in our regression specifications. Following Pulido and Swiecki (2018) we include controls for household size and household size squared, though our results are robust to excluding these controls. Additional details on the measurement, timeline, and scope of each of our panel surveys are given below.

2.4. Panel Tracking Surveys and Characteristics

Table 2 presents some basic characteristics of our surveys. There are at least two survey 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 smallest number of geographic divisions, and Tanzania, with the most divisions, is divided up into 409 communities. These surveys collectively cover a total of 183,169 individuals across countries. Each individual is surveyed on average 2.6 times, 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. People in all six countries are moving

rapidly to cities, and net rural-urban migration rates are positive everywhere. We compute, but do not report, the converse urban-to-rural migration rates, and these are everywhere smaller than rural-urban migration rates.

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

China. The China Family Panel Study (CFPS) is a longitudinal survey with waves in 2010, 2012, 2014, and 2016 (Institute of Social Science Survey, 2015). The baseline sampling frame included 16,000 households across 25 provinces and aimed to represent 95 percent 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 consists of just two comprehensive categories: purchased food and own-produced food items.

Ghana. The Yale EGC-ISSER Ghana Panel Survey includes a total of 5,009 households representative of all 10 regions of Ghana, with the data collected in two waves in 2009 and 2013. Migrant and split-off households were tracked in between subsequent waves, and 87 percent of first wave households were re-interviewed in the second wave. 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 consumption using differences in prices in rural and urban locations. Urban-rural status is reported as a single binary variable.

Indonesia. We use the Indonesia Family Life Survey (IFLS), and our variables used mirror those used by Hicks et al. (2017). Tracking was particularly successful in the IFLS, with re-interview rates across adjacent waves rounds exceeding 90 percent for all waves. Our cleaned panel differs slightly from the main dataset used in Hicks et al. (2017) in that we do not use any recall data. All data is measured contemporaneous to the survey date. For urban status, we use both the measure used by Hicks et al. (2017) which includes small towns and cities, and our preferred measure which includes only cities. Our main outcome of interest is consumption per adult but we also consider income per adult. We measure income per adult as wage income across all household members plus business income plus farm profits.

Malawi. The Malawi Integrated Household Panel Survey (IHPS) 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. The South African National Income Dynamics Study (NIDS) 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 initial wave sample remained in wave 5. The consumption module includes a total of 32 food items and 53 non-food items.

Tanzania. The Tanzania National Panel Survey (NPS) 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.

2.5. Results: Cross-Sectional Gaps and Observational Returns to Migration

Estimates of the rural-urban consumption gaps across all six countries are presented in Table 3. Column (1) shows the cross-sectional gap, defined as the coefficient on an urban dummy in a regression of log consumption on that urban indicator 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 41 log points in Ghana to 74 log points in South Africa, and the gap of 63 log points estimated for Indonesia is similar in magnitude to the 58 point estimate found by Hicks et al. (2017). The cross-sectional consumption gaps are generally smaller than the gaps in living standards reported in Gollin et al. (2014) and Young (2013).

Column (2) reports the same urban-dummy coefficients, but adds year fixed effects and individual fixed effects to the regressions. The urban premium is now identified from consumption changes for individuals that are observed in both urban areas and rural areas in different survey waves. The urban coefficients here thus represent one estimate of the observational return to migration that are at the heart of the Hicks et al. (2017) and Alvarez (2020) studies. Column (2) shows that the observational returns to migration are substantially lower than the cross-sectional urban gaps, which is consistent with those studies. The urban coefficient falls from 55 to 16 log points in China, from 41 to 15 log points in Ghana, from 52 to 5 log points in Malawi, from 74 to 21 log points in South Africa and from 67 to 11 log points in Tanzania. The Indonesian urban coefficient falls from 63 log points to 15 log points.

Even with the addition of individual fixed effects, meaningful urban gaps remain in many countries. The median urban coefficient estimate across our countries gap is 14 log points. In South Africa, a 21 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 indistinguishable from zero.

An important deficiency of column (2) is that it groups together rural-to-urban and urban-to-rural migrants. If people always move in search of a better life – i.e. when the net returns to moving are positive – then the coefficient on the urban dummy will be biased towards zero on account of the reverse urban-to-rural movers. To address this, columns (3) and (4) repeat the same fixed effects regressions, but simply split the sample into those who started out in the city in the first wave of data (column 3), versus those that start out in the rural area (column 4). Column (4) identifies the urban coefficient using only rural-to-urban migrants, which most closely resembles the conceptual underpinning of the spatial productivity gaps literature. Column (4) displays large gains to urban moves in each of the six countries in our sample. The smallest coefficient is the 17 percent in Indonesia, which was the subject of the Hicks et al study. Malawi and Tanzania are next at 19 and 21 percent, with the latter statistically significant at the one-percent level and the former insignificant. China has a statistically significant return of 23 percent and South Africa a significant return of 29 percent. Ghana is the largest at 34 log points. The unweighted average is 23 log points, and these positive returns are statistically distinguishable from zero with 95 percent confidence in five of the six countries. When pooling across all countries using a random-effects model that is standard in meta-analysis, this average return is significant at the 0.1 percent level.[1]

We see a very similar pattern when we replace our consumption measure with household income in Table 4. In Indonesia and South Africa, the two countries for which we could reliably construct household income, the gains look very similar. The per-capita-income gain for rural-to-urban movers in Indonesia is 15 log points in column (4), comparable to the 17 log points gain in consumption. The trend is the same in South Africa. The overall income gains of 25 log points in the fixed effects regression are entirely driven by rural-to-urban movers, who experience a 29 percent increase in household income per adult.

A word is in order about hours controls, which Hicks et al. (2017) include in their preferred specification. We do not. One reason is that some of our surveys do not have reliable data on hours worked. More importantly, 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 et al. (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. Both consumption and leisure of course matter

---

[1]Consistent with the simple theory that people move to better opportunities, the urban coefficient in column (3) – which is estimated based only on people who move *away from* cities - is often *negative* or zero. For example, the large negative urban coefficient for Malawi indicates that when the first wave of the survey observes an individual in the city she *gains* 35 percent in consumption when she subsequently chooses to move to a rural area.

for welfare. However, one may wonder whether the value of leisure is very high at the margin for rural workers near subsistence consumption levels in the poor countries and households in our sample. Empirically, the substantially higher average work hours in poor countries (Bick, Fuchs-Schündeln, and Lagakos, 2018) suggest that those households may not put much weight on leisure relative to consumption at the margin. If so, the *level* of consumption in the city may better capture the outcome of interest for a poor rural individual considering migrating to the city rather than consumption per hour worked.

Overall, the evidence in this section points to substantial observational returns for those moving from rural to urban areas. This is a substantively different conclusion from Hicks et al. (2017), and more similar to the observational returns found for Tanzania by Beegle et al. (2011), for South Africa by Garlick et al. (2016), or for India by Dercon et al. (2013).

2.6. Differences with Previous Literature

Our preferred estimates of the urban premia in Indonesia differ to some extent from those found in Hicks et al. (2017) and Pulido and Swiecki (2018), and we explain the reasons for these differences here. As in the analysis Pulido and Swiecki (2018), and unlike that of Hicks et al. (2017), we do not use recall data in any of our variables, and we examine household consumption and income per capita rather than per hour. The remaining differences between our estimates and those of Pulido and Swiecki (2018) and Hicks et al. (2017) can be accounted for by our choices of outcome measure and the types of urban destinations we include. These differences are shown in the different rows of Table 5, where we vary the types of urban destinations included (cities only vs. cities and small towns) and also the outcome measure (individual income vs. household income per capita vs. consumption per capita). Our approach in row 4 most closely match the strategy of Hicks et al. (2017) and Pulido and Swiecki (2018), and we see that our estimate of 6.5 log points in column (2) is similar to the main estimates of those papers.

In addition, we can see two trends clearly: the gains to consumption and household income are larger than the gains to individual income, and the gains to migration to cities are larger than the gains to migrating to small towns. The differences between household-measured outcomes and individually-measured outcomes may be due to a larger sample being included in the former. We lack individual wage income for many individuals who reside in families where subsistence agriculture is their only livelihood. It is plausible that there are larger gains for individuals who move from subsistence agriculture to wage work than for individuals who move from wage work in a rural area to wage work in an urban area. Indeed, Pulido and Swiecki (2018) find that rural-urban wage premia in Indonesia are largely driven by those who switch out of agriculture when they move.

10

3. Model of Heterogeneous Costs and Benefits of Migration

The results in Table 3 provoke the following questions: How should one interpret these observed returns to migration? And what do they say about the gains from lowering internal migration costs? In this section we attempt to answer these questions using a stylized model of migration in which individuals differ in returns to migration – as in most of the previous literature – as well as in costs of migration, which evidence (e.g. the studies cited in the introduction) suggests is important.

3.1. Model Environment

Regions and Migration Costs. The economy is composed of a unit measure of individuals who live for two periods. There are many regions, each with a competitive labor market. Amongst these regions there is one urban region, indexed by u. The remaining regions are rural and indexed by r_i. A share λ_u of individuals is born in the urban region; of those born in the rural area, a share λ_{ri} are born into rural region i. The total population is $\lambda_u + \sum_i \lambda_{ri} = 1$.

The rural regions differ in their costs of migrating out of the region. Individuals who choose to migrate out of rural region i lose a fraction of their income, m_i, where $0 < m_i < 1$. We assume that migration is only possible at the end of the first period, so that we can observe the income of migrants both before and after migration.

Endowments and Preferences. Individual utility is linear in consumption, which equals individual income net of migration costs. Individuals are also endowed with one unit of time in each period which they supply inelastically to the labor market in which they are located. Following Roy (1951), we assume that individuals are also endowed with a bivariate skill vector, $\{z, 1\}$, representing the efficiency of their labor in the urban and rural locations. The simplifying assumption here is that the skill vector—and hence the pattern of selection—is "one sided" in the sense that individuals are only heterogenous in productivity in one location.

The distribution of urban skills is given by a cumulative distribution function, $F(z)$, and probability density function, $f(z)$, with support on the positive reals. We make the standard assumptions on these distributions that moments exist and that they are continuous and differentiable. We will also express outcomes in terms of the conditional distributions given where $f(z|Z \geq \hat{z})$ is the conditional density function given threshold $\hat{z}$. In the discussion below, we present general results without appealing to distributional assumptions. We also provide examples using a Pareto distribution: $F(z) = 1 - z^{-\theta}$, where the minimum of the support is normalized to the value one, and θ controls the variation in skills, with a lower θ implying more dispersion in skills.

A final source of heterogeneity arises from place of birth. We also assume that individuals are

more productive in their location of birth, by a factor $\psi > 1$. This assumption is motivated by Bazzi et al.'s (2016) evidence that exactly this kind of "local knowledge" prevents skills from being transferable across regions. In our setting, individuals born in the urban area with a particular z draw will be more productive than rural-urban migrants with that same z draw by a factor ψ.[2]

Technology. Competitive firms in each region use labor services as the sole input to produce a single consumption good, Y. The production function in the urban and rural areas are given by:

$$Y_u = A\,L_u \qquad \text{and} \qquad Y_{ri} = L_{ri}, \tag{1}$$

where

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

Ω^l denotes the set of individuals employed in labor market l, L_l is the number of efficiency units employed in region l, and N_l is the total number of workers employed in region l. The appeal to linear production technologies simplifies matters in the sense that wages per efficiency unit will be pinned down by technology parameters.

3.2. Optimization, Equilibrium, and Social Welfare

Below we describe the choices of workers and the associated equilibrium conditions more formally.

Firms. Given the technologies described in (1) and the competitive market structure, optimization implies that each labor market is characterized by a wage per efficiency unit:

$$w_u = A \qquad \text{and} \qquad w_{ri} = 1 \qquad \text{for all } i. \tag{2}$$

In other words, wages per efficiency unit equal the marginal products of labor.

Workers. Workers make a migration decision at the end of period one to maximize their earnings in the second period. For simplicity, we assume that those in the urban location stay there. Because of the local knowledge each individual is born with, no rural individual will move to any other rural region. Thus, migration decisions are only about potential moves from rural areas to the urban area. An individual with skill vector $\{z, 1\}$ chooses their location to maximize their utility:

$$\max\{mw_u z\,,\ \psi\}. \tag{3}$$

---

[2]While the urban incumbent productivity advantage is assumed exogenously here, it could be captured in 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.

Market Clearing. In each period, market clearing in the labor and output markets requires that:

$$N_u + \sum_i N_{r,i} = 1 \quad \text{and} \quad Y_u + \sum_i Y_{r,i} = C, \tag{4}$$

so that labor is fully employed in production and aggregate output equals aggregate consumption.

Definition 1 *A Competitive Equilibrium consists of location choices and wages per efficiency units $\{w_u, w_{ri}\}$ that satisfy the following conditions:*

- *Wages per efficiency units $\{w_u, w_{ri}\}$ satisfy firm optimization, (2).*

- *Given wages, locations choices are optimal, (3).*

- *Markets clear, (4).*

We can now characterize some properties of the equilibrium. Note that the workers' location choices in (3) take a very simple form, where migration to the urban area is optimal if and only if:

$$m_i A z \geq \psi. \tag{5}$$

This implies that there is a single cutoff value $\hat{z}_i = \frac{\psi}{m_i A}$ determining which workers migrate to the city and which remain in the rural area.

This single cut-off property is useful because it allows us to characterize social welfare (under a utilitarian welfare function with equal weights) in a very simple fashion. Furthermore, because there is no optimization for the first period, we can focus on outcomes in the second period. Given these observations, the social welfare function is:

$$\mathcal{W}(\mathbf{m}, \hat{z}(\mathbf{m})) = \lambda_u \int \psi A z f(z) dz + \sum_i \lambda_{ri} \left\{ (1 - F(\hat{z}_i)) \left[\int_{\hat{z}_i} m_i A z f(z|Z \geq \hat{z}_i) dz \right] + \psi F(\hat{z}_i) \right\}, \tag{6}$$

which we will write as a function of the migration costs, $\mathbf{m}$, and the cutoffs $\hat{z}(\mathbf{m})$, which depend on the migration costs indirectly. This formulation will be useful below when we characterize the gains from lowering migration costs.

It is useful to walk through each term in (6) carefully. The first term represents the welfare of urban incumbents: the number of people born in the urban region multiplied by the integral over each urban incumbent's utility. The utility of each urban incumbent is their wage in the

urban area, which equals the acquired skill of that location, ψ, times aggregate productivity in the urban location, A, times the worker's innate skill in the urban area, z.

The second term in (6) represents the welfare of those in the rural regions, which is the sum of the utility of those who migrate and those who stay. For a given region, i, the term $(1 - F(\hat{z}_i))$ represents the mass of migrants who have productivity above $\hat{z}$, and hence migrate. This term is then multiplied by the average utility of these migrating agents (hence, the conditioning on values above $\hat{z}$). The utility of these migrating agents is the migration cost, m, times the wage in the urban area, A, times z. The second component inside the brackets represents the expected utility those who remain in the rural area, i.e. the mass of those with productivity below $\hat{z}$ times their utility, which is just ψ.

3.3. A Story of Three "Returns"

To what extent are commonly measured "returns" to migration informative for migration policy? To what extent do the gains from reducing migration costs relate to measures of potential and actual returns to migration that one may observe in cross-sectional or panel data? The model above allow us to think through these questions in the context of three different measures of returns to migration: (i) the cross-sectional urban-rural earnings gap, (ii) the observational returns to migration, estimated using panel fixed effects regressions, and (iii) the gains from lowering migration costs, which may be estimated in some settings using experimental evidence.

3.3.A. The Cross-Sectional Gap

The cross-sectional gap is the average consumption in the urban area compared to the average consumption in the rural area (as in Gollin et al. (2014), Young (2013) and Herrendorf and Schoellman (2018)). In a regression context, the cross-sectional gap is analogous to the OLS coefficient on an urban dummy in a regression of log consumption or income on the urban dummy and possibly other cross-sectional covariates. In period one, the cross-sectional gap (in logs) takes a simple representation in our model described below.

Proposition 1 *The **Cross-Sectional Gap** is given by*

$$\text{CSG} = \log A + \int \log z f(z) dz, \tag{7}$$

which with Pareto distributed skills is

$$= \log A + \frac{1}{\theta}. \tag{8}$$

Equation (7) tells us that the cross-sectional gap reflects two forces: (i) the effect of being in the urban area—captured by A—and (ii) the average skill level of those in the urban area, represented by the second term. This formulation exactly picks up the idea that the cross-sectional gap is a mixture of both the "treatment effect" of being in the urban area and inherent differences in the workers between urban and rural areas. Of course, the former is causal from the perspective of migration policy while the latter is not.

The example with Pareto distributed skills further develops this intuition. In the Pareto case, the parameter θ controls how dispersed skills are in the urban area. As θ becomes larger, workers become more homogeneous, and the cross-sectional gap reflects more of the pure urban productivity effect and less of the worker selection effect.

3.3.B. Observational Returns to Migration

To try to separate the effects of worker selection from the pure urban productivity effect, studies since Glaeser and Maré (2001) have turned to panel regressions that estimate the urban coefficient while controlling for individual worker fixed effects. This is the approach taken, in a developing country context, by Beegle et al. (2011), Dercon et al. (2013), Garlick et al. (2016), Hicks et al. (2017), and Alvarez (2020), among others, before the current study. The idea is that if one sees the same worker's outcomes in two locations, perhaps one can learn something about the intrinsic benefit of the urban area by "differencing out" the worker effect. We will refer to the return to migration estimated using such a fixed-effects strategy as the observational returns to migration.

We can use our model to interpret the observational returns to migration. To do so, we first define the return, R, for an individual worker as the difference in log consumption (or earnings) between the urban and rural areas:

$$R = \log A + (\log z - \log(1)) - \log \psi. \tag{9}$$

For those that actually work in the rural area in period one and the urban area in period two, this is an actual return to migration. For those that do not move this is a potential outcome. The first term in this return is the benefit of being in the urban area, A, relative to the rural area. The second term, $\log z - \log(1)$, is about the relative productivity advantage that the individual enjoys in the urban area, which reflects selection on comparative advantage. The third term reflects losses in the local knowledge that rural-urban migrants give up when they leave their home areas behind.

Looking at (9) reinforces why one would want to look at a return to migration for one individual—it partially reveals the productivity advantage of being in the urban area, A. The downside is that it is contaminated with other issues—selection on individual ability and skill loss. The

hope would be that these issues do not matter much and that taking an average across many migrants may be mostly informative about the pure urban effect. Averaging across all migrants in the model gives the observational return to migration, which we characterize below.

Proposition 2 *The **Observational Return to Migration** is given by*

$$\text{ORM} = \log A + \sum_i \frac{\pi_{uri}}{\pi_{ur}} \left(\int_{\hat{z}_i} \log z f(z | Z \geq \hat{z}_i) dz - \log \psi \right) \tag{10}$$

where $\pi_{uri} = \lambda_{ri}(1 - F(\hat{z}_i))$ is mass of rural to urban migrants from rural region i and $\pi_{ur} = \sum_i \pi_{uri}$ is the total mass of migrants. The observational return is then the Cross-Sectional Gap plus a skill adjustment for migrants:

$$= \text{CSG} + \sum_i \frac{\pi_{uri}}{\pi_{ur}} \underbrace{\left\{ \left(\int_{\hat{z}_i} \log z f(z | Z \geq \hat{z}_i) dz - \int \log z f(z) dz \right) - \log \psi \right\}}_{\text{skill adjustment for migrants}}. \tag{11}$$

The first expression for the observational return in Proposition 2 integrates over the return in (9) conditional on migrating. The observational return to migration reflects urban productivity, the selected productivity of those migrating, and then skill loss. The next result in Proposition 2 is the surprise. It says that the observational return is just the cross-sectional gap plus a skill adjustment for those who migrate in equilibrium (see Appendix B for a formal derivation). The skill adjustment reflects (i) how skilled the migrants are relative to the incumbents (the term inside the parenthesis), and (ii) the skill loss from moving away from the rural area.

The issue here is that the observational return to migration still reflects both the productivity effect of the urban area and the composition of who migrates in equilibrium. It is, in other words, no panacea for stripping out the pure urban effect from the cross-sectional gap. In fact, the observational return embeds the issues of treatment and selection in a way that is not much different than what shows up in the cross-sectional gap. The intuition for this outcome—and we discuss alternative perspectives on this below—is that, in our model, migrants are not marginal. Because of the assumption that people are not optimally located in period one, there are infra-marginal migrants who started out in the wrong place (the rural area). As these migrants move in to the urban area and their returns get measured, the entire distribution of gains are traced out up to a correction for selection and skill loss.

One perspective on these observational returns is that they should also have an implication of for the size of the migration cost (our discussant Schoellman (Forthcoming) articulates this view well). In particular, if the observed returns are small, then migration costs must be small and vice versa. Under the assumption of Pareto distributed skills we can get some more insight

about this perspective. Under this assumption we have

$$\text{ORM} = \frac{1}{\theta} - \sum_i \frac{\pi_{uri}}{\pi_{ur}} \log m_i, \tag{12}$$

which helps illustrate the general point that the observational return mixes several forces. Yes, the observational return does contain information about the migration costs, m_i, in the rural regions. However, the migration costs are intermingled with the important of selection, as the prominent role for parameter θ shows. In other words, it is not so easy to draw sharp inferences about the size of the migration costs without knowing more about the importance of selection.

With that said, the model does, however, predict a different use of the observational returns *by region*. For simplicity, suppose that there were simply two regions: a high cost location, indexed by h, and a low cost location, indexed by l. If we compare observational returns across these two regions we have

$$\text{ORM}_h - \text{ORM}_l = \int_{\hat{z}_h} \log z f(z|Z \geq \hat{z}_h) dz - \int_{\hat{z}_l} \log z f(z|Z \geq \hat{z}_l) dz \tag{13}$$

which says that relative observational return reveals the relative productivity of those migrating from one region to another. Comparing observational returns in two regions differences out the common component related to the cross-sectional gap. What remains are the differences in the extent to which migrants are selected or not across regions.

The standard intuition behind the conditional expectation in (13) is that it is increasing in the conditioning point. Thus, only most productive workers from high-cost region will sort into the urban area, and their observed returns will be high. In contrast, many relatively less productive workers from the low-cost region migrate urban area, and their observed returns will be low. Again, per se the observed returns are contaminated with many things, but the difference reveals the extent to which a region may be a high or low return region. Moreover, as we discuss below, inputs of this nature are useful in identifying the relative welfare gains from lowering migration costs in different regions.

With Pareto distributed skills, differences in observational returns exactly reveal differences in migration costs:

$$\text{ORM}_h - \text{ORM}_l = [\log m_l - \log m_h] > 0. \tag{14}$$

Returning to the discussion around (13), why would observed returns be relatively large in high moving cost locations? Because the cost of migration is large and only those who are most able to overcome the cost are observed to be migrating. Put differently, variation in observational returns to migration by region may be informative about the underlying differences in

migration costs and their relative importance.[3]

3.3.C. Gains from Lower Migration Costs

The specific policy question many care about is: what are the welfare gains from lowering internal migration costs? Should a policymaker turn to the observational return to migration for guidance on this matter? The short answer in our model is no. Below, we use the model and the social welfare function in (15) to derive a connection between observed flows of migrants and the extent to which migrants are selected.

Given our social welfare function in (6), we compute change in social welfare given a change in migration costs from region i. We express this derivative as:

$$\frac{\mathrm{d}\mathcal{W}}{\mathrm{d}m_i} = \frac{\partial W}{\partial m_i} + \frac{\partial W}{\partial \hat{z}_i}\frac{\partial \hat{z}_i}{\partial m_i}, \tag{15}$$

where the total derivative is decomposed into a direct effect from a reduction in migration costs and an indirect effect that works through the change in the equilibrium cut-off. This decomposition is useful from a mechanical perspective – it's easier to compute – and from an economic perspective, as it develops intuition.

Direct effect. The direct effect of a change in migration costs is:

$$\frac{\partial W}{\partial m_i} = \underbrace{\lambda_{ri}(1 - F(\hat{z}_i))}_{\text{rural to urban migrants}} \times \underbrace{\int_{\hat{z}} Azf(z|Z \geq \hat{z}_i)dz}_{\text{avg. prod. of migrants}}, \tag{16}$$

which is how many people are currently migrating (as we are holding $\hat{z}_i$ fixed) multiplied by the average productivity of the migrants.

Indirect effect. The indirect effect is how a change in migration costs changes the composition of those migrants. This component is then

$$\frac{\partial W}{\partial \hat{z}_i} = \lambda_{ri}\left(-m_i A\hat{z}_i f(\hat{z}_i) + \psi f(\hat{z}_i)\right)|_{\hat{z}_i} = 0. \tag{17}$$

where the first term in the parenthesis can be thought of as the marginal social cost of moving the marginal worker into the urban area versus the marginal social benefit of keeping that marginal worker in the rural area. When this trade-off is evaluated at the equilibrium cutoff of $\hat{z}_i$, it says that the indirect effect is zero. In other words, a reduction in migration costs provides no welfare gain through the change in the composition migrants.

---

[3]As we discuss in the Appendix, with parametric assumptions one can combine differential returns in (14) with differences in migration rates across regions (similar to gravity models in the trade literature, or Bryan and Morten (2019)) to estimate the parameters of the model and separate out A and θ.

It may be surprising to some readers, as it was to some of the authors, that the gains from real-location are zero in this model. The reason the model delivers zero indirect effects is because the model is efficient. What we mean by efficiency here is that the decentralized allocation satisfying the definition of the competitive equilibrium (Definition 1) corresponds with the centralized allocation of workers that would arise from a planner maximizing social welfare, (6).[4] What this implies is the allocation at which migration costs are changed already optimally balances the marginal cost of adding one more worker to the urban area versus the marginal benefits of keeping that worker in the rural area. Thus, the social gain from reallocating a worker into (or out of) the urban area is zero on the margin.

Combing (15), (16) and (17) allows us to summarize the gains from lower migration costs:

Proposition 3 *The **Gains from Lower Migration Costs** in location i are given by*

$$\text{GLC}_i = \underbrace{\lambda_{ri}(1 - F(\hat{z}_i))}_{\text{rural to urban migrants}} \times \underbrace{\int_{\hat{z}} Azf(z|Z \geq \hat{z}_i)dz}_{\text{avg. prod. of migrants}}, \tag{18}$$

which, with Pareto distributed skills, is

$$\text{GLC}_i = \pi_{uri} \times \frac{\theta}{\theta - 1} \times \frac{\psi}{m_i} \tag{19}$$

where $\pi_{uri} = \lambda_{ri}(1 - F(\hat{z}_i))$ is mass of rural-urban migrants from region i.

There are several observations to be made about the gains from lowering migration costs. First, notice that the first term in (18) or (19) is the share of migrants from region i into the urban location. Holding all else equal, areas with more rural-urban migrations will have larger benefits from lowering costs of migrating.

A second observation is the role that the migration costs play. Conditional on the same amount of flows, areas with higher migration costs (lower m_i) will have larger gains associated with them. The observation further motivates the idea discussed around equation (14), that identifying high migration cost regions is useful approach to identifying high gains from migration situations.

Proposition 3 has implications for the importance of selection and the associated policy prescriptions. How selected migrants are also matters and the dispersion in talents is a key force in shaping the importance of selection. In the context of the Pareto distribution, this observation comes out very clearly. If θ is large and there is not much dispersion in talent, then the

---

[4]More specifically, one can cast the social planning problem as the planner choosing cut-offs, $\hat{z}_i$ to maximize (6). This optimization problem leads to the first-order necessary condition that $\frac{\partial W}{\partial \hat{z}_i} = 0$ (which is the same as in 17) and, hence, the planner's cut-off corresponds with the cut off in the competitive equilibrium.

gains from migration will be small (holding all else constant). If θ is small and there is a lot of dispersion in talents, then the gains from migration are larger.

This last observation is important in the context of interpreting the role of selection in the cross-sectional gap or the observational return and how it maps to the gains from lowering migration costs. Hicks et al. (2017) articulate the convention wisdom well when they write: "if observed gaps are partly driven by selection, the policies to incentivize workers to move to urban areas ...would be a less attractive policy direction." The results in Proposition 3 suggest the opposite: the more important is selection, the *larger* are the gains from migration.

All these observations illustrate the mapping from observational returns in Proposition 2 to the gains from lowering migration costs in Proposition 3 is not one-to-one. Observational returns in Proposition 2 confound the cross-sectional gap and the average net-benefits for migrants. The welfare gains in Proposition 3 reflect migration flows and migrants' average productivity in the urban area.

With that said, one approach could be to measure the relative gains across regions by combining relative observational returns and relative migration flows. Specifically, the relative gains are

$$\frac{\mathrm{GLC}_h}{\mathrm{GLC}_l} = \frac{\pi_{urh}}{\pi_{url}} \times \exp\left\{\mathrm{ORM}_h - \mathrm{ORM}_l\right\}. \tag{20}$$

This approach would identify relatively beneficial situations when comparing different interventions across different regions.

Before moving on, there are two important caveats regarding Proposition (3). While these formula build a lot of intuition, they are local effects and, hence, may not fully capture all the gains associated with larger attempts to influence patterns of internal migration. A second caveat is that we have nothing to say regarding the costs and fiscal burden associated with lowering internal migration costs. This is, of course, an important part of any policy decision-making, but well beyond the scope of the paper.

3.4. Alternative Perspectives

While our model is a relatively standard Roy (1951) model, we must admit that there are important alternative perspectives on what observational returns mean and reflect and how they relate to the gains from lower migration costs. Below we discuss the perspective provided by our discussant (Schoellman, 2020). A key issue is that the interpretation on the observational returns depends critically on why migrants migrate. At a minimum, this validates the point that observational returns are very difficult to interpret. However, the gains to migration (at least in the contexts we discuss) do not.

3.4.A. The Model of Schoellman (2020)

Our discussant shares our view that starting from a Roy (1951) model makes sense. The key difference is that in our model there are some rural workers who start out in life in a rural location, but eventually move to the urban location. Some workers started out in the wrong place, in other words. In contrast, our discussant wants to think about a situation where everybody is already in the correct location, but that movers move because, say, there is some growth in urban productivity.

This alternative assumption about initial conditions leads to a different interpretation of the observation returns. Because workers are essentially in the correct place in this view, anyone who is observed to move must have been near indifferent and hence have productivity approximately equal to the cutoff, $\hat{z}$. The observational return in this view is:

$$\text{ORM}_i = \log A + \log \hat{z}_i - \log \psi = -\log m_i \tag{21}$$

which reflects the direct cost of migrating. Where are the extra selection terms that show up our equation for observed returns in (10)? The key difference here is that in our discussant's world, all movers are marginal. In contrast, migrants in our model are not marginal. Thus, some workers gain above and beyond the moving cost and that gain reflects the talent distribution in the urban area.

Does this matter for welfare and the gains to lowering migration costs? Not quite. To see this, social welfare in our discussant's world would be:

$$\mathcal{W} = \sum_i \lambda_i \left\{ (1 - F(\hat{z}_i)) \left[\int_{\hat{z}_i} m_i A z f(z | Z \geq \hat{z}_i) dz \right] + F(\hat{z}_i) \psi \right\} \tag{22}$$

where i indexes the origin location. So for each location there are some urban people who have arrived recently, or far in the past, and the first term in the brackets reflects the urban incumbents from i and the last term reflects the people remaining in the rural location. Comparing (22) and (6) reveals that they would lead to the same first order conditions and, in turn, the same results described Proposition 3. There would be an interpretation difference in that the stock of migrants in the urban area that came from location i matters, not the flow of migrants as in our model. But the same idea is there—the gains are described by the amount of people from a location adjusted by the productivity of people coming in. Overall, our discussant's comments, and these observations, suggest a need for dynamic frameworks in future work to better understand the welfare gains from lowering migration costs.

4. Observational Returns to Migration by Region

An important feature of our model is that it emphasizes different costs of rural-urban migration across different communities. The feature led us to the following insights. While the observational returns overall are hard to interpret, studying how observational returns vary across regions in the data is beneficial. In particular, the theory says that variation in observational returns is informative about situations in which the returns for some individuals may be relatively large. The theory suggests that these situations reflect high migration costs, and our formula for the welfare gains suggests that the gains from lowering migration costs may be high in some areas, even if observational returns are modest on average. A related testable prediction of the model is that regions with lower rural-urban migration rates have higher observational returns to migration.

To test this prediction of the model, we examine the observed returns to migration for migrants from rural communities with high urban migration rates and low urban migration rates separately. For each enumeration area in our initial data, we calculate the share of individuals that begin in that community who migrate to an urban area. We then split our sample of rural-urban migrants in half, based on whether the out-migration rate in their community of origin is above or below the median community.

In our data, we find support for our model's prediction that regions with higher migration costs have higher observational returns to migration. Table 6 shows that in three of our six countries (China, Indonesia, and Tanzania), migrants from low-migration communities have significantly higher consumption gains from migration than those from high-migration communities. In the remaining three countries we cannot reject that the returns are equal across our two groups. In the case of South Africa, we can reject any large differences in the returns between low and high-migration communities. In contrast, our confidence intervals for Ghana and Malawi, by virtue of the small sample sizes in these countries, are largely uninformative. In no country is the observed gains to migration significantly higher in high-migration communities. When we aggregate our results across countries using a random-effects model, we find that the difference in returns between low and high migration communities averages 12 log points, and that this difference is significant at the 5-percent level.[5]

---

[5]In this calculation we weight each country by $1/[se^2 + t^2]$, where se is the standard error of the estimate in that country and t^2 is an estimate of the variance of the true estimates across countries. Note that while there is substantial heterogeneity in the returns to rural-urban migration across our samples, this is driven entirely by returns in high-migration communities. The returns to migration from low-migration communities are very similar across countries and we cannot reject that there are no differences.

5. Observational versus Experimental Returns to Migration

While our results from Table 6 suggest that the observational returns to migration may understate the returns to migration for 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 et al. (2014), a randomized controlled trial that experimentally lowered the cost of seasonal rural-to-urban migration by offering travel subsidies to rural households in northern Bangladesh.

The subsidy, valued at roughly $8.50 USD, provided enough money for a round-trip bus fare to nearby cities conditional on the household sending at least one migrant 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.

These data give us the rare opportunity to compare the observational returns to migration to the experimental returns accruing to those who were induced to move by reduced costs. We can generate a local average treatment effect estimate of the consumption gains for induced migrants by instrumenting for the household's migration decision with the travel subsidy treatment. In addition, due to the panel nature of the datasets that the authors collected, 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 7 shows two estimates of the return 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 two estimates is an economically meaningful and statistically significant 27 log points. In other words, the returns to migration for policy-induced migrants in this context are some 300 percent greater than the returns for those that change from being a migrant to a non-migrant in the status quo.

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

sample involves only temporary rather than permanent migration, during a specific agricultural lean season, and includes households from only two districts 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 stemming from an external shock that reduced migration costs.

6. Conclusion

Recent evidence shows that observational returns to rural-urban migration are not far above zero in three large developing countries. One may be tempted to conclude from this evidence that those who haven't migrated yet would also gain little from migrating, and hence that policymakers should give up on trying to reduce internal migration costs. We argue that this conclusion would be premature. Empirically, we draw on panel evidence from six developing countries and find higher observational returns on average than those reported in the existing literature. We then show theoretically that such estimates of the observational returns, using panel fixed effects regressions, are not that informative about the gains from lowering migration costs in an economy with heterogeneous benefits and costs of migration. The reason is that lowering migration costs can help those with high potential benefits from migration, but who happen to face even higher costs. Complementary experimental evidence from Bangladesh highlights the fact that many rural households do indeed face high costs, and they gain a lot when those costs are reduced by a policy experiment. Moreover, running the panel fixed effects regression in that same sample heavily understates the "true" returns that can be extracted from the experimental design.

Our paper points to the need for more research to help identify the joint distribution of benefits and costs of rural-urban migration in developing countries. To inform migration policy, we need to develop a more refined understanding of the potential benefits of migration for those not already migrating, and the precise reasons why they are not migrating. Experiments can help, particularly in identifying the specific barriers that hold back internal migration. For example, in the field experiment of Baseler (2019), an information treatment given to a random group of rural Kenyan villagers increased the fraction of rural households sending a migrant to Nairobi from 20 percent to 28 percent. Information frictions appear to be a type of migration cost that is holding back migration in that setting. Using the construction of a network of new highways to connect the new capital of Brasilia, Brazil, as a natural experiment, Morten and Oliveira (2019) find that roads promoted internal migration in Brazil. Asher and Novosad (2020) also find changes in mobility and occupational choice in rural India associated with last-mile road connections to villages, identified by applying a regression discontinuity design to data on a massive village road building program. The lack of good transportation infrastructure is, in short, not just a barrier to trade in goods but to movements of workers internally. Future

work should explore other costs of internal migration 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. (2020): "The Agricultural Wage Gap: Evidence from Brazilian Micro-data," *American Economic Journal: Macroeconomics*, 12, 153–73.

ASHER, S. AND P. NOVOSAD (2020): "Rural Roads and Local Economic Development," *American Economic Review*, 110, 797–823.

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, 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, 170–199.

BRUECKNER, J. K. AND S. 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*, ed. by G. Duranton, J. V. Henderson, and W. Strange, vol. 5B, 1399–1451.

BRYAN, G., S. CHOWDHURY, AND A. M. MOBARAK (2014): "Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh," *Econometrica*, 82, 1671–1748.

BRYAN, G. AND M. MORTEN (2019): "The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia," *Journal of Political Economy*, 127, 2229–2268.

DERCON, S., P. KRISHNAN, AND S. KRUTIKOVA (2013): "Changing Living Standards in Southern Indian Villages 1975?2006: Revisiting the ICRISAT Village Level Studies," *Journal of Development Studies*, 49, 1676–1693.

FRENCH, E. AND C. TABER (2011): "Identification of Models of the Labor Market," in *Handbook of Labor Economics*, ed. by O. Ashenfelter and D. Card, Elsevier, vol. 4A, chap. 6, 537–617.

GARLICK, J., M. LEIBBRANDT, AND J. LEVINSOHN (2016): "Individual Migration and Household Incomes," NBER Working Paper 22326.

GLAESER, E. L. AND D. C. MARÉ (2001): "Cities and Skills," *Journal of Labor Economics*, 19, 316–342.

GOLLIN, D., D. LAGAKOS, AND M. E. WAUGH (2014): "The Agricultural Productivity Gap," *Quarterly Journal of Economics*, 129, 939–993.

HECKMAN, J. J. AND B. HONORÉ (1990): "The Empirical Content of the Roy Model," *Econometrica*, 58, 1121–1149.

HECKMAN, J. J. AND E. VYTLACIL (2005): "Structural Equations, Treatment Effects, and Econometric Policy Evaluation," *Econometrica*, 73, 699–738.

HERRENDORF, B. AND T. SCHOELLMAN (2018): "Wages, Human Capital, and Barriers to Structural Transformation," *American Economic Journal: Macroeconomics*, 10, 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.

INSTITUTE OF SOCIAL SCIENCE SURVEY (2015): "China Family Panel Studies," Peking University Open Research Data Platform, V37.

KHANNA, G., W. LIANG, A. M. MOBARAK, AND R. SONG (2019): "The Productivity Consequences of Pollution-Induced Migration in China," Unpublished Working Paper, Yale University.

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, 948–980.

LUCAS, R. E. B. (1997): "Internal Migration in Developing Countries," in *Handbook of Population and Family Economics*, ed. by M. R. Rosenzweig and O. Stark, vol. 1B, 721–798.

MEGHIR, C., A. M. MOBARAK, C. MOMMAERTS, AND M. MORTEN (2019): "Migration and Informal Insurance," NBER Working Paper 26082.

MORTEN, M. (2019): "Temporary Migration and Endogenous Risk Sharing in Village India," *Journal of Political Economy*, 127, 1–46.

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.

MUNSHI, K. AND M. ROSENZWEIG (2016): "Networks and Misallocation: Insurance, Migration, and the Rural-Urban Wage Gap," *American Economic Review*, 106, 46–98.

NAKAMURA, E., J. SIGURDSSON, 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, 645–652.

ROY, A. (1951): "Some Thoughts on the Distribution of Earnings," *Oxford Economic Papers*, 3, 135–46.

SARVIMÄKI, M., R. UUSITALO, AND M. JÄNTTI (2019): "Habit Formation and the Misallocation of Labor: Evidence from Forced Migrations," Unpublished Working Paper, Aalto University School of Business.

SCHOELLMAN, T. (Forthcoming): "Comment on "Migration Costs and Observational Returns to Migration in the Developing World"," *Journal of Monetary Economics*.

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

Table 1: Comparison of Consumption Aggregates

Variable	CHN	GHA	IDN	MWI	ZAF	TZA
Number of Food Categories	2	85	37	136	32	72
Number of Non-Food Categories	29	53	23	120	54	46
Own-Produced Food	Yes	Yes	Yes	Yes	Yes	Yes
Durable Goods	No	No	No	Yes	No	No
Rent/Mortgage	Yes	No	Yes	Yes	Yes	No
Rental Value of Owner-Occupied Housing	No	No	Yes	Yes	Yes	No
Spatially-Deflated	No	Yes	No	Yes	No	Yes

Note: This table summarizes the features of the consumption aggregates in China (CHN), Ghana (GHN), Indonesia (IDN), Malawi (MWI), South Africa (ZAF) and Tanzania (TZA).

Table 2: 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 list the number of survey waves, the number of communities (i.e. enumeration areas) surveyed, the number of individuals surveyed, and the total number of observations, for each country. Column 5 lists the fraction of adults living in a rural location in wave 1. Column 6 presents the annualized rural-urban migration rate for adults in wave 1.

Table 3: Observational Returns to Migration in Six Developing Countries

	(1)	(2)	(3)	(4)
China	0.545***	0.161***	0.012	0.226***
	(0.005)	(0.028)	(0.064)	(0.031)
Ghana	0.410***	0.148	-0.173	0.339**
	(0.013)	(0.122)	(0.220)	(0.148)
Indonesia	0.625***	0.145***	0.039	0.167***
	(0.009)	(0.019)	(0.031)	(0.029)
Malawi	0.520***	0.048	-0.350***	0.189
	(0.012)	(0.089)	(0.123)	(0.134)
South Africa	0.737***	0.212***	0.028	0.291***
	(0.006)	(0.022)	(0.044)	(0.026)
Tanzania	0.666***	0.112***	0.101**	0.213***
	(0.032)	(0.030)	(0.045)	(0.043)
Individual FE	No	Yes	Yes	Yes
Year FE	No	Yes	Yes	Yes
Sample	Full	Full	Start Urban	Start Rural

Note: This table presents the estimated coefficients of urban dummy variables from regressions of log consumption per adult on urban dummies and other covariates in the six countries. Column (1) presents the cross-sectional estimates, with no other controls. Column (2) adds year and individual fixed effects, plus quadratic controls for age and household size. Column (3) has year and individual fixed effects, plus quadratic controls for age and household size, and restricts the sample to only those starting in an urban location. Column (4) is the same model as in column (3), but restricts the sample to only those starting from a rural location. Robust standard errors, clustered at the level of the wave 1 household, are in parenthesis. $^*p < .1,^{**}p < .05,^{***}p < .01$

Table 4: Observational Returns to Migration: Income Measures

	(1)	(2)	(3)	(4)
Indonesia	0.541***	0.112***	-0.004	0.148***
	(0.013)	(0.023)	(0.035)	(0.037)
South Africa	0.649***	0.249***	0.035	0.289***
	(0.006)	(0.023)	(0.045)	(0.029)
Individual FE	No	Yes	Yes	Yes
Year FE	No	Yes	Yes	Yes
Sample	Full	Full	Start Urban	Start Rural

Note: This table presents the estimated coefficients of urban dummy variables from regressions of log household income per adult on urban dummies and other covariates in the six countries. Column (1) presents the cross-sectional estimates, with no other controls. Column (2) adds year and individual fixed effects, plus quadratic controls for age and household size. Column (3) has year and individual fixed effects, plus quadratic controls for age and household size, and restricts the sample to only those starting in an urban location. Column (4) is the same model as in column (3), but restricts the sample to only those starting from a rural location. Robust standard errors, clustered at the level of the wave 1 household, are in parenthesis. $^{*}p < .1,^{**}p < .05,^{***}p < .01$

Table 5: Results for Indonesia Under Alternative Specifications

	(1) Cross-Section	(2) Panel	(3) Rural-Only	(4) Work Moves
City, Ind. Earnings	0.638***	0.101***	0.081	0.016
	(0.019)	(0.038)	(0.066)	(0.134)
City, HH. Earnings	0.541***	0.112***	0.148***	0.272***
	(0.013)	(0.037)	(0.037)	(0.076)
City, HH. Consumption	0.625***	0.145***	0.167***	0.263***
	(0.009)	(0.019)	(0.029)	(0.056)
City + small towns, Ind. Earnings	0.739***	0.065**	0.044	0.071
	(0.014)	(0.030)	(0.034)	(0.073)
City + small towns, HH. Earnings	0.393***	0.094***	0.086***	0.142***
	(0.009)	(0.018)	(0.020)	(0.048)
City + small towns, HH. Consumption	0.597***	0.125***	0.110***	0.156***
	(0.006)	(0.014)	(0.015)	(0.036)
Individual FE	No	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Initial Location	All	All	Rural	Rural
Move Reason	Any	Any	Any	Work

Note: This table presents the estimated coefficients of urban dummy variables from regressions of welfare measures on urban dummies and other covariates for Indonesia. Column (1) presents the cross-sectional estimates, with no other controls. Column (2) adds year and individual fixed effects, plus quadratic controls for age and household size. Column (3) has year and individual fixed effects, plus quadratic controls for age and household size, and restricts the sample to only those starting in an rural location. Column (4) estimates the same model as column (3) and restricts the sample to only those starting from a rural location and also excludes migrant individuals who report the reason for their move being a non-work reason. Robust standard errors, clustered at the level of the wave 1 household, are in parenthesis. The rows vary in their definition of urban areas and in the dependent variables. The urban definition in rows 1-3 includes only locations coded as cities. The urban definition in rows 4-6 includes both towns and cities, as in Hicks et al. (2017). The dependent variable in rows 1 and 4 is individual earnings, as in Hicks et al. (2017). The dependent variable in rows 2 and 5 is household income per adult. The dependent variable in rows 3 and 6 is household consumption per adult. $^*p < .1,^{**}p < .05,^{***}p < .01$

Table 6: Observational Returns to Migration by Region

	(1) High Migration	(2) Low Migration	(3) p-value: (1) = (2)
China	0.108**	0.233***	0.043
	(0.046)	(0.043)	
Ghana	0.240	0.428*	0.507
	(0.174)	(0.228)	
Indonesia	0.074*	0.260***	0.001
	(0.040)	(0.041)	
Malawi	0.287	0.121	0.532
	(0.199)	(0.179)	
South Africa	0.303***	0.277***	0.616
	(0.039)	(0.037)	
Tanzania	0.071	0.357***	0.000
	(0.052)	(0.061)	
Individual FE	Yes	Yes	
Year FE	Yes	Yes	
Sample	High Migration	Low Migration	

Note: This table presents the estimated coefficients of urban dummy variables from regressions of log consumption per adult on urban dummies. Specifications are as in Table 3, Column (4), with year and individual fixed effects, plus quadratic controls in age and household size, and restricting the sample to only those starting from a rural location. The sample is divided by the rural-urban migration rate in the origin community, so that there are an equal number of rural-urban migrants in each group. Column (2) restricts the sample to include households from enumeration areas with rural-urban migration rates above the median rate for rural-urban migrants. Column (2) restricts the sample to include households from enumeration areas with rural-urban migration rates below the median rate for rural-urban migrants. Column (3) reports the p-value of the difference between the estimates in Column (1) and Column (2). Robust standard errors, clustered at the level of the wave 1 household, are in parenthesis. $^*p < .1,^{**}p < .05,^{***}p < .01$

Table 7: Observational vs. Experimental Returns to Seasonal Migration

	(1) Observational	(2) Experimental	(3) Difference
Seasonally Migrated	0.092*	0.357**	
	(0.053)	(0.156)	
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, Chowdhury, 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 assignment 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 1,000 bootstrap replications. $^*p < .1,^{**}p < .05,^{***}p < .01$

Appendix (for Online Publication)

A. Data Appendix: Spatial Price Deflation

Malawi. The Malawi Integrated Household Panel Survey (IHPS) created a spatial and temporal price index that is used to deflate consumption aggregates. Spatial deflation and within-year monthly deflation is based on the official monthly Malawi CPI during the duration of fieldwork across four regions: Urban, Rural North, Rural Centre, and Rural South. Urban prices were consistently higher than those in all other regions. Due to a rebasing of the Malawi CPI midway through the panel the official CPI could not be used to adjust for price differences across survey rounds. Instead, deflation across rounds of the survey rely on a Laspeyres index constructed using food price data from IHPS household surveys and non-food prices from IHPS community surveys. The bundle of items used to construct this index covered 58 percent of consumption in the sample.

Tanzania. The Tanzania National Panel Survey (NPS) price index adjusts for spatial and temporal price differences. The index relies on the prices of food items recorded in NPS household surveys. These differences in food prices were used to deflate both food and non-food consumption, with the assumption that trends in non-food prices were similar. The NPS created Fisher price indices using the median unit price of each food item. The index adjusts for spatial price differences across four strata: rural, Zanzibar, Dar es Salaam and other urban areas. The index adjusts for temporal price differences at the monthly level. It includes all items that had been purchased by at least 10 households by stratum and quarter, resulting in a bundle of items which covered 67 percent of total food consumption.

Ghana. We built a price index for Ghana that deflates prices for urban relative to rural areas. The index relies on the prices of food items recorded in our survey data. We use this index to deflate both food and non-food consumption, with the assumption that trends in food and non-food prices were similar across space. Food consumption was reported using a diverse mixture of standard and non-standard units, with 73 unique units used. We therefore include only unit-prices for units-item combinations with at least 10 observations in both rural and urban locations (which covers 98 percent of food consumption in both regions. We treat each item-unit pair as a separate good in our index, and calculate a Fischer price index for rural consumption with urban locations as the base.

B. Model Appendix

This appendix derives the results in Section 3.3. We begin with some mathematical preliminaries regarding the Pareto distribution that we use in our examples.

- The Pareto CDF is $F(z) = \text{Prob}\{Z \leq z\} = 1 - T^\theta z^{-\theta}$ and the associated PDF is $\theta T^\theta z^{-\theta-1}$. The expected value is $\frac{\theta}{\theta-1}T$. Throughout the examples we normalize the parameter $T = 1$.

- The expected value, conditional on being at or above the value $\hat{x}$, is $\frac{\theta}{\theta-1}\hat{x}$.

- The logarithm of $\frac{z}{\hat{x}}$ is distributed exponentially with parameter θ. This implies that the expected value of $\log\left(\frac{z}{\hat{x}}\right)$ is $\frac{1}{\theta}$.

The **Cross-Sectional Gap** is simply the average of consumption in the urban area minus the average of consumption in the rural area. We focus on period one for simplicity and with

$$\text{CSG} = \underbrace{\log A + \int \log z f(z)dz + \log \psi}_{\text{urban consumption}} - \left(\underbrace{\log 1 + \log \psi}_{\text{rural consumption}}\right), \tag{23}$$

$$\text{CSG} = \log A + \int \log z f(z)dz$$

In the case of Pareto distributed talents, notice that

$$= \log A + \log T - \log T + \int \log z f(z)dz. \tag{24}$$

and again as:

$$= \log A + \log T + \int \log\left(\frac{z}{T}\right)f(z)dz. \tag{25}$$

Then, applying the properties of the exponential distribution, we have:

$$= \log A + \log T + \frac{1}{\theta}. \tag{26}$$

Then with $T = 1$ (our normalization) the result in equation (8) follows.

The **Observational Return** is given by

$$\text{ORM} = \log A + \sum_i \frac{\pi_{uri}}{\pi_{ur}}\left(\int_{\hat{z}_i} \log z f(z|Z \geq \hat{z}_i)dz - \log \psi\right), \tag{27}$$

where $\pi_{uri} = \lambda_{ri}(1 - F(\hat{z}_i))$ is mass of rural to urban migrants from rural region i and $\pi_{ur} = \sum_i \pi_{uri}$ is the total mass of migrants. To connect this with the cross-sectional gap, do the fol-

lowing: add and subtract the unconditional distribution of talent

$$\text{ORM} = \log A + \int \log z f(z)dz - \int \log z f(z)dz + \sum_i \frac{\pi_{uri}}{\pi_{ur}} \left(\int_{\hat{z}_i} \log z f(z|Z \geq \hat{z}_i)dz - \log \psi \right),$$
(28)

and then pulling the minus term into the brackets by noting that the weights $\frac{\pi_{uri}}{\pi_{ur}}$ sum to one gives us

$$= \text{CSG} + \sum_i \frac{\pi_{uri}}{\pi_{ur}} \underbrace{\left\{ \left(\int_{\hat{z}_i} \log z f(z|Z \geq \hat{z}_i)dz - \int \log z f(z)dz \right) - \log \psi \right\}}_{\text{skill adjustment of migrants}}$$
(29)

which is (11) in the body of the paper. For Pareto distributed talents, first focus on the $\int_{\hat{z}_i} \log z f(z|Z \geq \hat{z}_i)dz$ term

$$\int_{\hat{z}_i} \log z f(z|Z \geq \hat{z}_i)dz = \log \hat{z}_i + \int_{\hat{z}_i} \frac{\log z}{\log \hat{z}_i} f(z|Z \geq \hat{z}_i)dz$$
(30)

$$= \log \hat{z}_i + \frac{1}{\theta}$$
(31)

then plugging the expression above into the ORM formula (27) we have

$$\text{ORM} = \log A + \sum_i \frac{\pi_{uri}}{\pi_{ur}} \left(\log \hat{z}_i + \frac{1}{\theta} - \log \psi \right).$$
(32)

Before further canceling terms, notice that if everybody is migrating, so $\hat{z}_i = 1$, then the observational return just reflects the cross-sectional gap $\log A + \frac{1}{\theta}$ net of skill loss $-\log \psi$. Assuming that we are away from this case, substitute in for $\hat{z}_i$ and we have

$$\text{ORM} = \log A + \sum_i \frac{\pi_{uri}}{\pi_{ur}} \left(\log \left(\frac{\psi}{mA} \right) + \frac{1}{\theta} - \log \psi \right).$$
(33)

After canceling terms we get

$$\text{ORM} = \frac{1}{\theta} - \sum_i \frac{\pi_{uri}}{\pi_{ur}} \log m_i$$
(34)

which is equation (12) in the body of the paper.

Several observations about the observational return.

- First, if we compare observational returns across two regions (say a high migration and

low migration region) we have

$$\text{ORM}_h - \text{ORM}_l = \int_{\hat{z}_h} \log z f(z|Z \geq \hat{z}_h) dz - \int_{\hat{z}_l} \log z f(z|Z \geq \hat{z}_l) dz \tag{35}$$

which the relative observational return reveals the relative productivity of those migrating from one region to another (note that this strongly relies upon common skill loss). The standard intuition about this is that the conditional expectation is increasing in the conditioning point. This intuition implies that workers from the high migration cost are relatively more productive than workers from the low migration cost region. The Pareto distributed talents case makes a stronger prediction illustrating that

$$\text{ORM}_h - \text{ORM}_l = [\log m_l - \log m_h] > 0 \tag{36}$$

high migration cost regions are where the migrants are of higher productivity and the relative return exactly reveals the relative migration costs.

- A second observation is that under some assumptions, the treatment effect A can be estimated. Under the assumption of Pareto and if one assumes that $\sum_i \frac{\pi_{uri}}{\pi_{ur}} \log m_i \approx 0$, then the observational return

$$\text{ORM} \approx \frac{1}{\theta} \tag{37}$$

reveals $\frac{1}{\theta}$ which is the selection part of the cross-sectional gap in (8). In other words, under this assumption

$$\text{CSG} - \text{ORM} \approx A \tag{38}$$

so the difference between the cross-sectional gap and the observational return tells us about A or the treatment effect of being in the urban area.

- A third observation is a relationship between relative flows and relative returns can be used to estimate the importance of selection θ. Under the assumption of Pareto distributed talents, notice that

$$\pi_{uri} = \lambda_{ri} \left(\frac{\psi}{Am_i} \right)^{-\theta} \tag{39}$$

which after taking logs and rearranging gives

$$\log \frac{\pi_{uri}}{\lambda_{ri}} = -\theta \left(\log \psi - \log A - \log m_i \right) \tag{40}$$

where $\frac{\pi_{uri}}{\lambda_{ri}}$ is the migration rate out of region i. Take differences across the high migration cost region and the low migration cost region

$$\log \frac{\pi_{urh}}{\lambda_{rh}} - \log \frac{\pi_{url}}{\lambda_{rl}} = \theta \log m_h - \theta \log m_i \tag{41}$$

$$= -\theta \left(\log m_i - \log m_h \right) \tag{42}$$

where common terms such as the A and the ψ cancel, then comparing (41) with the relative observational return under Pareto distributed talents reveals

$$\log \frac{\pi_{urh}}{\lambda_{rh}} - \log \frac{\pi_{url}}{\lambda_{rl}} = \theta \left(\mathrm{ORM}_h - \mathrm{ORM}_l \right) \tag{43}$$

which says that by connecting the flows and the relative returns, this is another approach to reveling θ and in turn A or the treatment effect from urban area. Or anther way to put this is that an estimate of A is

$$\mathrm{CSG} - \frac{(\mathrm{ORM}_h - \mathrm{ORM}_l)}{\log \frac{\pi_{urh}}{\lambda_{rh}} - \log \frac{\pi_{url}}{\lambda_{rl}}} = A. \tag{44}$$

The **Gains from Lower Migration Costs**. To compute this start with our social welfare function

$$\mathcal{W}(\mathbf{m}, \hat{z}(\mathbf{m})) = \lambda_u \int \psi A z f(z) dz + \sum_i \lambda_{ri} \left\{ (1 - F(\hat{z}_i)) \left[\int_{\hat{z}_i} m_i A z f(z | Z \geq \hat{z}_i) dz \right] + \psi F(\hat{z}_i) \right\}, \tag{45}$$

and then compute this in two steps. First, we will compute the direct effect, so a change in m holding fixed $\hat{z}$

$$\frac{\partial W}{\partial m_i} = \underbrace{\lambda_{ri}(1 - F(\hat{z}_i))}_{\text{rural to urban migrants}} \times \underbrace{\int_{\hat{z}} A z f(z | Z \geq \hat{z}_i) dz}_{\text{avg. prod. of migrants}}, \tag{46}$$

which just follows from straightforward differentiation. The next step is to compute the indirect effect. That is the change in social welfare through how a change in m changes $\hat{z}$. Before proceeding, note that

$$(1 - F(\hat{z}_i)) \int_{\hat{z}_i} m_i A z f(z | Z \geq \hat{z}_i) dz = \int_{\hat{z}_i} m_i A z f(z) dz \tag{47}$$

and using this substitution and computing the requisite derivative gives

$$\frac{\partial W}{\partial \hat{z}_i} = \lambda_{ri} \left(\frac{\partial}{\partial \hat{z}_i} \int_{\hat{z}_i} m_i A z f(z) dz + \psi f(\hat{z}_i) \right).$$ (48)

and then partial derivative inside the parenthesis is

$$\frac{\partial}{\partial \hat{z}_i} \int_{\hat{z}_i} m_i A z f(z) dz = -m_i A z f(z)$$ (49)

which follows from fundamental theorem of calculus. Then combining this we have

$$\frac{\partial W}{\partial \hat{z}_i} = \lambda_{ri} \left(-m_i A \hat{z}_i f(\hat{z}_i) + \psi f(\hat{z}_i) \right)$$ (50)

which equals zero when evaluated at $\hat{z}_i$. Thus, only the direct effect matters for the gains to lowering migration costs. That is the

$$\text{GLC}_i = \underbrace{\lambda_{ri}(1 - F(\hat{z}_i))}_{\text{rural to urban migrants}} \times \underbrace{\int_{\hat{z}} A z f(z | Z \geq \hat{z}_i) dz}_{\text{avg. prod. of migrants}}.$$ (51)

Finally the direct effect with the Pareto distribution then simply follows from the mathematical preliminaties. Applying these formula leaves

$$= \pi_{uri} \times \frac{\theta}{\theta - 1} \times \frac{\psi}{m_i}.$$ (52)

Some observations on connecting this with the observational return or the cross-sectional gap. Under Pareto distributed talents (and the common ψ) one can identify relative gain situations:

$$\frac{\text{GLC}_h}{\text{GLC}_l} = \frac{\pi_{urh}}{\pi_{url}} \times \frac{m_l}{m_h}$$ (53)

and then the important point is the following: the ratio of the migration costs just equals the difference in observational returns from those discussed above. So

$$\frac{\text{GLC}_h}{\text{GLC}_l} = \frac{\pi_{urh}}{\pi_{url}} \times \exp\left(\log m_i - \log m_h \right)$$ (54)

$$= \frac{\pi_{urh}}{\pi_{url}} \times \exp\left(\text{ORM}_h - \text{ORM}_l \right).$$ (55)