When Are There Returns to Rural-Urban Migration? Reconciling Experimental & Observational Data

Jacob Tasto *

Draft: March 25, 2026

Abstract

Experimental studies in the development literature consistently find positive causal returns to inducing rural-urban migration, but observational studies of self-selected migration find almost no average change in a worker's wage—how do we reconcile these conflicting findings? I revisit individual-level panel data in Kenya to document heterogeneity in the returns to self-selected rural-urban migration: returns are much larger for workers with persistent negative origin earnings premiums. I argue that the LATE identified in RCTs is properly upper-bounded by the observational returns for this class of workers. I then explore why returns are heterogeneous by origin premium. I find evidence that most of the heterogeneity is due to wage compression, specifically resulting from workers with high rural premiums *losing* their premium when they migrate to an urban labor market, so that rural-urban migration has equalizing effects on the wage-gap of rural-born workers.

1 Introduction

The rural-urban wage gap in developing countries is well-documented (Lagakos, 2020). A body of research has emerged to study this gap: why does it exist, does it reflect causal returns or selection (Alvarez, 2020; Hamory et al., 2020*a*; Lagakos et al., 2020), and if it reflects real returns, what are the mechanisms preventing higher migration rates from closing the gap (Bryan, Chowdhury and Mobarak, 2014; Baseler, 2023; Meghir et al., 2022; Blumenstock, Chi and Tan, 2025; Akram, Chowdhury and Mobarak, 2025; Miner, 2026). This literature has covered a wide range of settings, including India, Bangladesh, Kenya, Indonesia, and Brazil.

But out of this work, a puzzle emerged. RCTs that experimentally induce rural-to-urban migration (both permanent and seasonal) consistently find positive causal returns to migration (Bryan, Chowdhury and Mobarak, 2014; Bryan et al., 2023; Baseler, 2023). But, event studies of

*Tasto: Department of Economics University of Washington (jtasto@uw.edu). This paper has benefited from helpful comments from Rachel Heath, Jason Kerwin, Tynan Challenor, and Alison Campion. All errors are my own.

self-selected movers with individual fixed effects in both Kenya and Indonesia find no average change in a worker's wage following the rural-urban event (Hamory et al., 2020*a*). Said another way: there are returns to rural-urban migration among workers who *do not* self-select into migration (call them "non-movers"), but *no* returns among workers who *do* select into migration. This runs counter to a classic selection model—where those with higher returns are the ones that select into treatment—and raises a puzzling question. Why are there returns to rural-urban migration for non-movers but no returns for movers?

This paper makes two contributions towards explaining the puzzle. First, I revisit the observational data in Kenya and show, using a TWFE specification with individual and time fixed effects, that despite there being no average wage return, there is a positive return for monthly earnings. This finding aligns with other observational estimates of the consumption returns to rural-urban migration (Lagakos et al., 2020) and with the experimental literature, which typically studies the effects of induced migration on earnings and consumption rather than wages. Second, I document important heterogeneity in the returns to rural-urban migration based on a worker's position in their origin's conditional earnings distribution. I find that the returns to rural-urban migration — for both earnings and wages — are significantly higher for workers who persistently earn less than their rural peers (those in the origin of the same age and gender with the same level of education). In contrast, the returns are minimal for workers without a negative premium. I argue that the observational returns for migrants with negative origin premiums provide an upper-bound for the experimental returns, as expected under selection on outcomes. The exact experimental estimate falls between zero and this upper-bound based on the extent to which the randomized intervention induces migration among rural workers with high versus low earnings. Hence, the LATE identified in RCTs can be larger than the average return for self-selected migrants writ large, but never exceeds the returns for workers with negative origin earnings premiums. In sum, we can reconcile the RCT and observational findings with a simple explanation: the *average* return to rural-urban migration is minimal but the return for low-earning workers is large.

In the second part of this paper, I attempt to answer the natural follow-up question: *why* are the returns to rural-urban migration greater for workers with negative premiums? First, I find that mean reversion cannot explain the magnitude of the returns heterogeneity. If a low wage in the rural labor market reflects a bad idiosyncratic draw, rather than low productivity, then migra-

tion could generate larger returns for low-earnings workers because they started from an unlucky position. To approximate the scope of mean reversion in generating heterogeneous returns to migration, I present estimates of the returns to lateral migration — rural-to-rural and urban-to-urban — and demonstrate that the excess returns for workers with negative premiums is much less significant than what is observed in rural-urban migration. I interpret this as evidence that another distinct mechanism, not mean reversion, is the primary source of the returns heterogeneity in rural-urban migration. Second, I provide evidence for wage compression as an explanation. If the wage distribution is tighter in urban labor markets than rural ones, then even under constant effects on relative productivity — i.e., a 20th percentile worker in the rural labor market is a 20th percentile worker in the urban one — we would still expect low-earning workers to gain more from urban migration, simply because the distance to the mean is smaller in urban markets than rural ones. I run quantile regressions for the rural-urban earnings and wage gaps and demonstrate that the quantile effects are decreasing in quintiles. This indicates that among rural-born workers, the urban earnings distribution is tighter than the rural one and could subsequently generate returns heterogeneity without any heterogeneity in the effect of migration on productivity. However, the magnitude of wage compression among rural-urban migrants is not large enough to explain the magnitude of returns heterogeneity, implying rank is not perfectly preserved across a rural-urban migration event.

Next, I attempt to distinguish whether the "excess returns" for low-earning workers are in part due to these workers resolving a spatial skill mismatch. If workers have negative rural earnings premiums because they are skill-mismatched and resolve this spatial skill-mismatch by migrating to an urban labor market, then rural-urban migration should have a positive effect on their earnings premium. However, I find that rural-urban migration has essentially no long-term effect on the earnings premium of low-earning workers. Instead, the bulk of the returns heterogeneity seems to come from *high productivity* rural workers immediately and permanently *losing* their earnings premium when they migrate to an urban labor market. In effect, returns are heterogeneous because rural-born workers are very heterogeneous in rural productivity but relatively homogeneous in urban productivity. And importantly, this homogeneity concentrates around the lower end of the urban distribution. The resulting outcome is that urban migration has an equalizing effect on the wages of migrating workers, but the equalization stems from reductions at the

top.

I then study migrants who select into "rural return." I find first that low-earning migrants exhibit steeper rates of rural return in the first 6-10 months after arriving in the urban destination. I also find that migrants opting into rural return experience sharp declines in their earnings and wages relative to their urban levels, but that there is minimal heterogeneity in this decline by the original rural earnings premium. And, I find that the returns to urban migration are smaller for workers who ultimately return to a rural labor market than for those who stay urban. I interpret these results as evidence of (1) selection into rural return based on lower gains to urban employment, and (2) some permanency to the urban productivity effect: urban migration shrank the wage gap between high and low rural earners, and this effect persisted when these workers moved back to the rural location.

Assuming, as this paper argues, that the returns to rural-urban migration are greater for low-earning workers, I then discuss the barriers that prevent higher rates of migration for this type of rural worker. I focus on two observations: (1) there is a distinct upward trend prior to migration in the earnings and wages of migrants with negative premiums that is not present for high-earning workers, and (2) in the first urban year, migrants with negative premiums return to a rural location at a higher rate. I argue that both of these observations are consistent with binding credit constraints and insufficient buffer savings preventing migration for low-earning workers. A first-order difference between workers with and without a negative premium is their difference in income and (presumably) savings. So, the responsiveness of low-earners' migration decision to a rise in income suggests these workers were, if not constrained by their savings, at least more hesitant to migrate until their income grew. Similarly, the higher rural return rate is consistent with low-earning migrants having smaller buffer savings at the time of their urban move and consequently having shorter runways to insure against urban unemployment and job search.

The paper proceeds as follows. In Section 2, I detail the rural-urban migration puzzle that this paper sets out to resolve. In Section 3, I define the setting and introduce the data used. Section 4 presents evidence to reconcile the seemingly conflicting results in the experimental and observational literature. In Section 5, I explore the source of the heterogeneity in returns. Section 6 discusses how these results contribute to the literature on the barriers to rural-urban migration. Section 7 concludes.

2 The Puzzle

Across countries, the gap between rural and urban wages is wide (Lagakos, 2020; Gollin, Lagakos and Waugh, 2014). Motivated by this gap, a body of experimental research in development economics has sought to study why higher rates of migration do not arbitrage the divide away. These studies (a selection of which are summarized in Table 1) have found evidence for the role of information frictions, borrowing constraints, lack of insurance, and lack of destination networks. This literature also generally finds that inducing rural-urban migration generates large causal returns for "non-movers," i.e. those that do not self-select into migration without an incentive or other intervention.

TABLE 1: SUMMARY OF MIGRATION RCTS

Study	Setting	Intervention	Migration Effect	Returns
Bryan, Chowdhury and Mobarak (2014)	Bangladesh	$11.50 conditional migration incentive	8-10pp	Consumption: 30-35%
		Information	None	None
Baseler (2023)	Kenya	Information	40%	Income: 24%, Food Expend: 9%
Baseler et al. (2025)	Kenya	Connection to mentor	12%	Income: 7-9%
		Group meeting	12%	None
Miner (2026)	Kenya	Unemployment insurance	200%	Wage Earnings: 22%
		Cash Transfer	100%	None

Notes:

A separate strand of literature has attempted to parse out whether the rural-urban wage gap reflects causal returns or sorting. Early cross-sectional work noted a rural-urban gap in educational attainment (Gollin, Lagakos and Waugh, 2014), and the tendency for rural-urban migrants to be more educated than those who remain in rural areas (Young, 2013), both of which were interpreted as evidence of sorting. More convincingly, studies leveraging individual-level panel data in Kenya, Indonesia, and Brazil found that including individual fixed effects reduced the estimate of the observational return to rural-urban migration (or a move out of agriculture) to near zero

(Hamory et al., 2020*a*; Alvarez, 2020). Figure 1 presents the event study results from Hamory et al. (2020*a*) that appeared to be particularly decisive evidence against causal returns to rural-urban migration. Self-selected migrants in both Indonesia and Kenya experienced no average increase in their wage following an urban move.

The findings from these two strands presented a puzzle: why are there returns to rural-urban migration for non-movers but no returns for movers? Classic selection models would posit that individuals select into migration based on their expected returns, which should create *larger* returns for self-selected movers than induced movers. The presence of binding constraints can of course distort the selection mechanism away from a perfectly positive relationship between returns and migration. But that still leaves the question of why workers with the most to gain would also be the most likely to face a binding constraint, and why individuals with near zero nominal returns would select into migration at all, given the higher cost of living in urban locations.

Lagakos et al. (2020) complicated the narrative of zero observational returns to rural-urban migration. They also studied individual-level panel data but found positive average returns of 17-34 log points on *consumption measures* in China, Ghana, Malawi, South Africa, Tanzania, and Indonesia. This suggests the zero observational returns in Hamory et al. (2020*a*) may in part stem from a focus on log wages rather than log consumption — where consumption is arguably a more relevant measure of welfare returns, and is more comparable to the variables used in estimating the experimental returns — but leaves open the source of the direct discrepancy between the observational and experimental results for Kenya.

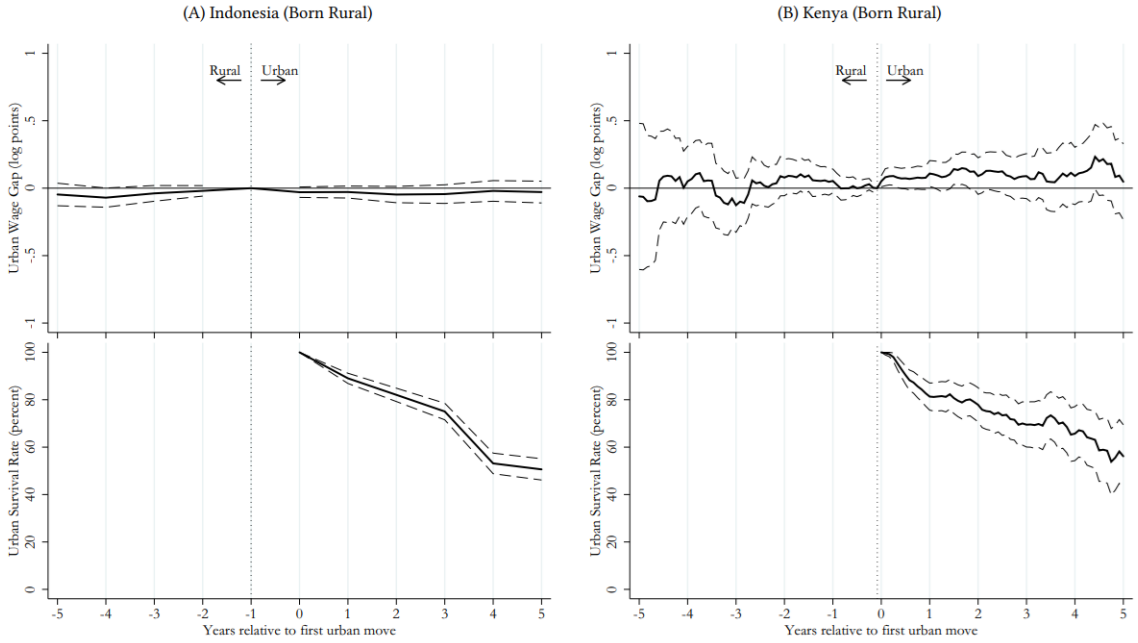
3 Setting & Data

The setting of interest is defined broadly as any geographic region (1) that is composed of distinct rural and urban locations, (2) where the labor markets are specialized between these two types (i.e. value different skills and human capital), and (3) where a large rural-urban wage gap exists[1]. This definition encompasses many developing countries, including those studied in both the experimental and observational papers summarized in Section 2.

To reconcile the returns to rural-urban migration puzzle, I will primarily study self-selected

[1]Here, "large" means the wage gap is sufficiently great to compensate for cost of living and amenity differences between the rural and urban areas. In other words, where there appear to be real returns to rural-urban migration.

FIGURE 1: EVENT STUDY OF RURAL-URBAN MIGRATION (WAGE)



Source: Hamory et al. (2020*a*)

migration in observational individual-level panel data. Following Hamory et al. (2020*a*), there are two candidate datasets: the Kenya Life Panel Survey (KLPS) and the Indonesian Family Life Survey (IFLS). I will focus throughout this paper on the KLPS for its richer panel data. The KLPS gathers data at a higher frequency (monthly as compared to yearly in the IFLS) and, more importantly, consistently records the current location of each individual at a sufficient granularity (village/city, as compared to district in the IFLS). This allows me to both identify a person's position in their origin's earnings distribution, where that distribution is either for a rural or urban labor market, and to identify lateral moves[2]. To study rural-urban migration in the IFLS requires relying on a self-reported indicator for location type (rural, town, or urban) and does not provide location data at a level conducive to my analysis. Specifically, the IFLS records an individual's location along three dimensions: province, kabupaten (district), and kecamatan (sub-district). Both provinces and kabupatens are too broad for the analysis in this paper, as they encompass both rural and urban locations. Kecamatans, conversely, are too granular and frequently include just a handful of individual-year observations. The IFLS also stopped recording recalled income data between waves after its third wave in 2000.

[2] "Lateral" moves are being defined as a rural-to-rural or urban-to-urban migration event.

3.1 Data

The main analysis in this paper builds off Hamory et al. (2020*a*)'s replication data (Hamory et al., 2020*b*). Using the KLPS, the authors constructed an unbalanced individual-level panel dataset of monthly earnings, hours worked, and wages (including formal and informal employment) spanning March 1998 to October 2014. The panel included residential location at the village/city level for all moves lasting at least four months. And it included several demographic variables (age, gender, and years of education).

The KLPS sample of 8,999 individuals was drawn from participants in two school-based RCTs that covered 200 rural primary schools in Busia district: a scholarship program for upper primary school girls in 2001 and 2002 (Kremer, Miguel and Thornton, 2009) and a deworming intervention for primary school students in 1998-2002 (Miguel and Kremer, 2004). The survey followed these students into adulthood, with an effective tracking rate of 85% in each KLPS round. Importantly, the study construction means all individuals in the sample were born in a rural location.[3]

The resulting dataset contains 148,716 individual-month observations for 6,123 individuals, of whom 799 migrate from a rural to urban location during the sample period. I winsorize values of income and wages at the 99th percentile in each year to avoid outsized impacts of outliers.[4]

4 Heterogeneity in the Returns to Rural-Urban Migration

What is the appropriate measure to estimate "returns" to migration? Hamory et al. (2020*a*) present results for wages—a classic estimate of productivity but which effectively abstracts from the importance of unemployment and underemployment in estimating the welfare effects of migration. Lagakos et al. (2020) focus instead on consumption measures, which more accurately reflect the benefits of migration, and find these average returns are larger than the wage returns estimated in Hamory et al. (2020*a*). In development settings, a focus on consumption measures is particularly relevant because a potentially large source of the return to urban migration is lower volatility of labor demand than in rural agricultural labor markets. A higher wage does not trans-

[3]This implies extra caution in interpreting the results of Section 5.4, as the earnings premiums in urban areas are necessarily relative only to the earnings of rural-born workers and not the urban labor market more broadly. Further discussion in Section 5.4

[4]My results hold if I skip the winsor step.

late to higher welfare if there are significant transition rates into and out of employment within a month.

Unfortunately, the KLPS does not track a broad consumption measure across waves. But, we can still incorporate employment volatility by estimating the effect of migration on total monthly earnings instead of wages. Reassuringly for an interpretation of earnings as a proxy for consumption, Lagakos et al. (2020) find that for the subset of their countries with both consumption and earnings data, their estimates for the returns to migration are similar across the two measures.

Figure 2 presents the event study of rural-urban migration for monthly earnings in Kenya, using the following two-way fixed effects regression specification:

$$y_{it} = \alpha_i + \gamma_t + \sum_{\tau=-60}^{-2} \beta_\tau T_{i,t+\tau} + \sum_{\tau=0}^{60} \beta_\tau T_{i,t+\tau} + \lambda_{pre} T_{i,t+\tau<-60} + \lambda_{post} T_{i,t+\tau>60} + \epsilon_{ivt} \quad (1)$$

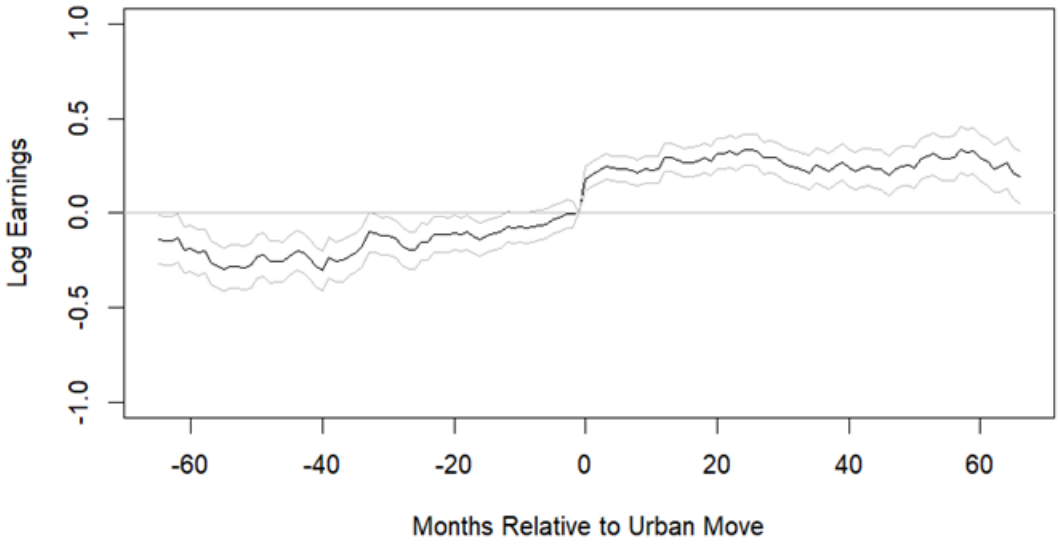
This specification includes individual and year fixed effects, uses the month prior to the urban move as the reference time period for the event study, pools observations more than five years pre or post the urban move, and clusters standard errors at the village-month level.

What is immediately obvious is that despite there being no average change in a worker's wage following a rural-urban migration event, there is a positive, significant, and sustained return for monthly earnings. This reduces some of the puzzle described in Section 2 and is consistent with the finding of positive returns for consumption and earnings in other settings (Lagakos et al., 2020).

A positive effect on earnings coupled with no effect on wages suggests that there are important employment dynamics surrounding rural-urban migration. Indeed, Figure 3 presents the event study using hours of work as the outcome variable. As expected, we see a level shift in the number of hours worked following an urban move, as well as a reduction in the volatility of hours month-to-month. This is consistent with rural-urban migration providing an important employment-smoothing benefit that carries over into smoother consumption patterns: labor demand is more consistent in urban labor markets, which generates a return to urban migration even if a worker experiences no change in their wage.

That said, is the only return to rural-urban migration an employment-smoothing one? Do

FIGURE 2: EVENT STUDY OF RURAL-URBAN MIGRATION (MONTHLY EARNINGS)



workers really experience no change in their wage despite moving across labor markets with markedly different wage distributions? Or does the average change in the wage identified by Hamory et al. (2020*a*) mask important heterogeneity in the productivity returns to rural-urban migration?

4.1 Heterogeneity by Origin Earnings Premium

I am primarily interested in whether the returns to rural-urban migration are greater for workers with a "negative origin earnings premium," i.e. those earning less than others in their village of the same age, gender, and education. This lens is motivated because this group of observational workers better reflects the compliers in most experimental research.[5]

To investigate this heterogeneity, I begin by dividing migrating workers into two groups: those with a negative origin earnings premium and those without. To do this, I first estimate time and location specific conditional earnings expectations by regressing earnings on year-month fixed effects, village fixed effects, age, gender, and years of education for the subset of my sample who are non-movers. I then estimate each migrant's expected monthly earnings using the fitted values from this regression. I compute each worker's earnings premium in each time period as the

[5]E.g. Bryan, Chowdhury and Mobarak (2014) constructs their sample by explicitly conditioning on rural households with small land holdings that had previously been exposed to hunger during the lean season in Bangladesh.

FIGURE 3: EVENT STUDY OF RURAL-URBAN MIGRATION (LABOR HOURS)



difference between their observed earnings and this estimated conditional mean. Lastly, I classify workers as having a negative origin earnings premium if their earnings premium is negative for at least 75% of observed periods prior to urban migration.

Table 2 presents estimates from the following TWFE specification:

$$y_{it} = \alpha_i + \gamma_t + \beta_1 PostUrban_{it} + \beta_2 PostUrban_{it} \times Below_i + \lambda_1 T_{i,t>60} + \lambda_2 T_{i,t>60} \times Below_i + \epsilon_{ivt} \tag{2}$$

α_i is an individual fixed effect, γ_t is a year-month fixed effect, $PostUrban_{it}$ is an indicator for periods in the first five years after a rural worker first migrates to an urban area, $Below_i$ is an indicator for workers with a negative origin earnings premium, and $T_{i,t>60}$ is an indicator for periods more than five years after urban migration. Note that the $PostUrban_{it}$ and $T_{i,t>60}$ indicators are absorbing; the estimates for β_1 and β_2 can be viewed as "ITT"-type estimates of the change in a worker's outcome after their first rural-urban migration event, *regardless of whether they subsequently return to a rural labor market*.[6]

In Table 2, we see that the observational returns to rural-urban migration are significantly

[6]In Section 5.5 I look explicitly at the rural-return rate, and at what happens to workers if they return to a rural location.

higher for workers with persistent negative origin earnings premiums. This conclusion holds in logs and levels, and for earnings and wages.[7] Specifically, the return to an urban move for migrants without a negative origin earnings premium is 11% in monthly earnings; but the return is 72% for migrants *with* a negative origin earnings premium. And, while non-negative premium migrants have almost no return in their log wage (or a small negative return in levels), negative premium migrants have a 36% return in their wage. The positive wage returns for negative premium migrants is particularly noteworthy if wages are taken as a good proxy for productivity.

TABLE 2: HETEROGENEITY IN THE RETURNS TO RURAL-URBAN MIGRATION

	(1)	(2)	(3)	(4)
	Monthly Earnings	Log Monthly Earnings	Wage	Log Wage
PostUrban	430.109	0.105***	-2.192*	0.026*
	(320.981)	(0.035)	(1.216)	(0.016)
PostUrban × NegativePremium	1,938.517***	0.386***	10.236***	0.283***
	(350.019)	(0.052)	(1.425)	(0.023)
Observations	176,150	148,415	143,108	137,237
R-squared	0.881	0.997	0.844	0.977
Pre-Mig Mean for Non-Neg Prem	6,745.452	8.492	44.859	3.263
Pre-Mig Mean for Neg Prem	2,597.811	7.625	22.403	2.655

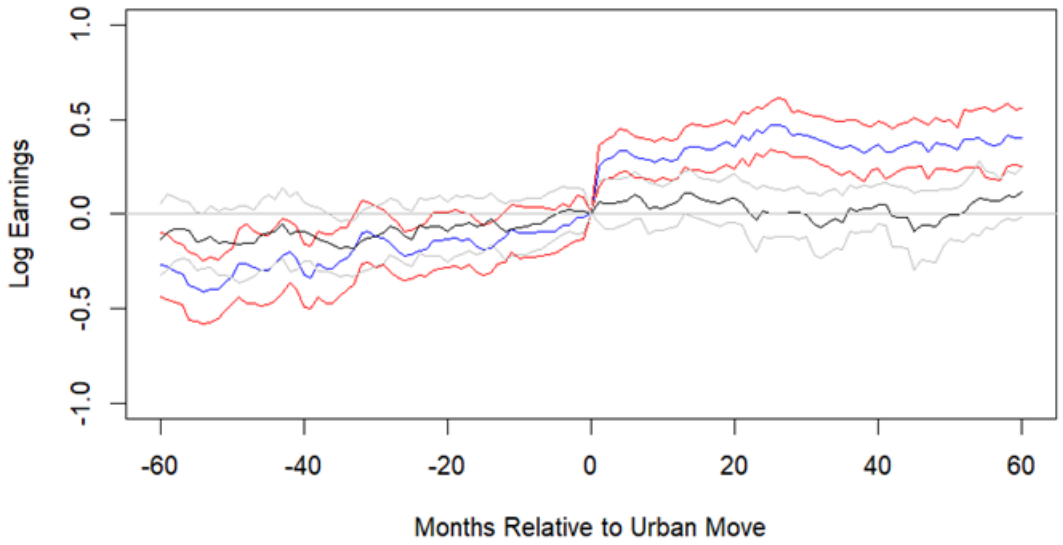
Notes: Estimates from regression Equation (2). Standard errors clustered at the village-month level. *** p<0.01, ** p<0.05, * p<0.1.

Figure 4 unpacks the earnings heterogeneity by origin premium further by presenting separate event studies of rural-urban migration for workers with persistent negative origin earnings premiums (point estimates in blue) and those without (point estimates in black). In line with the pooled post-urban estimates in Table 2, we see a much larger shift for rural workers with a negative earnings premium, and the shift is stable across the post-move period. Migrants who earn around their origin village's conditional average, meanwhile, experience no nominal returns, implying the real return (after compensating for cost of living and amenity differences) is likely negative. Comparing Figures 4 and 5, it appears the heterogeneity in the returns for earnings is driven almost entirely by heterogeneity in the returns for wages. Indeed, Figure 6 shows there are effectively no differences in the dynamics for hours worked (with the exception of possibly

[7]My preferred specification considers earnings and wages in logs because it is less sensitive to large outliers and the estimated regression is more plausibly linear in logs. Although it is impossible for both logs and levels to be linear in the same regression equation, I include the results in levels to show the estimates are robust to the inclusion of 0s, which are necessarily dropped for the log specification.

an initial positive effect on hours in the first few months post-move for migrants with a negative origin earnings premium).

FIGURE 4: EVENT STUDIES OF RURAL-URBAN MIGRATION: HETEROGENEITY BY ORIGIN EARNINGS PREMIUM (LOG MONTHLY EARNINGS)

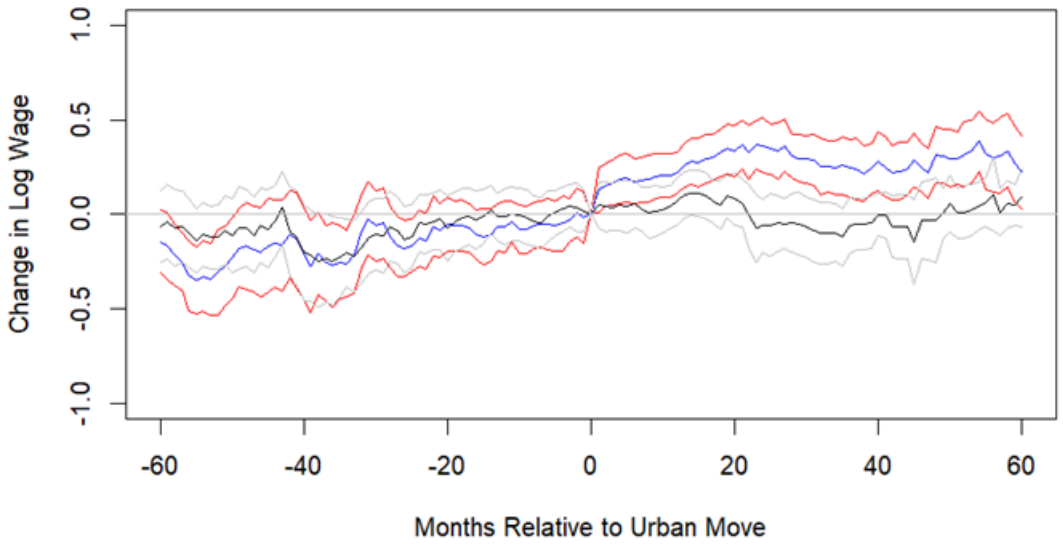


Notes: Event study for migrants with negative origin earnings premiums in blue (point estimates) and red (95% CI, standard errors clustered at the year-month level). Event study for migrants without negative origin earnings premiums in black (point estimates) and grey (95% CI, standard errors clustered at the year-month level). Estimates are from the TWFE specification in Equation 1, estimated separately for each group of migrants.

4.2 Reconciling with Experimental Data

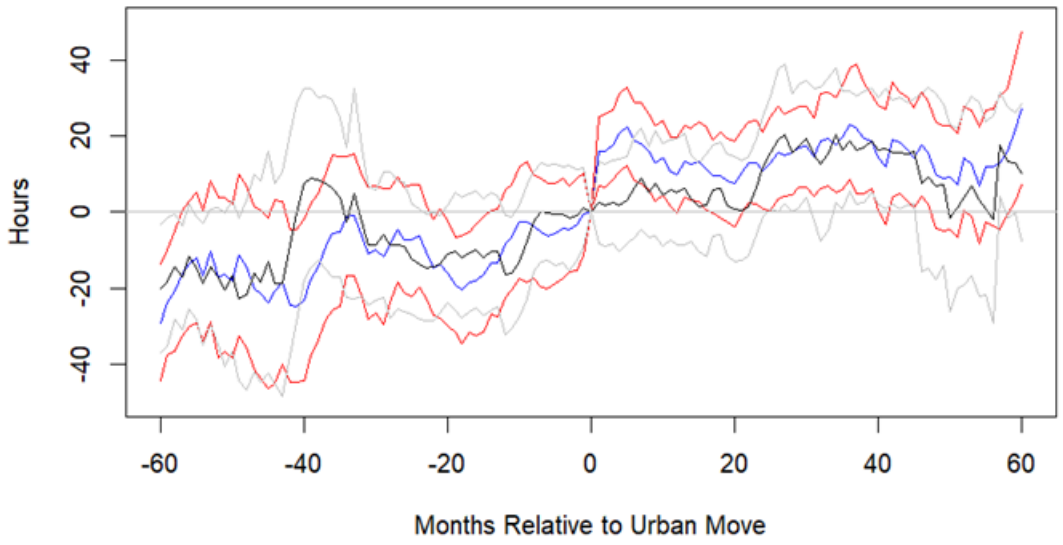
The results from Section 4.1 help resolve the puzzle posed in Section 2 for several reasons. First, I showed that despite there being no average observational return for wages, there is a positive average observational return for earnings. This aligns with the experimental results which focus on total earnings and consumption rather than wages, and further supports re-assessments of the returns to earnings and consumption in other observational settings (Lagakos et al., 2020). Second, I show that the average observational return masks important heterogeneity based on a worker's position in their origin's earnings distribution. The returns to rural-urban migration are magnitudes greater for workers earning less than their rural peers at baseline. And importantly, the estimated observational return for this class of worker exceeds all experimental returns discussed in Section 2 — i.e., the return for self-selected movers is greater than the return for non-

FIGURE 5: EVENT STUDIES OF RURAL-URBAN MIGRATION: HETEROGENEITY BY ORIGIN EARNINGS PREMIUM (LOG WAGE)



Notes: Event study for migrants with negative origin earnings premiums in blue (point estimates) and red (95% CI, standard errors clustered at the year-month level). Event study for migrants without negative origin earnings premiums in black (point estimates) and grey (95% CI, standard errors clustered at the year-month level). Estimates are from the TWFE specification in Equation 1, estimated separately for each group of migrants.

FIGURE 6: EVENT STUDIES OF RURAL-URBAN MIGRATION: HETEROGENEITY BY ORIGIN EARNINGS PREMIUM (HOURS WORKED)



Notes: Event study for migrants with negative origin earnings premiums in blue (point estimates) and red (95% CI, standard errors clustered at the year-month level). Event study for migrants without negative origin earnings premiums in black (point estimates) and grey (95% CI, standard errors clustered at the year-month level). Estimates are from the TWFE specification in Equation 1, estimated separately for each group of migrants.

movers, consistent with classic selection models. The experimental returns then lie somewhere between the zero returns for high-earning rural workers and the upper bound set by the observational returns for low-earning workers. The experimental estimate falls closer to the upper bound based on the extent to which compliers induced to migrate in RCTs more heavily weight towards low-earning rural workers.

For example, consider the seminal experimental estimate in Bryan, Chowdhury and Mobarak (2014). Their RCT design specifically samples rural households with small land holdings and previous exposure to hunger in the lean season. This selection criteria already begins targeting low-earning rural workers. And although they do not collect earnings data at baseline, so we cannot assess exactly where their compliers fell in the origin earnings distribution, they note that baseline income (proxied for with wall material) had a negative coefficient in the first-stage regression for treatment take-up. This provides suggestive evidence that their LATE is weighted more heavily by low-earning workers, explaining why it corresponds more closely to the observational returns for migrants with negative origin earnings premiums.

Consider also the experimental estimates in Kenya (Baseler, 2023). Here, the sample pool does not select for workers earning less than their peers but does collect one month of baseline earnings data. This allows us to assess where the intervention's compliers' fell in their origin's distribution. Among those induced to migrate in the treatment group, 62% had a negative earnings premium the month prior to treatment. This is not an exact replication of my procedure for identifying low-earning rural workers — it is based on just a single month's earnings whereas I select for workers with persistent negative premiums — but again is suggestive that the experimental LATE is composed more heavily of the returns for low-earning workers.

5 Sources of Heterogeneity

In the previous section, I documented that the returns to rural-urban migration in Kenya are greater for low-earning rural workers. This aligns well with the recent assessment by Cenci, Kleemans and Tjernstrom (2026) of selection and heterogeneity in the consumption returns to migration in Indonesia, China, and Tanzania. Using a group random coefficient model to flexibly estimate heterogeneity in the observational consumption returns, they found migration is a pro-poor

strategy, in that those with the lowest baseline consumption had the most to gain from migration. I have independently drawn the same conclusion for earnings using a different approach in a different setting. Consequently, it appears returns heterogeneity is a pervasive feature of rural-urban migration across the developing world.

Having established stark heterogeneity in the observational returns to rural-urban migration, I now turn to a more interesting question: *why* are the returns greater for low-earning workers? A tempting interpretation, particularly in light of the heterogeneity for wages, is that this class of worker leverages migration to resolve a spatial skill mismatch. Under this interpretation, workers can be divided into two skill types: those with skills that lead to high productivity in the specializations of rural labor markets and those with skills that lead to high productivity in urban ones. Workers with negative premiums in their origin are poorly matched to the rural labor market but are relatively better matched to an urban labor market. Hence, the effect of urban migration on their productivity is much greater than the effect on the productivity of workers well-matched in the rural market. The returns are heterogeneous because the urban place effect is heterogeneous.

However, there are several explanations that could also generate the observed returns heterogeneity even without any spatial skill arbitrage. In particular, 1) mean reversion, 2) compression in the income distribution, and 3) differential selection into urban destinations.

5.1 Mean Reversion

Suppose we conceptualize a worker's monthly earnings as a random draw from the conditional earnings distribution in their current location. Further, assume workers are completely homogeneous in productivity. Under this framework, a "negative earnings premium" then reflects persistent "bad draws" in the origin, rather than low rural productivity. A migration event, rural-to-urban or otherwise, triggers a worker to pull a new draw from a new location's distribution. Consequently, even if the expected value of earnings in each location was equivalent for all workers (i.e. no heterogeneity in place effects), we would estimate higher returns for workers who started towards the lower tail of the rural distribution because high and low earning rural workers converge to the same mean under a new urban draw. In this way, the heterogeneity results in Section 4 could be generated mechanically by segmenting workers into lucky and unlucky

pre-migration pools, but which both have the same average earnings in the urban destination.

One implication of this mechanism is that *lateral migration* — rural-to-rural or urban-to-urban — would also exhibit returns heterogeneity by baseline earnings. This leads naturally to approximating the role of mean reversion by estimating returns heterogeneity in lateral migration events. There is much less opportunity to resolve spatial skill mismatch in lateral migration because workers do not move between labor market types. Arguably, the excess returns for those with negative origin premiums then isolates the contribution of mean reversion to returns heterogeneity.

Table 3 presents estimates of the observational return to lateral migration. As expected, the returns are greater for workers with negative origin earnings premiums. But the difference relative to average workers is much smaller than in rural-urban migration. I estimate that the earnings return to lateral migration for workers without negative premiums is 7%, while it is 15% for low-earning workers. Similarly, the returns to wages are 4.8% for migrants without negative premiums and 14% for low-earning ones. Recall that the observational return to rural-urban migration is more than 6x greater for low-earning workers, significantly beyond the 2x estimated difference for lateral migration.

Figure 7 plots separate event studies of lateral migration for workers with and without negative premiums. Compared to rural-urban migration in Figure 2, lateral migration does not exhibit clear heterogeneity in the post-move period. The comparison between rural-urban and lateral migration is of course imperfect, but as a rough approximation the results here suggest that much of the heterogeneity in the returns to rural-urban migration *cannot* be explained by mean reversion.

5.2 Earnings Compression

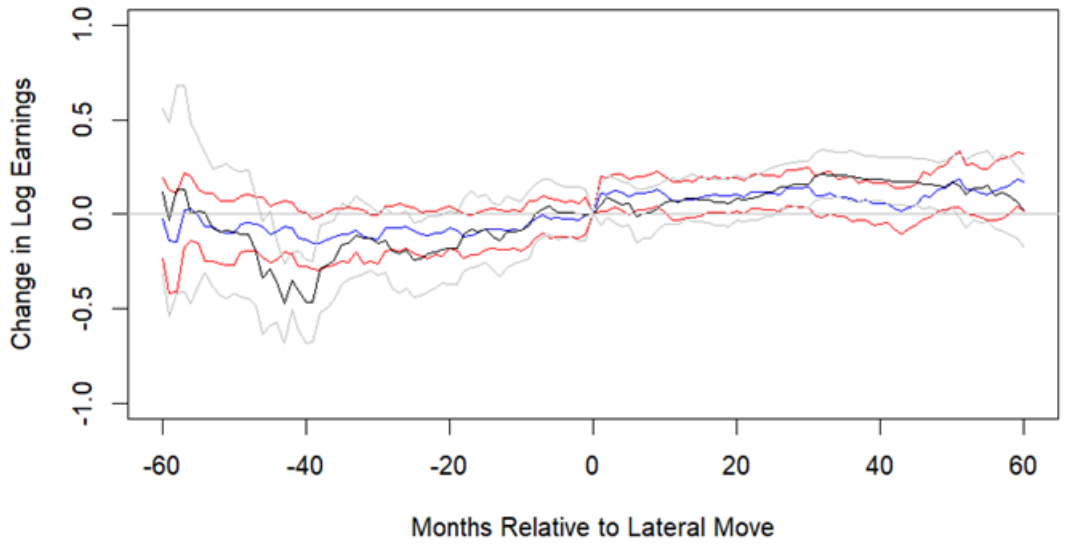
The earnings compression explanation proceeds as follows. Suppose there is no heterogeneity in the productivity effect of urban migration. Workers with relatively low productivity in the rural area have relatively low productivity in the urban area. Taken to an extreme, a worker who falls in the ith percentile of their rural origin's earnings distribution will fall in the ith percentile of the urban destination's earnings distribution. If the earnings distribution is *tighter* in the urban labor market than the rural one, then we would nonetheless observe the type of heterogeneity in the returns to rural-urban migration described in Section 4. Put simply, a move from the 20th

TABLE 3: HETEROGENEITY IN THE RETURNS TO LATERAL MIGRATION

	(1)	(2)	(3)	(4)
	Monthly Earnings	Log Monthly Earnings	Wage	Log Wage
PostMove	-430.572	0.070***	-2.512	0.047***
	(339.130)	(0.017)	(2.096)	(0.017)
PostMove × NegativePremium	1,259.160***	0.073***	8.085***	0.083***
	(357.450)	(0.022)	(2.315)	(0.023)
Observations	27,419	25,853	24,523	23,860
R-squared	0.839	0.996	0.9244	0.974
Mean Pre-Mig for Non-Negative Prem	11,063.74	8.866	64.506	3.551
Mean Pre-Mig for Negative Prem	3,606.56	7.854	28.669	2.776

Notes: Estimates from regression Equation (2) modified so that the PostUrban indicator is instead an indicator for periods after a rural-rural or urban-urban move. Standard errors clustered at the year-month level. *** p<0.01, ** p<0.05, * p<0.1.

FIGURE 7: EVENT STUDY OF LATERAL MIGRATION: HETEROGENEITY BY ORIGIN EARNINGS PREMIUM



Notes: Event study for migrants with negative origin earnings premiums in blue (point estimates) and red (95% CI, standard errors clustered at the year-month level). Event study for migrants without negative origin earnings premiums in black (point estimates) and grey (95% CI, standard errors clustered at the year-month level). Estimates are from the TWFE specification in Equation 1, modified so that the event study indicators are instead relative to an individual's first rural-rural or urban-urban move, estimated separately for each group of migrants.

percentile of the rural distribution to the 20th percentile of the urban distribution is greater than the move from the 75th to the 75th percentile.

I test for earnings compression by running quantile regressions of log earnings and log wages on an indicator variable for urban observations. Under a null of no compression, there would be no difference in the urban coefficient at different quintiles, so that the urban distribution is simply a level shift of the rural one. Table 4 provides evidence against this null. Specifically, the rural-urban earnings and wage gaps are decreasing in quintiles. This pattern holds with controls and when restricting the sample to rural-urban migrants (although the pattern weakens for the earnings regressions restricted to rural-urban migrants). This implies there is both earnings and wage compression at the bottom of the distribution that could explain larger returns for low-earning rural workers without any heterogeneous urban productivity gains.

However, the magnitude of the earnings compression cannot fully account for the returns heterogeneity. Focusing on the quantile regressions for rural-urban migrants in columns (5) and (6), under perfect rank preservation we would expect at most a 20 log point difference between the returns for low-earners and high-earners. But the estimates in Table 2 show that the returns for low-earners are 30 to 50 log points greater. This suggests that rank is not perfectly preserved when workers migrate from a rural to urban labor market: low-earners appear to move up the distribution and high-earners down the distribution.

5.2.1 Mean Reversion with Compressed Distributions

In light of the earnings compression between rural and urban labor markets, it is worth revisiting the results of Section 5.1. Section 5.1 argued that mean reversion cannot fully explain the heterogeneity in returns to rural-urban migration because the difference in returns for high versus low baseline earners is much smaller in lateral migration than rural-urban migration. But that conclusion assumes lateral migration — where workers migrate between similar earnings distributions — provides a good approximation for the scope of mean reversion in rural-urban migration. How much heterogeneity can be generated by mean reversion if instead workers migrate to a more compressed earnings distribution?

I conduct a simple Monte Carlo experiment to explore this question. First, I construct separate

TABLE 4: QUANTILE REGRESSIONS OF THE RURAL-URBAN GAP

	(1)	(2)	(3)	(4)	(5)	(6)
	Log Monthly Earnings	Log Wage	Log Monthly Earnings	Log Wage	Log Monthly Earnings	Log Wage
10th Percentile	1.240	0.782	1.107	0.806	0.431	0.321
	(0.036)	(0.015)	(0.014)	(0.012)	(0.000)	(0.029)
20th Percentile	1.139	0.729	1.072	0.714	0.420	0.306
	(0.026)	(0.014)	(0.008)	(0.008)	(0.021)	(0.028)
30th Percentile	1.022	0.591	0.964	0.622	0.465	0.236
	(0.020)	(0.004)	(0.007)	(0.007)	(0.024)	(0.019)
40th Percentile	0.898	0.619	0.892	0.542	0.480	0.224
	(0.019)	(0.010)	(0.006)	(0.007)	(0.017)	(0.018)
50th Percentile	0.944	0.581	0.796	0.484	0.449	0.180
	(0.023)	(0.009)	(0.006)	(0.006)	(0.019)	(0.015)
60th Percentile	0.847	0.506	0.746	0.423	0.437	0.122
	(0.016)	(0.007)	(0.005)	(0.007)	(0.014)	(0.019)
70th Percentile	0.811	0.511	0.669	0.378	0.401	0.141
	(0.020)	(0.006)	(0.006)	(0.006)	(0.016)	(0.025)
80th Percentile	0.693	0.551	0.594	0.303	0.390	0.163
	(0.023)	(0.010)	(0.007)	(0.010)	(0.017)	(0.032)
90th Percentile	0.715	0.405	0.498	0.255	0.228	0.131
	(0.019)	(0.019)	(0.009)	(0.010)	(0.023)	(0.027)
Observations	148,415	137,237	148,415	137,237	29,260	27,045
Controls	N	N	Y	Y	Y	Y
Origin Fixed Effects	N	N	N	N	Y	Y
Sample	All	All	All	All	Migrants	Migrants

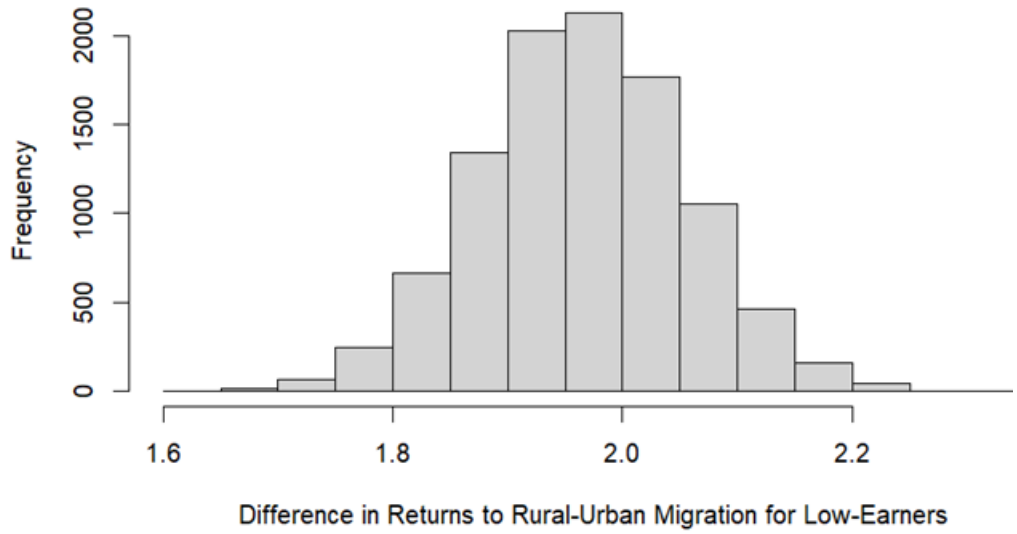
Notes: Estimates from quantile regressions of the outcome variable on an indicator for urban. Controls include age, gender, and years of education. *** p<0.01, ** p<0.05, * p<0.1.

estimates of the rural and urban log earnings distributions by selecting parameters of a normal distribution to maximize the likelihood of observing the rural and urban empirical CDFs. I then simulate an extreme form of mean reversion, whereby workers draw log earnings from the rural and urban distributions *independently*, and compute the difference in mean "return" between workers who drew a below-average rural income and those who drew an above-average rural income: an analog to $\hat{\beta}_2$ in regression equation (2). Figure 8 presents the histogram of results across 10,000 iterations, where each iteration draws a sample of 800 individuals (the number of rural-urban migrants in the observational data). Figure 9 presents results from the same experiment but where the empirical CDFs only include earnings observations for movers, to better match the earnings distributions of rural-urban migrants. Both Monte Carlo experiments find that if earnings states were independently drawn, then mean reversion across the empirically compressed earnings distributions would generate an estimated $\hat{\beta}_2$ on the order of 1.8 to 2.1. This is significantly greater than the empirically estimated 0.4. Figure **??** presents results from the same Monte Carlo experiment, but assuming workers can only mean revert within types: the urban earnings distributions are estimated separately for high and low earning rural migrants, and workers draw an independent urban income conditional on whether their rural income was above or below the mean. The histogram of returns heterogeneity shifts further to the left under this assumption, but centers around 120 log points — still four times greater than the empirical 0.4. In fact, the empirical estimate falls so far outside the distribution of Monte Carlo iterations that it appears a low rural earnings state likely reflects individual differences in ability that translate across labor markets, rather than unlucky idiosyncratic draws from the rural earnings distribution.

5.3 Differential Selection into Migration

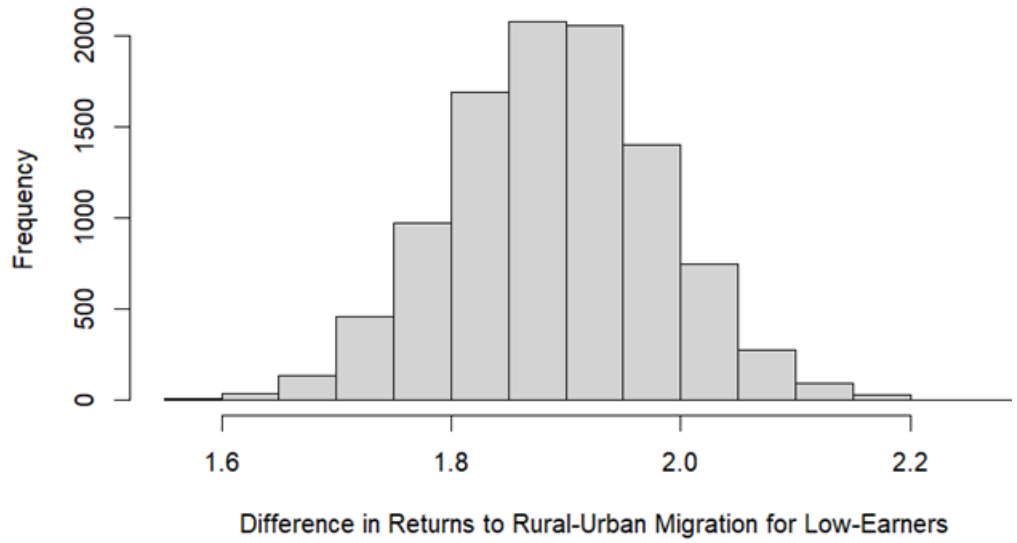
Another driver of the heterogeneous returns might be differences in the motivation for migration. Specifically, low-earning rural workers are likely more motivated by income gains than high-earning workers, who may instead migrate for non-labor-market reasons. This would generate the observed returns pattern through differential selection on outcomes. I test for whether the returns heterogeneity is generated by differences in destination choice by running a "first-stage" regression of migration to each of Kenya's five largest cities on a dummy for a baseline negative

FIGURE 8: MONTE CARLO RESULTS: FULL LOG EARNINGS DISTRIBUTIONS



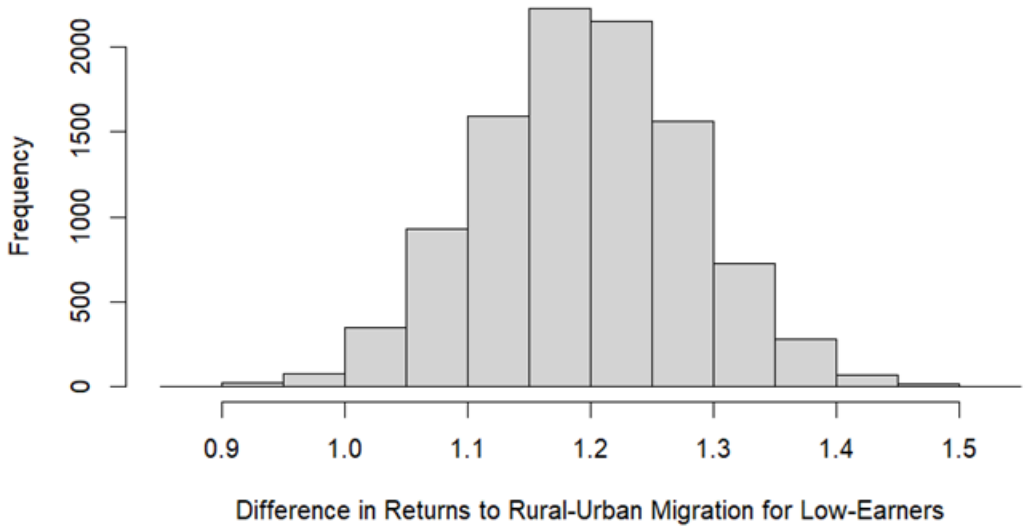
Notes: Histogram of Monte Carlo Results for the Difference in the Average Return to Rural-Urban Migration between Low Earning and High Earning Rural Workers. Assumes workers' draws from the rural and urban distributions are independent. Rural and urban log earnings distributions are constructed by assuming a normal approximation and selecting the mean and standard deviation parameters that maximize the likelihood of all empirical observations in the KLPS.

FIGURE 9: MONTE CARLO RESULTS: MOVERS' LOG EARNINGS DISTRIBUTIONS



Notes: Histogram of Monte Carlo Results for the Difference in the Average Return to Rural-Urban Migration between Low Earning and High Earning Rural Workers. Assumes workers' draws from the rural and urban distributions are independent. Rural and urban log earnings distributions are constructed by assuming a normal approximation and selecting the mean and standard deviation parameters that maximize the likelihood of the rural-urban migrants' subset of empirical observations in the KLPS.

FIGURE 10: MONTE CARLO RESULTS: MOVERS' LOG EARNINGS DISTRIBUTIONS, MEAN REVERSION WITHIN TYPES



Notes: Histogram of Monte Carlo Results for the Difference in the Average Return to Rural-Urban Migration between Low Earning and High Earning Rural Workers. Assumes workers' draws from the rural and urban distributions are independent, but that workers draw from separate urban earnings distributions conditional on type (below vs above average draw in the rural distribution). Rural and urban log earnings distributions are constructed by assuming a normal approximation and selecting the mean and standard deviation parameters that maximize the likelihood of the rural-urban migrants' subset of empirical observations in the KLPS.

earnings premium. I also include origin village fixed effects and a set of observables in a separate specification. The results are presented in Table 5. I find no statistically significant difference in the choice of urban destination between high and low earning rural workers, other than a slightly greater propensity among low earners to migrate to Kisumu (which accounts for just 4% of rural-urban migration). Although this does not negate the possibility that low-earners differentially select into migration across time (i.e. when their idiosyncratic returns are highest), it does imply that regardless of motivation, low-earning rural workers gain more from urban migration than high-earning workers in the same destination.

TABLE 5: TEST FOR DIFFERENTIAL URBAN DESTINATION CHOICE

	(1)	(2)
	Low Earners	Low Earners
Nairobi	-0.021	-0.028
	(0.034)	(0.038)
Mombasa	-0.030	-0.017
	(0.023)	(0.024)
Kisumu	0.034***	0.027**
	(0.008)	(0.012)
Nakuru	0.012	0.008
	(0.011)	(0.012)
Eldoret	0.007	-0.003
	(0.010)	(0.013)
Observations	829	829
Controls	N	Y
Origin Fixed Effects	N	Y

Notes: Estimates from linear regressions of migration to each destination on an indicator for low-earning rural migrants. Controls include age, gender, and years of education. Standard errors clustered at the origin village level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

5.4 Resolving a Skill Mismatch

If the source of heterogeneous returns is that rural workers with negative premiums are skill-mismatched and they resolve this spatial skill-mismatch by migrating to an urban labor market, then we would expect to observe a positive effect on their earnings premium post-move: we should see negative premium migrants move up the conditional earnings distribution, as they are better matched in the urban labor market.

Table 6 presents estimates of the effect of rural-urban migration on a worker's earnings premium (how much they earn in excess of the mean worker in the same village/city of the same gender, age, and education). For migrants without a negative origin premium, urban migration *reduces* their earnings premium by 67 log points, on a pre-migration mean of 0.710. This implies rural-urban migration effectively eliminates the premium these migrants enjoyed in their rural origin, and on average they earn the conditional mean among workers in their urban destination. For migrants with a negative origin premium, urban migration increases their earnings premium by 11 log points, on a pre-migration mean of -0.671. Although the effect is positive, the average migrant in this group still earns well below the conditional mean in their urban destination.

Figure 11 plots event studies for the change in log earnings premiums. Rural workers without a negative origin premium (point estimates in black) experience a distinct drop in their premium in the first urban month with no subsequent change. Low-earning rural workers, meanwhile, experience a comparatively mild drop in the first urban month, but slowly erode at this effect, so that two years post-move their premium is equivalent to their premium just before migrating.

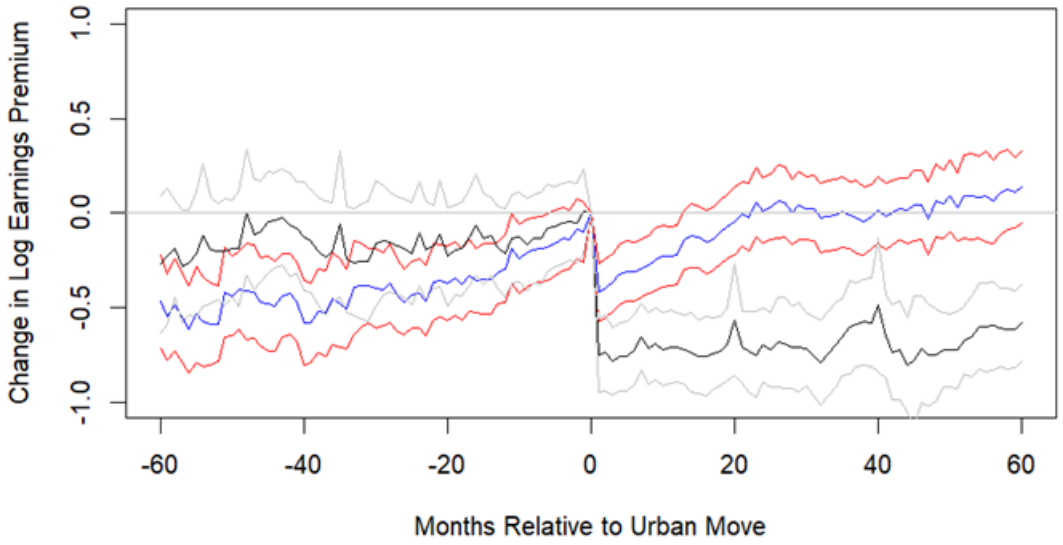
The event study in particular suggests an inversion of the framing of this paper. Returns to rural-urban migration are heterogeneous **not** because low-earning rural workers resolve a skill mismatch and gain more in the urban market; rather, they are heterogeneous because high-earning rural workers *give up their productivity advantage* when they migrate. In some sense, this suggests that a large source of the returns heterogeneity in Section 4 is that rural-born workers are 1) very heterogeneous in rural productivity, but 2) relatively homogeneous in urban productivity. If we return to the results in Table 2 column (4), and add the returns to migration to each group's pre-migration mean, we see that rural-urban migration significantly closes the productivity gap between migrants with and without a negative origin premium. There is a 0.60 difference in log wages pre-migration, but just a 0.315 difference post-migration. Combined with the effects of migration on earnings premiums, much of the decline in this gap appears to come from negative productivity effects at the top rather than positive effects at the bottom.

TABLE 6: HETEROGENEITY IN THE RETURNS TO RURAL-URBAN MIGRATION (EARNINGS PREMIUM)

	(1)	(2)
	Earnings Premium	Log Earnings Premium
PostUrban	-2711.690***	-0.595***
	(117.359)	(0.022)
PostUrban × NegativePremium	2590.481***	0.775***
	(137.790)	(0.027)
Observations	176,149	142,320
R-squared	0.740	0.766
Pre-Mig Mean for Non-Negative Premium	3752.218	0.668
Pre-Mig Mean for Negative Premium	-1716.908	-0.642

Notes: Estimates from regression Equation (2). Standard errors clustered at the village-month level. *** p<0.01, ** p<0.05, * p<0.1.

FIGURE 11: EVENT STUDY OF RURAL-URBAN MIGRATION (EARNINGS PREMIUM)



Notes: Event study for migrants with negative origin earnings premiums in blue (point estimates) and red (95% CI, standard errors clustered at the year-month level). Event study for migrants without negative origin earnings premiums in black (point estimates) and grey (95% CI, standard errors clustered at the year-month level). Estimates are from the TWFE specification in Equation 1, estimated separately for each group of migrants.

5.5 Rural Return

The analysis so far, in an ITT-style, treated all observations after the first rural-urban migration event as "post-urban." In this section I present some descriptive statistics for migrants who choose to return to a rural labor market, with a focus on heterogeneity by original rural earnings premium.

Figure 12 plots the urban survival rate for workers with a negative origin earnings premium in red and for those without in black. Workers with a negative premium are particularly more likely to return to a rural location 6-10 months after migrating to the urban location. There is not much distinct heterogeneity in rural return trends more than a year out. Five years after migrating, approximately 60% of workers with a negative rural earnings premium and 75% of workers without a negative premium never left the urban location.

Table 7 presents estimates of the change in income and wages for the migrants who return to a rural labor market. These estimates are from a TWFE regression of the form specified in Equation 2 but restricts the panel to post-urban observations among individuals who ever selected into rural-urban migration. In effect, the coefficients are the average change in an individual's earnings or wage after returning to a rural location, net of a weighted average of individual changes among those who remained urban around the time of the return. Heterogeneity is based on having a negative earnings premium in the *rural* location *prior* to urban migration — so workers are in the same classification that they belonged to for the analysis in Sections 4 and the rest of 5.

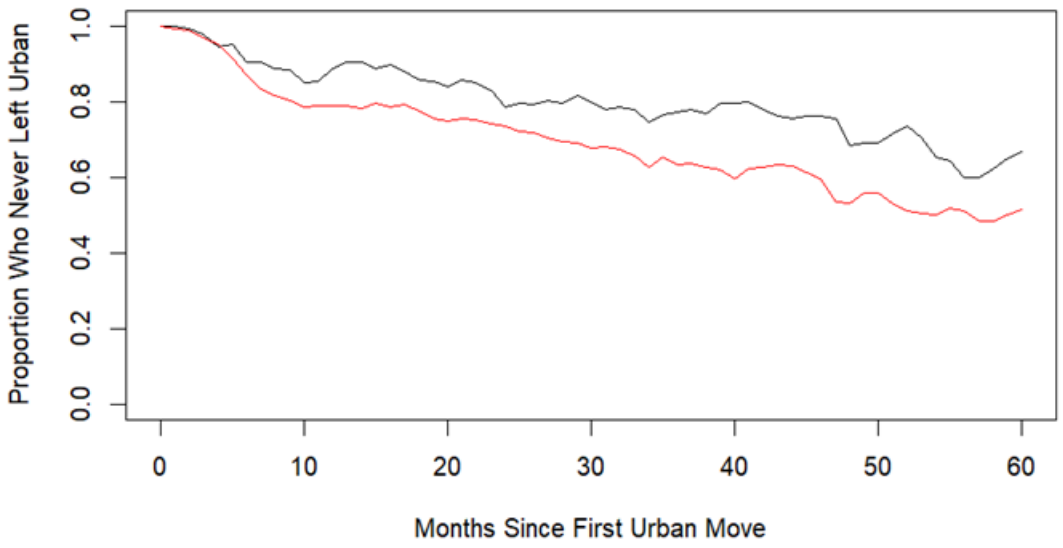
Both types of workers experience a drop in earnings and wages after returning to the rural labor market. There is no statistically significant heterogeneity for the change in wage. The drop in earnings is smaller for negative premium workers in levels but greater in logs. This discrepancy in sign is due either to the effect of positive outliers in the negative premium group or a higher rate of 0 earnings in the non-negative premium group. But reading the table as a whole suggests there is not systematic heterogeneity in the negative change in income and wages following a rural return based on a migrant's original rural earnings premium. Relative to what they earned in the urban labor market, migrants who conduct rural return experience a 19-27 log point decline in monthly earnings and a 17 log point decline in their wage.

It is not clear how to interpret the lack of heterogeneity in rural return. On its face, it is suggestive of somewhat permanent productivity effects from working in the urban labor market:

workers started with relatively large gaps in rural productivity (i.e. wages), the gap in productivity shrank when they migrated to the urban location, and the gap did not widen again when they returned to the rural market. However, the subset of workers who return to the rural location are a self-selected slice of the original rural-urban migrant pool. In principle, those selecting to return had lower gains from urban migration and may not have been all that heterogeneous in rural productivity to begin with.

To disentangle these two interpretations, Table 8 presents estimates of the observational returns to rural-urban migration for the subsample of migrants who ultimately return to the rural location. First, it is noteworthy that the pre-migration means for workers with negative and non-negative premiums are similar in levels and differences to those reported in Table 2. This suggests that the pool of rural returners did not systematically differ, as compared to urban survivors, prior to urban migration. However, consistent with selection on outcomes, the returns to rural-urban migration are lower for those who ultimately migrate back to a rural labor market than for those who remain in the urban one. Recall in Table 2 the estimated returns on log earnings and log wages were 10 and 3 log points, respectively, for workers with non-negative premiums, and were 49 and 31 log points for workers with negative premiums. Among those who later migrate back to a rural location, these estimates fall to 6 and -7 log points for workers with non-negative premiums, and 28 and 14 log points for workers with negative premiums. So, those selecting into rural return do appear to have had lower gains to urban migration than those who remained urban. But, these results suggest there is some permanency to the urban productivity effect. Looking narrowly at the the evolution of the effect of migration on workers who complete a full migratory cycle, we see that rural-urban migration shrank the gap between the wages of workers with and without negative rural premiums, and the subsequent urban-to-rural migration did not reverse this. However, it is possible that this evolution is due to a form of the mean reversion discussed in Section 5.1: with each wage draw in a new location, workers move closer and closer to the average for that labor market.

FIGURE 12: HETEROGENEITY IN THE URBAN SURVIVAL RATE BY ORIGIN EARNINGS PREMIUM



Notes: Plots the proportion of rural-urban migrants without a rural observation x months after their first urban migration event. Urban survival rate for migrants with a negative origin earnings premium in red, and for those without a negative premium in black. The survival rate can increase due to changes in the sample composition at different months-post-move.

TABLE 7: HETEROGENEITY IN RURAL RETURN

	(1)	(2)	(3)	(4)
	Monthly Earnings	Log Monthly Earnings	Wage	Log Wage
RuralReturn	-3525.534***	-0.1901***	-13.439***	-0.174***
	(373.520)	(0.028)	(4.083)	(0.042)
RuralReturn × NegativeOriginPremium	1395.2806***	-0.084**	1.667	0.048
	(324.734)	(0.041)	(3.940)	(0.048)
Observations	19,134	16,886	16,558	15,819
R-squared	0.814	0.997	0.921	0.980

Notes: Estimates from regression of the form in Equation 2, modified to treat a rural return as the post-treatment period and excluding the $T_{i,t>60}$ terms. Panel restricted to the post-urban-migration periods for individuals that ever migrated from a rural to urban labor market. Standard errors clustered at the year-month level. *** p<0.01, ** p<0.05, * p<0.1.

TABLE 8: HETEROGENEITY IN THE RETURNS TO RURAL-URBAN MIGRATION AMONG ULTIMATE RETURNERS

	(1)	(2)
	Log Monthly Earnings	Log Wage
PostUrban	0.055***	-0.068***
	(0.018)	(0.024)
PostUrban × NegativePremium	0.229***	0.203***
	(0.024)	(0.028)
Observations	130,868	120,498
R-squared	0.997	0.978
Pre-Mig Mean for Non-Neg Prem	8.568	3.379
Pre-Mig Mean for Neg Prem	7.536	2.669

Notes: Estimates from regression Equation (2), with the migrant pool limited to those who later return to a rural location. Standard errors clustered at the year-month level. *** p<0.01, ** p<0.05, * p<0.1.

6 Barriers to Rural-Urban Migration

If the returns to rural-urban migration are larger for workers with negative origin earnings premiums, what prevents more low-earning workers from leveraging the spatial arbitrage opportunity?

As discussed in Section 2, a robust literature in development economics has studied the barriers to rural-urban migration in experimental settings. This literature has shown rural workers respond to incentives (Bryan, Chowdhury and Mobarak, 2014; Bryan et al., 2023), information (Baseler, 2023), networks (Blumenstock, Chi and Tan, 2025), risk insurance (Miner, 2026), and relaxed budget constraints (Meghir et al., 2022; Akram, Chowdhury and Mobarak, 2025). Although I do not attempt to add causal evidence to this literature, the observational data in this paper offers some suggestive "clues" as to why low-earning rural workers do not migrate at higher rates.

Returning to Figure 4, we see a distinct upward trend in the earnings of workers with negative premiums leading up to the urban migration event. This is consistent with low-earning workers being unable to afford a move until their income rises and they accumulate more savings to fund migration. "Funding" migration likely includes both paying direct moving costs and stockpiling sufficient buffer savings to insure against uncertain urban labor market prospects. The pre-migration upward earnings trend is especially telling because it is unique to low-earning migrants. We do not see a rise in earnings precipitate migration for high-earning workers. Al-

though other systematic differences certainly exist between these two groups of migrants, there is a clear income gap. So, a pre-migration upward trend in income for low-earners combined with no upward trend for high-earners is suggestive of insufficient savings being an important barrier to migration. The importance of accumulating buffer savings before selecting into migration is further supported by the relatively steep increase in rural return for low-earning workers 6-10 months after migrating (Figure 12). Workers with persistent negative premiums (in red) likely have shorter insurance runways (due to a smaller buffer of savings), and so are forced to abort their urban migration attempt earlier if they struggle to find work or do not experience a large enough increase in their wage to justify the higher cost of living.

Assuming buffer savings and binding budget constraints are a salient barrier to rural-urban migration for the workers with the most to gain, an important question is whether the observed pre-migration rise in income is endogenous or exogenous. Do workers decide they want to migrate, then endogenously increase their earnings in anticipation of moving; or, do they have to wait for a sufficient exogenous shock to their earnings to trigger migration? Assuming a worker's wage is largely exogenous,[8] returning to Figure 5 suggests that exogenous shocks are an important contributor. Although less distinct than the pre-migration rise in earnings, there is an upward trend in the wages of low-earning workers who select into rural-urban migration. Hours of labor are less clearly exogenous for a worker and may reflect more endogenous decision-making in anticipation of migration. Figure 6 shows that hours worked trends upwards before migration for workers with both negative and non-negative origin premiums. To the extent that labor supply is inelastic, this upward trend can be read as further evidence for rural workers requiring positive labor market shocks to induce a move. However, labor supply is unlikely to be perfectly inelastic, and so this trend certainly leaves room for endogenous increases in hours either 1) in anticipation of funding migration, or 2) in response to the higher wage seen in Figure 5. Ultimately, experimental research offers more compelling insight into the most pertinent barriers to rural-urban migration. But, the observational evidence is at least consistent with the importance of binding budget constraints, regardless of whether the pre-migration rise in earnings is primarily an endogenous or exogenous phenomenon.

It is worth noting that the conclusion I draw here — that binding credit constraints and insuf-

[8] This exogeneity is supported by studies of rural labor markets in developing countries. See [insert citations].

ficient buffer savings restrict rural-urban migration among those with the highest returns — appears to contrast with Lagakos, Mobarak and Waugh (2023) and Kleemans and Magruder (2018). Lagakos, Mobarak and Waugh (2023) estimate a dynamic incomplete-markets model to match the experimental data in Bryan, Chowdhury and Mobarak (2014) and conclude that migration decisions are not very constrained by a lack of credit. Rather, seasonal migration is, as they put it, "an act of desperation" in *response* to uninsurable negative shocks in the rural location, and is not a "move to opportunity." And Kleemans and Magruder (2018) find that rural households with low asset levels in Indonesia send more temporary migrants in response to negative rainfall shocks. Two responses. First, the nature of my dataset focuses my analysis on permanent rather than seasonal migration. Although seasonal migration may respond positively to negative origin shocks and may not be as concerned with binding credit, the decision to permanently migrate may be more sensitive to credit constraints and insurance due to higher migratory costs and more consequential labor market placement in the destination. Second, consistent with both Lagakos, Mobarak and Waugh (2023) and Kleemans and Magruder (2018) I also find that migration rates are higher for low-earning rural workers (on the order of 2 percentage points). But, this simple statistic misses important nuance in the time-path of earnings prior to migrating. As seen in the event studies in Figures 4-6, the decision to migrate is preceded by an upward earnings and wage trajectory. So, low-earning workers are more inclined to migrate, but only after experiencing positive income shocks. This insight aligns with the model of Kleemans (2023), which highlights that a migratory response can be triggered by both contemporaneous negative and preceding positive income shocks.

7 Conclusion

In this paper, I have documented that, consistent with experimental research, there are causal returns to rural-urban migration and wage gaps do not simply reflect selection. But not all workers experience positive gains. The largest returns to rural-urban migration accrue to workers with low income in the rural origin; whereas workers with rural earnings premiums gain little from urban migration. Importantly for policy, my findings suggest that a primary source of the returns heterogeneity in Kenya is wage compression, resulting from negative premium effects at the top of

the migrant distribution. Whether we should expect the results from Kenya to generalize to other developing economies then depends crucially on whether other settings match this characteristic of comparatively narrow urban job opportunities for rural migrants. In these settings,policies that induce migration among higher-earning rural workers should expect to do little for welfare. But policies that relax the budget constraint of low-earning workers will likely generate larger gains.

References

Akram, Agha Ali, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak. 2025. "Experimental Evidence on the Network Effects of Migration."

Alvarez, Jorge A. 2020. "The Agricultural Wage Gap: Evidence from Brazilian Micro-data." *American Economic Journal: Macroeconomics*, 12(1): 153–173.

Baseler, Travis. 2023. "Hidden Income and the Perceived Returns to Migration." *American Economic Journal: Applied Economics*, 15(4): 321–352.

Baseler, Travis, Zachary Barnett-Howell, Thomas Ginn, and Stepan Gordeev. 2025. "Reaching the Novice or Nudging the Expert? Networks, Information, and the Experimental Returns to Migration."

Blumenstock, Joshua E, Guanghua Chi, and Xu Tan. 2025. "Migration and the Value of Social Networks." *The Review of Economic Studies*, 92(1): 97–128.

Bryan, Gharad, Shyamal Chowdhury, Ahmed Mushfiq Mobarak, Melanie Morten, and Joeri Smits. 2023. "Encouragement and distortionary effects of conditional cash transfers." *Journal of Public Economics*, 228: 105004.

Bryan, Gharad, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak. 2014. "Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh." *Econometrica*, 82(5): 1671–1748. Publisher: The Econometric Society.

Cenci, Eduardo, Marieke Kleemans, and Emilia Tjernstrom. 2026. "Selection and Heterogeneity in the Returns to Migration."

Gollin, Douglas, David Lagakos, and Michael E. Waugh. 2014. "Agricultural Productivity Differences across Countries." *American Economic Review*, 104(5): 165–170.

Hamory, Joan, Marieke Kleemans, Nicholas Y Li, and Edward Miguel. 2020*a*. "Reevaluating Agricultural Productivity Gaps with Longitudinal Microdata." *Journal of the European Economic Association*, 19(3): 1522–1555. _eprint: https://academic.oup.com/jeea/article-pdf/19/3/1522/53253861/jvaa043.pdf.

Hamory, Joan, Marieke Kleemans, Nicholas Y. Li, and Edward Miguel. 2020*b*. "Replication Data for: âReevaluating Agricultural Productivity Gaps with Longitudinal Microdataâ."

Kleemans, Marieke. 2023. "Migration Choice under Risk and Liquidity Constraints."

Kleemans, Marieke, and Jeremy Magruder. 2018. "Labour Market Responses To Immigration: Evidence From Internal Migration Driven By Weather Shocks." *The Economic Journal*, 128(613): 2032–2065. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/ecoj.12510.

Kremer, Michael, Edward Miguel, and Rebecca Thornton. 2009. "Incentives to Learn." *The Review of Economics and Statistics*, 91(3): 437–456. Publisher: The MIT Press.

Lagakos, David. 2020. "Urban-Rural Gaps in the Developing World: Does Internal Migration Offer Opportunities?" *Journal of Economic Perspectives*, 34(3): 174–192.

Lagakos, David, Ahmed Mushfiq Mobarak, and Michael E. Waugh. 2023. "The Welfare Effects of Encouraging RuralâUrban Migration." *Econometrica*, 91(3): 803–837. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA15962.

Lagakos, David, Samuel Marshall, Ahmed Mushfiq Mobarak, Corey Vernot, and Michael E. Waugh. 2020. "Migration costs and observational returns to migration in the developing world." *Journal of Monetary Economics*, 113: 138–154.

Meghir, Costas, A Mushfiq Mobarak, Corina Mommaerts, and Melanie Morten. 2022. "Migration and Informal Insurance: Evidence from a Randomized Controlled Trial and a Structural Model." *The Review of Economic Studies*, 89(1): 452–480.

Miguel, Edward, and Michael Kremer. 2004. "Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities." *Econometrica*, 72(1): 159–217. Publisher: [Wiley, Econometric Society].

Miner, Gwyneth. 2026. "Reducing Migration Uncertainty: The Impact of an Income Smoothing Program for Kenyan Migrants." *Working Paper*.

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