

NBER WORKING PAPER SERIES

REEVALUATING AGRICULTURAL PRODUCTIVITY GAPS WITH LONGITUDINAL
MICRODATA

Joan Hamory Hicks
Marieke Kleemans
Nicholas Y. Li
Edward Miguel

Working Paper 23253
<http://www.nber.org/papers/w23253>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2017

We would like to thank David Albouy, Jorge Alvarez, Lori Beaman, Ben Faber, Fred Finan, Cecile Gaubert, Doug Gollin, Seema Jayachandran, Supreet Kaur, Andrés Rodríguez-Clare, Todd Schoellman, John Strauss, Chris Udry, and seminar participants at U.C. Berkeley, University of Illinois, Northwestern University, University of Southern California, and the IMF for many helpful discussions and suggestions. We are grateful to Brian Feld and Vedika Ahuja for excellent research assistance. Hicks and Miguel gratefully acknowledge support from NIH grant R01-HD044475, which also supported creation of the KLPS dataset. All errors remain our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Joan Hamory Hicks, Marieke Kleemans, Nicholas Y. Li, and Edward Miguel. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Reevaluating Agricultural Productivity Gaps with Longitudinal Microdata
Joan Hamory Hicks, Marieke Kleemans, Nicholas Y. Li, and Edward Miguel
NBER Working Paper No. 23253
March 2017
JEL No. J43,O13,O15,R23

ABSTRACT

Recent research has pointed to large gaps in labor productivity between the agricultural and non-agricultural sectors in low-income countries, as well as between workers in rural and urban areas. Most estimates are based on national accounts or repeated cross-sections of micro-survey data, and as a result typically struggle to account for individual selection between sectors. This paper contributes to this literature using long-run individual-level panel data from two low-income countries (Indonesia and Kenya). Accounting for individual fixed effects leads to much smaller estimated productivity gains from moving into the non-agricultural sector (or urban areas), reducing estimated gaps by over 80 percent. Per capita consumption gaps between non-agricultural and agricultural sectors, as well as between urban and rural areas, are also close to zero once individual fixed effects are included. Estimated productivity gaps do not emerge up to five years after a move between sectors, nor are they larger in big cities. We evaluate whether these findings imply a re-assessment of the current conventional wisdom regarding sectoral gaps, discuss how to reconcile them with existing cross-sectional estimates, and consider implications for the desirability of sectoral reallocation of labor.

Joan Hamory Hicks
CEGA, Center for Effective Global Action
207 Giannini Hall,
University of California, Berkeley
Berkeley, CA 94720-3310
jrhamory@berkeley.edu

Marieke Kleemans
Department of Economics
University of Illinois
214 David Kinley Hall
1407 West Gregory Drive
Urbana, IL 61801
kleemans@illinois.edu

Nicholas Y. Li
Evans Hall, #3880
University of California
Berkeley, CA 94720-3880
nicholasli@econ.berkeley.edu

Edward Miguel
Department of Economics
University of California, Berkeley
530 Evans Hall #3880
Berkeley, CA 94720
and NBER
emiguel@econ.berkeley.edu

1. Introduction

The shift out of agriculture and into other more “modern” sectors (e.g., manufacturing) has long been viewed as central to economic development. This structural transformation was a focus of influential early scholarship (including Rosenstein-Rodan 1943; Lewis 1955; Rostow 1960; Pack 1972; Kuznets 1973; Johnston and Kilby 1978; Schultz 1988) with the issue even stretching back to Soviet debates over whether to “squeeze” farmer surplus to hasten industrialization (Preobrazhensky 1921).

A more recent macroeconomic empirical literature has revived interest in these issues, often using data from national accounts (Gollin, Parente, and Rogerson 2002; Caselli 2005; Restuccia, Yang, and Zhu 2008). This body of work has documented several important patterns that help shed light on the sources of income differences across countries. First, it shows that the share of labor in the agricultural sector correlates strongly with levels of per capita income: most workers in the poorest countries work in agriculture while only a small share do in wealthy countries. Importantly, while income per worker is only moderately larger (on average) for non-agricultural workers in wealthy countries relative to poor countries, agricultural workers are many times more productive in rich countries. This creates a double disadvantage for poor countries: agricultural work tends to be far less productive in low-income countries, yet the workforce is concentrated in this sector.¹ Studies that explore the closely related gap between the urban and rural sectors reach similar conclusions.

Several recent studies have examined the extent to which these productivity gaps across sectors can reasonably be viewed as *causal impacts* rather than mainly reflecting *worker selection*. By a causal impact of sector, we mean that a given worker employed in the non-agricultural (or urban) sector is

¹ See Appendix Figures A1 and A2, respectively, for an illustration of these differences using cross-country data.

more productive than the same worker employed in the agricultural (rural) sector. In contrast, worker selection would reflect differences driven by the fact that workers of varying ability and skill levels are concentrated in particular sectors. This paper seeks to disentangle these two competing explanations by estimating sectoral wage gaps using unusually long-run individual-level panel data from two low-income countries, Indonesia and Kenya.

If there are causal impacts of sector, the large share of the workforce employed in the agricultural sector in low-income countries could be viewed as a form of input misallocation along the lines of Hsieh and Klenow (2009) and Restuccia and Rogerson (2008). The resolution of this econometric identification issue, namely, distinguishing causal effects from selection, is not solely of scholarly interest: the existence of causal sectoral productivity gaps would imply that the movement of population out of rural agricultural jobs and into other sectors could durably raise living standards in low-income countries, narrowing cross-country differences. The existence of large causal sectoral productivity gaps also raises questions about the nature of the frictions that limit individual movement into more productive employment, and the public policies that might promote such moves (e.g., Tanzania's attempts to move rural populations into towns in the 1970s, Stren, Halfani, and Malombe 1994) or hinder them (e.g., China's *hukou* urban residential permit system, Au and Henderson 2006).

Gollin, Lagakos, and Waugh (2014, henceforth GLW) and Young (2013) are two important recent studies that explore this identification issue. GLW examine labor productivity gaps in non-agricultural employment versus agriculture using a combination of national accounts and repeated cross-sectional data from micro-surveys, and document a roughly three-fold average productivity gap across sectors. In their main contribution, GLW show that accounting for differences in hours worked and average worker schooling attainment across sectors—thus partially addressing worker selection—reduces the average estimated agricultural productivity gap by a third, from roughly 3 to 2. They also

find that agricultural productivity gaps and per capita consumption gaps based on household data remain large but tend to be somewhat smaller than those estimated using national labor surveys, possibly in part due to differences in how each source measures economic activity.

GLW remain agnostic regarding the causal interpretation of the large agricultural productivity gaps that they estimate. If individual schooling captures the most important dimensions of worker skill and thus largely addresses selection, GLW's estimates would imply that the causal impact of moving workers from agriculture to the non-agricultural sector in low-income countries would be to roughly double productivity, a large effect. Of course, to the extent that educational attainment alone fails to capture all aspects of individual human capital, controlling for it would not fully account for selection.

Young (2013) examines the related question of urban-rural differences in consumption (as proxied by measures of household asset ownership, education, and child health), rather than productivity, and similarly finds large cross-sectional gaps.² Young's interpretation differs from GLW in emphasizing the role of selective migration across sectors. Using Demographic and Health Surveys that have retrospective information on individual birth district, Young shows that rural-born individuals with more years of schooling than average in their sector are more likely to move to urban areas, while urban-born individuals with less schooling tend to move to rural areas. Young makes sense of this pattern through a model which assumes that there is more demand for skilled labor in

² While Young (2013) focuses on urban-rural gaps, he sometimes uses data on non-agricultural vs. agricultural differences when urban-rural data is missing; GLW similarly use urban-rural data when they lack data on agriculture.

urban areas, shows that this could generate two-way flows of the kind he documents, and argues that he can fully explain urban-rural consumption gaps once he accounts for sorting by education.³

The current study directly examines the issue of whether measured productivity gaps are causal or mainly driven by selection using long-term individual-level longitudinal (panel) data on worker productivity. Use of this data allows us to account for individual fixed effects, capturing all time-invariant dimensions of worker heterogeneity, not just educational attainment (as GLW do). We focus on two country cases – Indonesia and Kenya – that have long-term panel micro data sets with relatively large sample sizes, rich measures of earnings in both the formal and informal sector, and high rates of respondent tracking over time. The datasets, the Indonesia Family Life Survey and Kenya Life Panel Survey (henceforth “IFLS” and “KLPS”), are described in greater detail below.⁴

For both countries, we start by characterizing the nature of selective migration between non-agricultural versus agricultural economic sectors, and between urban versus rural residence. Like Young (2013), we show that individuals born in rural areas who attain more schooling are significantly more likely to migrate to urban areas and are also more likely to hold non-agricultural employment, while those born in urban areas with less schooling are more likely to move to rural areas and into agriculture. We exploit the unusual richness of our data, in particular, the existence of measures of cognitive ability (a Raven’s Progressive Matrices score), to show that those of higher ability in both

³ Porzio (2016) argues that a model of worker sorting can explain a large share (roughly 40%) of intersectoral productivity gaps, considering agriculture as well as a range of non-agricultural sectors. Lagakos and Waugh (2013) similarly model how worker sorting across sectors could generate sectoral productivity differences in equilibrium.

⁴ There are other panel data sets where similar approaches could be employed, for instance, the Mexican Family Life Survey; we leave this for future work. It is worth noting that Mexico is a member of the OECD and is considerably richer in per capita terms than Indonesia or Kenya. In related work, Alvarez (2015) finds substantial narrowing of productivity gaps in Brazil with the inclusion of individual fixed effects, albeit only using formal sector wage data, and Herrendorf and Schoellman (2016) employ cross-sectional microdata to assess sectoral differences in human capital.

Indonesia and Kenya are far more likely to move into urban and non-agricultural sectors, even conditional on educational attainment. This is a strong indication that conditioning on completed schooling is insufficient to fully capture differences in average worker skill levels across sectors.

We next estimate sectoral productivity differences, and show that treating the data as a repeated cross-section generates large estimated sectoral productivity gaps, echoing the results in existing work. In our main finding, we show that the inclusion of individual fixed effects reduces estimated sectoral productivity gaps by over 80 percent. This pattern is consistent with the bulk of the measured productivity gaps between sectors being driven by worker selection rather than causal impacts.

Specifically, we first reproduce the differences documented by GLW for Indonesia and Kenya, presenting both the unconditional gaps as well as adjusted gaps that account for worker labor hours and education (see Figure 1). These are large for both countries, with raw gaps of around 130 log points, implying roughly a doubling of productivity in the non-agricultural sector. When we treat our data as a series of repeated cross-sections, the gaps remain large, at 60 to 80 log points. These are somewhat smaller than GLW's main estimates, though recall that GLW's estimates using household survey data (like ours) also tend to be smaller. Conditioning on individual demographic characteristics (age and gender) as well as hours worked and educational attainment narrows the gap, but it remains large at between 30 and 60 log points. Finally, including individual fixed effects reduces the agricultural productivity gap in wages to 4.7 log points in Indonesia and to 13.4 log points in Kenya, and neither effect is statistically significant. Analogous estimates show that productivity gaps between urban and rural areas are also reduced substantially, to zero in Indonesia and 13.2 log points in Kenya. The estimated productivity gaps in GLW are an order of magnitude larger than our estimates.

We obtain similar results for the gap in per capita consumption levels across sectors where this is available for Indonesia. This is useful since consumption measures may better capture living standards in less developed economies than earnings measures, given widespread informal economic activity. Furthermore, we show that the productivity gap is not simply a short-run effect by demonstrating that gaps do not emerge even up to five years after an individual moves to an urban area. We also find that productivity gaps are no larger even when considering only moves to the largest cities in Indonesia and Kenya (the capitals of Jakarta and Nairobi, respectively).

Our methodological approach is related to Hendricks and Schoellman (2016), who use panel data on the earnings of international migrants to the United States, including on their home country earnings. Mirroring our main results, the inclusion of individual fixed effects in their case greatly reduces the return to international migration (by roughly 60 percent). Similarly, McKenzie et al. (2010) show that cross-sectional estimates of the returns to international immigration (to New Zealand) exceed those using individual panel data or those derived from a randomized lottery. Bryan et al. (2014) estimate positive gains in consumption (of roughly 30 percent) in the sending households of individuals randomly induced to migrate within Bangladesh, although no significant gains in total earnings. Bazzi et al. (2016) argue that cross-sectional estimates of productivity differences across rural areas within Indonesia are likely to overstate estimates derived from panel data using movers. Other related studies on the nature of selective migration include Chiquiar and Hanson (2005), Yang (2006), Beegle et al. (2011), Kleemans (2016), and Rubalcava et al (2008), among others.

A limitation of the current study is that we focus on two countries, in contrast to the scores of countries in GLW and Young (2013). This is due to the relative scarcity of long-run individual panel data sets in low-income countries that contain the rich measures necessary for our analysis. That said,

the finding of broadly similar patterns in both countries, each with large populations (250 million in Indonesia and 45 million in Kenya) in two different world regions, suggests some generalizability.

Another important issue relates to the local nature of our estimates, namely, the fact that the fixed effects estimates are derived from movers, those with productivity observations in both the non-agricultural and agricultural (or urban and rural) sectors. It is possible that productivity gains could be different among non-movers, an issue we discuss in Section 2 below. There we argue that, to the extent that typical Roy (1951) model conditions hold and those with the largest net benefits are more likely to move, selection will most likely produce an upward bias, leading our estimates to be upper bounds on the true causal impact of moving between sectors. However, absent additional knowledge about the correlation between individual preferences, credit constraints, and unobserved productivity shocks, it is in principle possible that selection could bias our estimates downward instead. Similarly, it is possible that very long-run and even inter-generational “exposure” to a sector could persistently change individual productivity due to skill acquisition, and this opens up the possibility that selection and causal impacts are both important. We return to these important issues of interpretation in the conclusion, including ways to reconcile our estimates with existing empirical findings.

The paper is organized as follows. Section 2 presents a conceptual framework for estimating sectoral productivity gaps, and relates it to the core econometric issue of disentangling causal impacts from worker selection. Section 3 describes the two datasets (IFLS and KLPS); characterizes the distinctions between the non-agricultural and agricultural sectors, and urban vs. rural areas; and presents evidence on individual selection between sectors. Section 4 contains the main empirical results on productivity gaps, as well as the dispersion of labor productivity across individuals by sector, consumption gaps, dynamic effects up to five years after migration, and effects in big cities versus other urban areas. The final section presents alternative interpretations of the results, and concludes.

2. Theoretical Framework

We present a development accounting framework to disentangle explanations for the aggregate productivity gap across sectors. We consider both observable and unobservable components of human capital, and whether intrinsic worker preferences for sector may bias direct measurement of the productivity gap. A standard model suggests that worker selection is most likely to bias sectoral productivity gaps upward when estimated among those moving into non-agriculture (from agriculture) but lead to a downward bias when estimated among those moving into agriculture.

2A. The Agricultural Productivity Gap through the Lens of an Aggregate Production Function

Following Hendricks and Schoellman (2016), we denote production in sector k as $Q_k = K_k^\alpha (A_k H_k L_k)^{1-\alpha}$. Dropping subscripts for convenience, a representative firm in sector k will solve:

$$\max_{K,HL} K^\alpha (AHL)^{1-\alpha} - R(1 + \tau^K)K - Z(1 + \tau^H)HL$$

where R and Z represent returns per unit of physical capital K and a labor aggregate (comprised of the product of human capital per unit of labor, H , and quantity of labor, L), respectively, and τ_K and τ_H represent wedges that prevent factors from receiving their marginal product.

Solving the first order condition with respect to the quantity of the labor aggregate yields:

$$Z = \frac{1 - \alpha}{1 + \tau^H} \left(\frac{K}{Q}\right)^{\alpha/1-\alpha} A$$

An individual's income in sector k is given by $Y_{ik} = Z_k H_{ik} L_{ik}$. Denoting logs in lower case, the average log-income gap across the non-agricultural (n) and agricultural (a) sectors is:

$$\bar{y}_n - \bar{y}_a = \underbrace{(z_n - z_a)}_{\text{residual income gap}=\beta} + \underbrace{(\bar{l}_n - \bar{l}_a)}_{\text{labor supply gap}} + \underbrace{(\bar{h}_n - \bar{h}_a)}_{\text{human capital gap}} \quad (1)$$

The agricultural productivity gap is comprised of a labor supply gap, a human capital gap, and a productivity residual, β , the key parameter of interest. The residual gap β captures not only wedges that directly prevent equalization of marginal products of labor between sectors but also wedges that may impact wages indirectly by causing misallocation of capital. These wedges are the focus of some models of structural transformation (Restuccia, Yang, and Zhu 2008; Graham and Temple 2006) and may serve as a summary parameter for an economy's degree of underdevelopment. In what follows, we do not take a stand on the specific components contributing to any sectoral productivity gap.

We assume that individual human capital takes the Mincerian form, $H_i = \exp[\mathbf{x}'_i \mathbf{b} + \eta_i]$ where \mathbf{x}_i is a vector of observed characteristics (e.g., years of schooling) with corresponding returns \mathbf{b} , and η_i represents unobserved skill. Substituting into the wage equation, log wages in sector k are:

$$y_i = z_a + \beta 1[[k = n]]_i + l_i + \mathbf{x}'_i \mathbf{b} + \eta_i \quad (2)$$

The agricultural productivity gap becomes:

$$\bar{y}_n - \bar{y}_a = \beta + (\bar{l}_n - \bar{l}_a) + (\bar{\mathbf{x}}_n - \bar{\mathbf{x}}_a)' \mathbf{b} + (\bar{\eta}_n - \bar{\eta}_a) \quad (3)$$

It is evident that any differences in unobserved components of worker human capital will be absorbed into the residual wage gap here, and an OLS estimate of β will be biased.⁵

⁵ This model can be generalized to allow for sector specific human capital with $h_{ik} = \exp[\mathbf{x}'_i \mathbf{b}_k + \eta_i]$ yielding an urban-rural gap described by $\bar{y}_n - \bar{y}_a = \beta + (\bar{l}_n - \bar{l}_a) + (\bar{\mathbf{x}}_n - \bar{\mathbf{x}}_a)' \mathbf{b}_a + (\bar{\mathbf{x}}_a)' (\mathbf{b}_n - \mathbf{b}_a) + (\bar{\eta}_n - \bar{\eta}_a)$ which motivates a Oaxaca-style decomposition where $\mathbf{b}_n - \mathbf{b}_a$ represent different returns paid to observable characteristics in non-agriculture. The main specifications below focus on human capital differences such as those described in equation 2 rather than this more flexible formulation of human capital. We also assume that the agricultural productivity gap does

There are two immediate approaches for obtaining better estimates of β . First, one can obtain a richer set of observable characteristics \mathbf{x}_i , reducing the scope for unobserved (to the econometrician) ability to determine income. Second, one can utilize panel data and estimate *within person* wage differences over time to purge the estimation of the time-invariant components of unobserved individual characteristics. While our estimation explores both avenues, our preferred estimates use the second approach, using fixed effects panel data estimation.

In a dynamic setting, the Mincerian human capital equation changes slightly to become: $H_{ikt} = \exp[\mathbf{x}'_i \mathbf{b} + \eta_i + \omega_{ikt}]$. Here, η_i is again unobserved individual skill, and ω_{ikt} is a mean zero, individual, sector-specific, time-varying shock. An individual's time-invariant human capital (which we estimate below as an individual fixed effect) is thus $\theta_i = \mathbf{x}'_i \mathbf{b} + \eta_i$. Equation 2 becomes:

$$y_{it} = z_a + \beta 1[[k = n]]_{it} + l_{it} + \theta_i + \omega_{it} \quad (4)$$

where $\omega_{it} = \omega_{iat} 1[[k = a]]_i + \omega_{int} 1[[k = n]]_i$ and the analogue of equation 3 is:

$$\overline{y_{nt}} - \overline{y_{at}} = \beta + (\overline{l_{nt}} - \overline{l_{at}}) + (\overline{\theta_n} - \overline{\theta_a}) + (\overline{\omega_{nt}} - \overline{\omega_{at}}) \quad (5)$$

Here, the time-varying, sector-specific components of human capital ω_{ikt} , are potential sources of omitted variable bias. Equation 4 is the key estimating equation; we explore potential limitations and pitfalls to this approach in what follows.

not change over time, which precludes time-varying frictions or production functions. This contrasts with longer-term views of development (e.g., Herrendorf, Rogerson, and Valentinyi 2014), but seems sensible given that the decadal time scale of our analysis is still dwarfed by that of historical economic development.

Estimating the agricultural productivity gap via equation 4 allows us to tentatively explore Lagakos and Waugh's (2013) hypothesis that comparative and absolute advantage in a Roy (1951) model of self-selection can explain sectoral productivity gaps within countries. To do so, we also allow for a richer formulation of the dynamic Mincerian human capital equation: $H_{ikt} = \exp[\theta_{ik} + \omega_{ikt}]$, where $\theta_{ik} = \mathbf{x}'_i \mathbf{b}_k + \eta_{ik}$ allows for different returns to observable human capital and unobserved ability by sector. Correspondingly, using a panel dataset containing multiple individual observations over time, we compute individual time-invariant human capital in each sector (θ_{ik}), and examine the joint distribution of these productivities, to explore whether those who have an absolute advantage in both sectors also tend to have a comparative advantage in the non-agricultural sector.

2B. Estimation Issues Related to Worker Selection into Sector

Departing from the general equilibrium model specified above, consider an agent i facing a choice of working in agriculture or non-agriculture. The utility v obtained by working in sector k is given by:

$$v_{ikt} = f(y_{it}, \mathbf{x}_i) + \xi_{ikt}$$

where ξ_{ikt} is an independent idiosyncratic preference shock for sector k in time t . For now, assume these preference shocks are uncorrelated with individual sectoral wage innovations, ω_{ikt} . We further assume that the non-stochastic component of utility is linearly separable, as $f(y_{it}, \mathbf{x}_i) = y_{it} + \mathbf{x}'_i \Gamma_k$.

Substituting in equation 4 for the individual specific productivity term, an individual chooses to work in the non-agriculture sector $k = n$ if and only if $v_{int} - v_{iat} > 0$; the probability of this occurring (abstracting away from labor supply differences across sectors for parsimony) is given by:

$$\begin{aligned} \Pr\{v_{int} - v_{iat} > 0\} &= \Pr\{\beta + (\omega_{int} - \omega_{iat}) \\ &\quad + (\Gamma_n - \Gamma_a)' \mathbf{x}_i + (\xi_{int} - \xi_{iat}) > 0\} \end{aligned} \tag{6}$$

The possible bias here is classic simultaneity: wage innovations ω_{ikt} simultaneously determine the worker's sectoral choice and her wage. In other words, receiving a positive productivity shock in non-agriculture ω_{int} is both positively correlated with the indicator for non-agricultural work and positively correlated with wages in that sector, while a positive productivity shock in agriculture ω_{iat} is negatively correlated with an indicator for non-agriculture and positively correlated with wages.⁶

The key threats to econometric identification in this panel data setting are time-varying shocks. The requirements for a convincing instrumental variable to remove selection bias in this context are thus relatively stringent. Such an instrument would ideally affect preferences for migration but be excludable from wage determination; this rules out using local rainfall shocks as an instrument, for instance. Both the IFLS and KLPS provide stated reasons for migration (subsequent to the move), but for these reasons to be used as instruments, the data would also need to provide reasons for staying, because not moving is also a choice. The dearth of credible natural experiments in migration makes the experimental variation in Bryan et al. (2014) and McKenzie et al. (2010) all the more valuable.

In a richer formulation of human capital with comparative advantage, the modified aggregate productivity gap in equation 5 (again abstracting away from labor supply differences) is:

$$\overline{y_{nt}} - \overline{y_{at}} = \beta + (\mathbf{E}\{\theta_{in} + \omega_{int} | v_{int} > v_{iat}\} - \mathbf{E}\{\theta_{ia} + \omega_{iat} | v_{iat} > v_{int}\}) \quad (7)$$

The selection expression in equation 6 suggests that we would only observe those employed in non-agriculture who would benefit from it, i.e., $\beta + \theta_{in} + \omega_{int} - \theta_{ia} - \omega_{iat} > c_{it}$, where c_{it} is the utility cost of moving for individual i in period t (captured by the other terms in equation 6). While this may

⁶ Explicitly, estimates of the agricultural productivity gap are biased if $\mathbf{E}\{\omega_{int} | v_{int} > v_{iat}\} - \mathbf{E}\{\omega_{iat} | v_{iat} > v_{int}\} \neq 0$.

be the average causal effect for this population—analogue to a local average treatment-on-the-treated in the program evaluation literature—extrapolating this effect to the non-movers may be problematic. This is especially relevant in fixed effects estimation, which estimates the productivity gap β using the wages of the movers, namely, those with productivity observations in both sectors.

Note that it is also possible that one might observe positive migration flows into non-agricultural employment even in the case where the true average productivity gap, β , was negative; in such a case, movers would consist of those with particularly large and positive individual returns to non-agricultural relative to agricultural employment in that time period, or perhaps those who face sufficiently large idiosyncratic preferences for the move, say, those with c_{it} negative.⁷

By this logic, fixed effects estimates (among the movers into non-agricultural work) will be generally larger than the average population treatment effect. This suggests that estimated gaps based on those who were initially in the agricultural (or rural) sector are likely to be *upper bounds* on the magnitude of the true average productivity gap in the population as a whole. Hendricks and Schoellman (2016) make a closely related point, arguing that their estimates of the returns to international migration are likely to be upper bounds.

In this study, this will likely be the case with the Kenya data (KLPS) where the entire sample lived in rural areas at baseline. In the Indonesia data (IFLS), which features sorting in both directions (since sample individuals were born in both urban and rural areas), it is in theory possible to observe

⁷ It is in theory possible to observe little or no migration when the true average productivity gap is positive. To generate this in the context of our model, the utility cost $\xi_{int} - \xi_{iat}$ must be both negatively correlated with the difference in time-varying sector specific wage innovations $\omega_{int} - \omega_{iat}$ and large enough in magnitude to dominate the β term. Note that such a correlation does not include highly productive individuals who simply prefer employment in agriculture, characteristics which would be captured in the fixed effects.

a non-agricultural premium every time an individual selects into non-agriculture and an agricultural premium every time an individual selects into agriculture. By a parallel logic to above, the selection equation in equation 6 suggests that among those initially working in the non-agricultural (urban) sector, we would only observe moves among those that benefit from working in agriculture, i.e., $-\beta + \theta_{ia} - \theta_{in} + \omega_{iat} - \omega_{int} > -c_i$. The resulting estimates would then serve as *lower bounds* on the magnitude of the true average productivity gain to non-agricultural employment.

The IFLS provides an ideal testbed to understand the role of these biases in estimating the related urban-rural gap. In the spirit of Young’s (2013) observation that migration flows in both directions, the data allow us to condition on individual birth location and measure the dynamic impacts on wages after migration. The bounding argument above predicts that the estimated urban-rural productivity gap would be larger when estimated for movers from rural to urban areas than it is when estimated for movers from urban to rural areas. We take this prediction to the data and find suggestive evidence for it. This model of selection implies that the true sectoral productivity gap in Indonesia is bounded by these two estimates, generated by movers in each direction.

3. Data

This paper uses detailed panel (longitudinal) data from Indonesia (IFLS) and Kenya (KLPS) to estimate worker productivity gaps between the non-agricultural and agricultural sectors, as well as the closely related question of gaps between workers in urban and rural areas. The data we use from both countries is unusually rich, and the long-term panel data structure features high rates of respondent tracking over time. At 250 million, the Southeast Asian country of Indonesia is the fourth most populous in the world, and Kenya is among the most populous Sub-Saharan African countries with approximately 45 million inhabitants. These countries are fairly typical of other low income countries

with respect to their labor shares in agriculture, estimated agricultural productivity gaps using national accounts data, and the relationships between these variables and national income levels.⁸

The high tracking rates of the datasets we employ allow us to construct multiyear panels of individuals' location decisions. Moreover, both datasets include employment information on both formal and informal sector employment. The latter is difficult to capture in standard administrative data sources yet often employs a large share of the labor force in low-income countries. If informal employment is more common in rural areas and in agriculture, and is partially missed in national accounts data, this may generate an upward bias in measured sectoral productivity gaps.

3A. Indonesia

Detailed individual and household-level data were collected in four rounds of the Indonesia Family Life Survey between 1993 and 2008 (Strauss et al., 2009). The survey is representative of 83 percent of the country's population who lived in 13 of the 27 provinces that existed in 1993. While the original sample consisted of 22,347 individuals, efforts to track them even when they had moved outside of the original study area, as well as the inclusion of members from split-off households during subsequent rounds (1997-98, 2000, 2007-08), results in a sample of 44,103 individuals. Attrition is often high in panel data; however, with an intensive focus on respondent tracking over space and time, the IFLS is unusually well-suited to study migration. In particular, re-contact rates between any two rounds are above 93 percent, and 90 percent of the original households were contacted in all four rounds (Strauss et al. 2009).⁹

⁸ See Appendix Figures A1 and A2 based on data from GLW. In both figures, the values for both Kenya (KEN) and Indonesia (IDN) are close to the best fitting regression line.

⁹ Thomas et al. (2012) contains a detailed discussion of tracking and attrition in the IFLS.

Detailed employment data were collected during each survey round. In addition to current employment, the survey included questions on previous employment, allowing us to create up to a 21-year annual employment panel at the individual level from 1988 to 2008. Employment status and sector of employment are available for each year, but in the fourth IFLS round, earnings were collected only for the current job. Therefore, the panel has annual data on employment status and sector of employment from 1988 to 2008, and earnings data annually from 1988 to 2000 and in 2007-08.

The IFLS includes information on the respondent's principal as well as secondary employment. Respondents are asked to include any type of employment, including wage employment, self-employment, temporary work, and unpaid family work.¹⁰ In addition to wages and profits, individuals are asked to estimate the value of their compensation in terms of share of harvest, meals provided, transportation allowance, housing and medical benefits, and credit; the main earnings measure we use is thus the comprehensive sum of all wages, profits, and benefits.¹¹

Individuals are asked to describe the sector of employment for each job. The single largest sector is "agriculture, forestry, fishing, and hunting": 34% of individuals report it as their primary employment sector, and 47% have secondary jobs in this sector. Agricultural employment is primarily rural: 42% versus 3% of rural and urban individuals, respectively, report working primarily in agriculture (Table 1, Panel A). Other common sectors are wholesale, retail, restaurants, and hotels (21% of main employment); social services (19%); manufacturing (15%); and construction (5%). These non-agricultural sectors are all more common in urban than rural areas. Men are more likely

¹⁰ Respondents are explicitly asked to include work on a family owned farm or non-farm business.

¹¹ De Mel, McKenzie, and Woodruff (2009) argue that self-reported profits give a more accurate depiction of firm profits in microenterprises than reconstructed measures.

than women to work in agriculture (38 vs. 27%) and less likely to work in wholesale, retail, restaurants, and hotels; and social services. Smaller male-dominated sectors include construction (7% of male employment vs. 0.8% for females) and transportation, storage, and communications (7 vs. 0.3%).

In the analysis that follows, we employ an indicator variable for non-agricultural employment, which equals 1 if a respondent's main employment is not in agriculture and 0 if main employment is in agriculture. The main analysis sample includes all individuals who are employed and have positive earnings and positive hours worked to ensure that the main variable of interest, the log wage, is defined. The sample includes 18,211 individuals and 115,897 individual-year observations.¹²

In addition to studying wage gaps, we explore consumption gaps to get a broader sense of welfare differences. IFLS consumption data were collected by directly asking households the value in Indonesian Rupiah of all food and non-food purchases and consumption in the last month, similar to consumption data collection in the World Bank's Living Standards Measurement Surveys.¹³ In contrast to the retrospective earnings data in the IFLS, the consumption data are all contemporaneous to the survey. Consumption data were collected at the household level, which we divide by the number of household members to obtain a per capita measure. The consumption sample includes 38,280 individual-year observations from 19,695 individuals in IFLS rounds 1–4. In the consumption analysis, we expand the sample to also include individuals without current earnings data; we also perform a robustness check on the consumption analysis using the main productivity sample.

¹² The panel is unbalanced due to attrition, death, and to limiting observations to respondents at least 16 years old.

¹³ Note that for a small number of frequently-consumed items, information was collected for the last week, and for a few low-frequency items, data was collected for the last year.

Data were collected on the respondent's location at the time of the survey, and all rounds of the IFLS also collected a full history of migration within Indonesia. All residential moves across sub-districts ("kecamatan") that lasted at least six months are included. Figure 2, Panel A presents a map of Indonesia with each dot representing an IFLS respondent's residential location. While many respondents live on Java, we observe considerable geographic coverage throughout the country. The IFLS also asked respondents for the main motivation of each move. Family-related reasons are most common at 46%, especially for women (50%), who are more likely than men to state they migrated for marriage. The second most common reason to migrate is for work (34%), with little difference by gender, while migrating for education is less common. We combine data across IFLS rounds to construct a 21-year panel, from 1988 to 2008 with annual information on the person's location, in line with the employment panel; refer to Kleemans (2016) and Kleemans and Magruder (2016) for more information on the construction of the IFLS employment and migration panel.

We utilize a survey-based measure of urban residence: if the respondent reports living in a "village", we define the area to be rural, while they are considered urban if they answer "town" or "city." We present the correspondence between urban residence and employment in the non-agricultural sector in Table 1, Panel A. In 66 percent of individual-year observations, people are employed in the non-agricultural sector, and in 21 percent of the observations, they live in urban areas. One can see that a substantial portion of rural employment is in both agriculture and non-agricultural work, while urban employment is almost exclusively non-agricultural, as expected.

Given the migration focus of the analysis, it is useful to report descriptive statistics both for the main analysis sample, as well as separately for individuals in four mutually exclusive categories (Table 2, Panel A): those who always reside in rural areas throughout the IFLS sample period ("Always Rural"), those who were born in a rural area but move to an urban area at some point ("Rural-

to-Urban Migrants”), those who are “Always Urban,” and finally, the “Urban-to-Rural Migrants” (born urban but move to a rural area at some point). As discussed above, the fixed effects analysis is driven by individuals who move between sectors during the sample period.

In the main IFLS analysis sample, 80 percent of adults had completed at least primary education, and a quarter had completed secondary education, while tertiary education remain quite limited, at less than 10 percent. Among those who are born in rural areas in columns 2 and 3 (of Table 2, Panel A), we see that migrants to urban areas are highly positively selected in terms of both educational attainment, and in terms of cognitive ability, with Raven’s Progressive Matrices exam scores roughly 0.2 standard deviation units higher among those who migrate to urban areas, a substantial effect.¹⁴ Migration rates do not differ substantially by gender.

These relationships are presented in a regression framework in Table 3, Panel A (columns 1 to 5), and the analogous relationships for moves into non-agricultural employment are also evident (Table 4, Panel A). Importantly, the relationship between higher cognitive ability and likelihood of migrating to urban areas holds even conditional on schooling attainment and demographic characteristics (column 6 of both tables), at 99% confidence. This indicates that sorting on difficult-to-observe characteristics is relevant in understanding sectoral productivity differences in this context.

It is worth noting that if we ignore migrants, individuals who are born and remain in urban areas are far more skilled than those who stay in rural areas. “Always Urban” individuals score over 0.4 standard deviation units higher on Raven’s matrices and have triple the rate of secondary schooling and six times the rate of tertiary education relative to “Always Rural” individuals. The urban-to-rural

¹⁴ Raven’s Matrices were administered to a subset of individuals in IFLS 3 and 4, namely those 7 to 24 years old. The Raven’s Matrices test is designed to capture fluid intelligence.

migrants in Indonesia are also negatively selected relative to those who remain urban residents, which corroborates Young's (2013) claim. These patterns emerge in Table 2, Panel A, where the urban-to-rural migrants score lower on all skill dimensions relative to those who remain urban; appendix Tables A1 and A2 report results analogous to Tables 3 and 4, among those individuals born in urban areas.

3B. Kenya

The Kenya Life Panel Survey (KLPS) includes information on 8,999 individuals who attended primary school in western Kenya in the late 1990s and early 2000s, following them through adolescence and into adulthood. These individuals are a representative subset of participants in two primary school-based randomized interventions: a scholarship program for upper primary school girls that took place in 2001 and 2002 (Kremer, Miguel, and Thornton 2009) and a deworming treatment program for primary school students during 1998–2002 (Miguel and Kremer 2004). In particular, the KLPS sample contains information on individuals enrolled in over 200 rural primary schools in Busia district at the time of these programs' launch. According to the 1998 Kenya Demographic and Health Survey, 85% of children in Western Province aged 6–15 were enrolled in school at that time, and Lee et al. (2015) show that this area is quite representative of rural Kenya as a whole in terms of socioeconomic characteristics. To date, three rounds of the KLPS have been collected (2003–05, 2007–09, 2011–14).

KLPS data collection was designed with attention to minimizing bias related to survey attrition. Sample individuals who had left the original study area were tracked throughout Kenya (as well as into neighboring Uganda and beyond, although we exclude international migrants from the present analysis).¹⁵ Respondents were sought in two separate “phases” of data collection: the “regular tracking

¹⁵ The results presented below are robust to the inclusion of international migrants (not shown).

phase” proceeded until over 60 percent of respondents had been surveyed, at which point a representative subset of approximately 25 percent of the remaining sample was chosen for the “intensive tracking phase” (and remaining unfound individuals no longer sought). These “intensive” individuals receive roughly four times as much weight in the analysis, to maintain representativeness with the original sample. The effective tracking rate for each KLPS round is around 85 percent.¹⁶

Similar to the IFLS, the KLPS includes detailed information on educational attainment, labor market participation, and migration choices. Employment data was collected in a wage employment module and a self-employment module, which both are designed to include both formal and informal employment. Most individuals were quite young (typically teenagers) during data collection for KLPS round 1, and few had wage employment or self-employment to report. Full employment histories, including more detailed questions, were collected during rounds 2 and 3, and it is from these rounds that we draw the data on individual earnings, hours worked, and wages used in the present analysis.

The Kenya agricultural productivity data deserves detailed discussion. Whenever total household annual agricultural sales were sufficiently large, exceeding 40,000 Kenyan Shillings (approximately 400-500 US dollars), full agricultural production and profit information was collected in the self-employment module and included in the present analysis. Agricultural wage employment is also common, and these data are always included. Limited questions on subsistence agricultural production were collected in KLPS rounds 1 and 2, but these are insufficient to create an individual productivity measure. More detailed information on agricultural productivity (in the previous 12 months) is contained in round 3, and this is included in the present analysis. To create a measure of

¹⁶ Baird et al. (2008) contains a detailed explanation and calculation of the effective tracking rate.

individual productivity comparable with other sectors, we focus on agricultural activities (e.g., growing a particular crop) in which the respondent provided all reported labor hours; we also restrict attention to activities in which the respondent reports being the main decision-maker, since it seems likely that they are most knowledgeable about such activities (although results are not sensitive to this restriction, not shown). The profit in an agricultural activity is the sum of all crop-specific production – valued either through actual sales or at the relevant crop price (collected in regular local market price surveys) if consumed directly – minus all input costs and hired labor costs. The individual wage divides this net profit by the labor hours the respondent supplied to the activity.

KLPS respondents reported industry for all wage and self-employment. Most individuals are engaged in relatively low-skilled work. The most common industry for wage employment is services, at 58% overall and 74% for females (with many women employed in domestic services). In rural areas, the most common industries for wage employment are services and agriculture (58 and 11%, respectively), while in urban areas they are services, and manufacturing and construction (63 and 11%). The largest self-employment industries are retail and services (38 and 31%).¹⁷

KLPS round 3 collected detailed consumption expenditure data for a subset of individuals. However, because it was only collected for this round, we are unable to utilize it in panel estimation. Instead, in the panel analysis we utilize a proxy for consumption, the number of meals eaten in the previous day, which is available in both KLPS rounds 2 and 3. Reassuringly, meals eaten is strongly

¹⁷ For wage employment, respondents also report occupation, and these tell a similar story regarding skill level. The most common occupations fall in the “unskilled trades” category (31%), followed by “skilled and semi-skilled trades” (20%), “retail and commercial” (19%), “professionals” (17%), and “agriculture” (13%). Agricultural wage employment is more common for men than for women (17 compared to 6%, respectively), and as expected, agricultural employment is far higher in rural than urban areas (28 vs. 3%). Common urban occupations are “unskilled trades” (37%), “skilled and semi-skilled trades” (23%) and “retail and commercial” (21%).

correlated with our primary measures of labor productivity as well as consumption expenditures per capita (in KLPS-3); see Appendix Table A3. As with Indonesia, in the meal consumption analysis, we are able to expand the sample to also include individuals without current earnings data.

KLPS respondents provide a history of residential locations since their last interview, and this data includes residential district, town, and village, allowing us to classify individuals who lived in towns and cities as urban residents. The KLPS includes information on all residential moves that lasted at least four months in duration, a slightly more permissive definition than in the IFLS, and we are able to construct a monthly residential panel from March 1998 to October 2014.¹⁸ Combined with the retrospective labor productivity data, the main analysis sample is a monthly panel with 128,439 individual-month observations for 4,537 individuals.

Figure 2, Panel B presents a map of Kenya, with each dot representing a respondent residential location during 1998–2014. Most residences in western Kenya are located in Busia district (where the sample respondents originally resided), with substantial migration to neighboring areas as well as to cities. Appendix Table A4 presents the list of main towns and cities, and shows that 70 percent of urban residential moves are to Kenya’s five largest cities, namely, Nairobi, Mombasa, Kisumu, Nakuru, and Eldoret. According to survey reports, men are slightly more likely than women to migrate for employment reasons (61% of moves compared to 58% for females) while women are more likely to migrate for family reasons, including marriage (12% vs. 1% for men). A smaller share of moves (approximately 6%) are for education.

¹⁸ Similar to the IFLS, the panel is unbalanced due to attrition, death, and inclusion of individuals 16 and older.

Summary statistics on employment sector and urban residence for KLPS respondents are presented in Table 1, Panels B and C. Panel B presents data for the main analysis sample; as described above, this contains subsistence agricultural information where available (from KLPS-3). The employment shares in agriculture is much higher in rural areas (25.9%) than urban (5.3%), as expected, but the share in rural areas is somewhat lower than expected, likely because subsistence agricultural activities were not captured in earlier KLPS rounds. For a more complete portrait, Panel C focuses on data from the 12 months prior to the KLPS-3 survey, which contains detailed information on subsistence agriculture, and here the agricultural employment share in rural areas is much higher.

Recall that the Kenya sample is all rural at baseline (they were originally attending rural schools). Similar patterns emerge regarding positive selection into urban migration, with educational attainment and normalized Raven's matrix scores both far higher among those who migrate to cities (Table 2, Panel B). In particular, there is a raw gap of nearly 0.3 standard deviation units in Raven's matrix scores between urban migrants and those who remain rural. Overall migration rates in Kenya are similar for females and males. Tables 3 and 4 (Panel B) report these patterns in terms of regression estimates, for urban migration and employment in non-agricultural work, respectively. As with Indonesia, controlling for educational attainment and gender, the Raven's score is strongly positively correlated with urban migration (at 95% confidence).

4. Results

4A. Main Agricultural and Urban Productivity Gap Estimates

GLW estimate raw and adjusted agricultural productivity gaps of 138 and 108 log points in Indonesia, respectively (Figure 1, Panel A). The estimate of this raw gap from the IFLS is somewhat smaller at 62 log points (Table 5, Panel A). The most straightforward explanation for this discrepancy is an issue

of measurement. GLW observe that, in an analysis of 10 countries, the average agricultural productivity gap was 17 log points smaller when estimated in Living Standards Measurement Study (LSMS) data that is similar to the IFLS, and which is more likely to capture earnings in informal employment.¹⁹ That said, the raw gap we estimate in the IFLS remains substantial.

Inclusion of control variables similar to those used by GLW to adjust macro data gaps reduces the estimated agricultural productivity gap in the IFLS to 51 and 32 log points (in Table 5, columns 2 and 3). Estimating on the subsample for which we have scores from Raven’s matrix tests, the gap is reduced slightly, although note the smaller sample size in this case.

Limiting the analysis to those who have productivity measurements at some point in time in both agricultural and non-agricultural employment, the productivity gap drops to 16 log points (column 5), suggesting that the selection on unobservable characteristics alluded to in Section 2 may play a meaningful role. Inclusion of fixed effects reduces the gap further (column 6), and using our preferred labor productivity measure, the log wage (namely, the log of total earnings divided by hours worked), as the dependent variable nearly eliminates the gap altogether: the coefficient estimate falls to 0.047 (standard error 0.031) in column 7, and further to 0.045 when considering the real log wage (adjusting for higher urban prices, column 8).

We follow a similar approach for Kenya, where the raw agricultural productivity gap falls from 79 log points to 56 with the inclusion of GLW’s controls (Table 5, Panel B, columns 1–4), and to 32.6 log points when including an individual fixed effect. Using the preferred hourly wage measure reduces

¹⁹ This comes from log transformed values from the “Average” row of GLW, Table 4, i.e. $\ln 2.6 - \ln 2.2 = 0.167$.

the gap to 13.4 log points (column 7), it falls further when adjusted with an urban price deflator (column 8), and neither fixed effects wage estimate is significant at traditional levels of confidence.

Comparing column 1 (the raw gap) to column 7 (the preferred fixed effects estimate) in Table 5, the agricultural productivity gap is reduced by 92 percent in Indonesia and by 83 percent in Kenya. The standard errors are somewhat larger for Kenya, so the upper end of the 95% confidence interval includes a sizable gap of 37 log points, consistent with some non-trivial productivity gains to non-agricultural employment. That said, even this value remains far lower than the 108 and 71 log point effects that GLW estimate for Indonesia and Kenya, respectively, once they condition on observable labor characteristics (namely, hours worked and educational attainment). As noted in the introduction, these results for Indonesia and Kenya are presented graphically in Figure 1, Panels A and B and compared to GLW's estimated productivity gaps.²⁰

Table 6 presents the closely related exercise of estimating the labor productivity gap between residents of urban and rural areas. While the existing empirical literature has sometimes conflated these two gaps, Table 1 shows that employment in rural areas is not exclusively characterized by agriculture. To the extent that residential migration is costlier than shifting jobs (but not homes), and the urban and non-agricultural wage premia are related but distinct parameters, one might suspect that an urban wage premium might even be more pronounced than the non-agricultural wage premium.

The microdata estimates from Indonesia and Kenya appear to be consistent with this view, at least at first glance: the raw gap reported in column 1 of Table 6 (Panels A and B) are 63 and 85 log points for Indonesia and Kenya, respectively. Similar to the agricultural productivity gap, the urban-

²⁰ Similar patterns are obtained when using alternative definitions of non-agricultural employment, namely, classifying simultaneous work in both sectors as agriculture or as non-agriculture; see Appendix Table A5.

rural productivity gap falls when additional explanatory variables are added in columns 2, 3 and 4, but remains substantial and statistically significant. Focusing the analysis only on those who have earnings measures in both urban and rural areas (column 5) leads to a further reduction. Finally, the urban-rural earnings gap falls to 1.8 log points with the inclusion of individual fixed effects in Indonesia, and -0.7 log point for the preferred log wage measure (column 7). The analogous urban productivity effect estimate for Kenya is slightly larger at 13.2 log points (column 7). Thus, the productivity gap in Indonesia falls by 100 percent in Indonesia (to zero), and the reduction for Kenya is 84 percent (from 85.1 to 13.2 log points, across columns 1 and 7) with the inclusion of individual fixed effects. Once again, these results are summarized in Figure 1 (Panels C and D).²¹ Urban productivity gaps in real wage terms (that account for higher urban prices) are further reduced in both countries (column 8).

The selection model (in Section 2) predicts that estimated productivity gaps would be higher among rural-to-urban migrants than for urban-to-rural migrants, given plausible patterns of selection bias. Table 7 explores this hypothesis in Indonesia by separately conditioning on birth location; Panel A contains those born rural and Panel B those born in urban areas. The same pattern of declining productivity gaps in each subsample is observed for non-agriculture (first four columns) and urban (columns 5-8) as additional controls are included. In the preferred log wage specifications in columns 4 and 8, productivity gaps are indeed somewhat larger for those born in rural areas, as predicted by the sorting model. The difference between estimates for those born in rural versus urban areas is small,

²¹ Appendix Table A6 further explores the relationship between these two gaps by conditioning on observations in rural areas. The raw agricultural productivity gap and subsequent decline with the inclusion of controls and fixed effects is quite similar in Indonesia in rural areas (panel A), and even more pronounced in Kenya. In Appendix Table A7, we present estimates for Indonesia on a sample of individuals who are at most 30 years old, for greater comparability with the Kenya sample, which consists of young adults; the estimates remain similar. Appendix Tables A8 and A9 report results separately for wage earnings and self-employment earnings, respectively, and generate similar results. Appendix Table A10 reports results for Kenya including subsistence agriculture even when the respondent is not the main decision maker for an activity (Panel A), and excluding subsistence agriculture entirely (Panel B), and results remain robust.

suggesting rather tight bounds. For instance, the estimated productivity gain to non-agricultural employment is 5.9 log points for those born in rural areas (Panel A, column 4) and -0.8 for those born urban (Panel B). While suggestive, note that the difference between these estimates is not significant.

4B. Productivity versus Living Standards

The discussion above establishes at least an 80 percent reduction in estimated sectoral productivity gaps once individual fixed effects are included in the analysis (Figure 1). The wage measures presented thus far are closely related to the labor productivity parameters that are the focus of most existing macroeconomic empirical literature. However, productivity and “utility” may diverge for many reasons, including price differences across regions, amenities, unemployment, and other factors. For instance, there could be considerable individual heterogeneity in the taste for rural versus urban amenities, e.g., comforts of home, ethnic homogeneity, better informal insurance, etc., in rural areas versus cosmopolitan cities’ better public goods and more novelty (but downsides too, such as crime).

Although it is impossible to fully capture these factors and convincingly measure individual welfare, to get somewhat closer to differences in living standards, we draw on consumption data from the IFLS. As described in Section 3, four rounds of the IFLS included questions on the value of household consumption which can be converted to per capita consumption. In the main specification, we include all individuals who have such consumption data, even if they lack earnings measures.

The initial consumption gap between non-agriculture and agriculture is large and similar the productivity gap at 54 log points (Table 8, Panel A). The gap falls considerably when including time fixed effects and control variables in column 2, and falls to only 3.0 log points when also including individual fixed effects in column 3, an effect that is not significant. A similar pattern is presented for the urban-rural consumption gap in columns 4, 5, and 6: the gap declines from 53 log points to 3.0 log

points (not significant). Note that prices may be higher in urban areas and we do not yet adjust for such differences, though a price adjustment in urban areas would presumably only lead this estimated gap to be smaller (or even slightly negative), strengthening the finding.²²

The consumption proxy measure in the KLPS tells a similar story. The raw gap in meals eaten in Kenya between those in non-agriculture versus agricultural employment is positive and statistically significant, though smaller than the earnings gap (Table 8, Panel B); differences in magnitude are difficult to interpret given the different nature of the meals measure. Mirroring the broad pattern observed for labor productivity measures, this gap falls by almost half when including controls, and is negative though not statistically significant when including individual fixed effects (columns 1-3); a similar pattern holds for the urban-rural gap (columns 4-6).

Another dimension of welfare relates to patterns of unemployment. Appendix Table A13 explores whether there are differences in unemployment rates and search behavior between urban and rural areas for Kenya, where this data is available. We find that unemployment (measured several ways) is either similar in urban and rural areas (Panel A, column 3) or somewhat higher in urban than in rural areas conditional on individual fixed effects (Panel A, column 6, and Panel B), strengthening the main finding that movers to urban areas do not experience large gains in total earnings.

Finally, up to this point we have only considered *mean differences* in productivity or consumption across sectors, but *variability of outcomes* could also be a determinant of individual wellbeing, as well as of migration choices. We test whether the variability of earnings in the

²² Appendix Table A11 shows the gap in both food and non-food consumption (Panels A and B, respectively). The gaps in both components of consumption see reductions of over 90 percent when including individual fixed effects. Appendix Table A12 repeats the consumption analyses on the main analysis sample (i.e., those with earnings data) for total consumption (Panel A) and by food and non-food consumption (Panels B and C, respectively), and results are similar.

agricultural sector is different than variability in the non-agricultural sector (conditional on individual fixed effects) and find no statistically significant differences in either Indonesia or Kenya (not shown).

4C. Sector-specific Productivity — Absolute and Comparative Advantage

In the conceptual framework, the richest model of human capital allowed for individual sector-specific productivity θ_{ik} . Analysis of these productivities has been given renewed focus in Lagakos and Waugh (2013), who argue that self-selection on the basis of comparative advantage could play an important role. In their model, comparative advantage is positively correlated with absolute advantage, meaning that the most productive workers have the most to gain from selecting into non-agriculture.

Utilizing panel data, we estimate a modified version of equation 4 replacing the individual fixed effect with an individual-sector fixed effect.²³ We recover these estimates, and then normalize the mean of the fixed effects of permanent rural residents (non-movers) to be zero. Figure 3 presents the joint distributions of these estimated individual productivities by sector. Panel A includes Indonesians born in rural areas. It is apparent that rural-to-urban migrants are positively selected relative to non-migrants, with an average rural wage approximately 18 log points higher than non-migrants. These individuals experience only a 1 log point average increase in their wage upon migration to an urban area. Panel B presents the same exercise with Indonesians born in urban areas. Here, there appears to be negative selection into rural migration, with the average mover having 10 log points lower wages when still in urban areas, and an increase of only 3 log points in rural wages among moving. Panel C presents results in Kenya (all of whom were rural residents as children) that

²³ This procedure is similar in spirit to the correlated random coefficient models utilized to analyze heterogeneous returns to hybrid seed adoption (Suri 2011) and labor unions' effects on wages (Card 1996, Lemieux 1998), although our approach makes fewer assumptions and is meant to be more descriptive.

are analogous to panel A. Compared to Indonesia, there appears to be even more positive selection among urban migrants in Kenya (at 42 log points) as well as a moderate positive urban premium of roughly 17 log points, which is slightly larger than the regression adjusted estimate presented above.

Note that the realizations of roughly half of migrants fall below the 45 degree line in the three panels of Figure 3, which taken literally means that they experience higher earnings in rural than urban areas. This is consistent with the empirical finding of zero or small positive sectoral productivity gaps.

This exercise is meant to be descriptive, and we interpret the relationships between the estimated individual urban and rural productivities with caution here, in part because the estimates are subject to measurement error and thus the fitted regression line may experience attenuation bias. With these caveats in mind, note that all three of these plots appear to show that absolute advantage plays a role in wage determination, with positive and similar slopes across settings and different subsets of individuals. Slopes are less than 1 in the three cases, which suggests that in a relative sense, those with absolute advantage in productivity may be gaining somewhat less from migrating (in proportional terms) than those who remain in the birth region, although once again this remains highly speculative.

4D. Dynamics of the Productivity Gap and Big City Effects

In unpacking the main result, we examine if dynamics and experience effects produce productivity gains that do not materialize right away. In particular, while the main specification includes time fixed effects which would account for overall growth of wages as the sample ages or year specific shocks, individuals may begin to earn more after spending time in urban areas. Figure 4 presents event study analyses of whether individuals earn more after migrating, where we estimate regressions of the form:

$$\begin{aligned}
y_{it} = & \theta_i + \delta_t + \mathbf{X}'_{it}\mathbf{b} + \sum_{\tau=-5}^{-2} \beta_{\tau}1[[k = U]]_{i,t+\tau} + \sum_{\tau=0}^5 \beta_{\tau}1[[k = U]]_{i,t+\tau} + \gamma_{pre} \sum_{\tau \leq -6} 1[[k = U]]_{i,t+\tau} \\
& + \gamma_{post} \sum_{\tau \geq 6} 1[[k = U]]_{i,t+\tau} + \varepsilon_{it} \tag{8}
\end{aligned}$$

These regressions are estimated on an unbalanced panel of individual-time periods and include individual fixed effects θ_i , time fixed effects δ_t , squared age as a time-varying covariate \mathbf{X}_{it} , and additional indicator variables for time periods exceeding five years pre- and post-move, γ_{pre} and γ_{post} , respectively. The Indonesian and Kenyan analyses both condition on individuals being born rural.

The β_{τ} parameters of primary interest are coefficients on indicators for time periods relative to the individuals' move to an urban area at $\tau = 0$. Estimates are relative to the year or month prior to the individuals' move in Indonesia and Kenya, respectively; we exclude an indicator for the period immediately prior to the individuals' move. These coefficients are identified by individuals who have adjacent productivity measures in both the period that they move to urban and the period immediately prior. We do not enforce a requirement that individuals are observed in every period five years prior- and post-move. If the extensive margin decision to exit the labor force entirely or attrit from the sample is correlated with urban labor market experiences, the results may be biased and we thus interpret them with caution. Nonetheless, the richness of the panel dataset is novel and worth exploring.

These parameters represent the difference in mean wages between movers and non-movers net of the difference that existed in the period prior to the urban move. An advantage of this approach is that it also allows us to assess wage dynamics *prior* to the move, which may give some clues about what precipitated the move – e.g., whether rural individuals are more likely to move following a negative earnings shock – and they also allow us to examine whether urban experience leads to gradually rising earnings there.

In Indonesia, urban wages do not change substantially relative to the year prior to moving, and even five years after the urban move, migrants see no average wage gain (Figure 4, Panel A). There are broadly similar results in Kenya relative to the month prior to the move; there is some suggestive indication of slightly rising wages in the first two years of residence in an urban area, but these are small (Panel B). There is no indication of meaningful pre-move trends in either country.

In this analysis, we consider wages for individuals who made an urban move regardless of whether they remained in cities or towns, or later moved back to rural areas. The bottom halves of both panels A and B show a “survival” rate in urban areas of between 50 to 60% after five years (in both countries), suggesting substantial return migration to rural areas. Naturally, one might suspect that those with the worst economic outcomes in urban areas might return home, yet this does not appear to be the case: Appendix Figure A3 separately plots post-move wages for those who remain in urban areas and those who return to rural areas, and we find no evidence of a significant divergence in earnings between these two groups (although note that there is limited statistical power to distinguish between them). This suggests a direction for future research in uncovering the reasons for these moves, including whether non-economic factors, including family reasons and heterogeneity in the taste for urban living, are often decisive factors.²⁴

Other scholars have argued that job experience is particularly valuable in big cities and that residence in these cities may boost individual productivity over time (see de la Roca and Puga 2016 for the Spanish case). We examine this issue, first repeating the main urban productivity gap analysis (from Table 6) but including a breakdown into the five highest population cities in each country, in

²⁴ We carry out an analogous event study of moves to rural areas among those born urban in Indonesia, and similarly find no evidence of significant dynamic impacts (see Appendix Figure A4).

Table 9. In Indonesia, all five cities are larger than 2 million inhabitants, with the capital Jakarta at 10 million. Kenya's capital Nairobi has 3.4 million people, the second largest city (Mombasa) has nearly one million, while the other three cities in Kenya are smaller. The capitals are also the largest destinations for urban migrants in each country.

The main finding is that there is no evidence for significantly larger effects in the largest cities of Jakarta and Nairobi (column 4 of Panels A and B, respectively). There is some evidence of significant positive urban productivity gains in certain cities (albeit not the largest), namely, Bandung in Indonesia and Mombasa in Kenya, for which we have found no immediate explanation; we leave an exploration of these differences between particular cities for future research.

While this analysis does not find evidence of an overall big city effect, we also assess whether large city productivity effects might manifest over a longer time horizon by repeating the event study analysis over a five year time horizon separately for Jakarta and Nairobi (Appendix Figure A6). These figures show no clear evidence of positive dynamic effects in these large capital cities, and if anything moving to Jakarta and Nairobi appears to lead to somewhat more negative wage realizations over time, although differences with other cities are imprecisely estimated and generally not significant. The bottom line is that there is no clear evidence for larger productivity gains for movers to the biggest cities in either Indonesia or Kenya, either immediately or over a five year time horizon.

5. Conclusion

Several influential recent studies document large sectoral productivity gaps in low-income countries and highlight an apparent puzzle, namely, “why so many workers remain in the agricultural sector, given the large residual productivity gaps with the rest of the economy” (Gollin, Lagakos, and Waugh 2014, p. 941). This study makes two main contributions using data from low-income countries with

large populations (Indonesia and Kenya) located in two different regions. First, we show that estimating sectoral productivity gaps—both across non-agricultural and agricultural sectors, and across urban and rural areas—using panel data and including individual fixed effects leads to a reduction of over 80 percent in the estimated gaps. The second main empirical contribution lies in demonstrating that there is extensive individual selection across sectors, both along relatively easily observable dimensions such as educational attainment as well as measures of skill (here, a measure of cognitive ability) that most standard economic datasets lack.

Taken together, the findings point to the importance of individual selection in driving observed sectoral gaps in productivity and living standards, and call into question strong causal interpretations. As a result, the puzzle of why the share of workers in rural agriculture remains high may not be as much of a puzzle as previously thought. Similarly, if gaps are mainly driven by selection, then policies to incentivize workers to move to urban areas (and out of agriculture), based on the logic of input misallocation, would not appreciably raise aggregate living standards and would not appear to be an appropriate policy direction.

An historical episode illustrates some of the potential risks of pro-urbanization policies. In the 1970s, Tanzania's authoritarian socialist government sought to move its rural population into larger villages and towns to speed up economic modernization. The underlying idea was that the provision of public services and the shift into non-agricultural work (including manufacturing) would be hastened if households would only leave their traditional homesteads, which were often highly spatially dispersed. After initial rhetorical encouragement and incentives by the government led to few moves, the government resorted to forced migration in certain regions in 1973, in the so-called "Operation Vijiji". The resulting economic and social dislocation is today widely viewed as a policy disaster within Tanzania (Stren, Halfani, and Malombe 1994). While one could argue that observers

are unable to assess the true economic effects of the policy in Tanzania, since the forced moves were quickly abandoned (within a year) in the face of large-scale resistance, at a minimum, the Tanzanian case indicates that it can sometimes be very costly from a welfare perspective to rapidly induce a large share of the population to move out of traditional rural agriculture.

As noted above, our main productivity gap estimates are derived from individual movers, namely, those with productivity measured in both sectors. Thus a logical way to reconcile our finding of small or even zero sectoral gaps with the existing macroeconomic empirical evidence of large average gaps is the possibility that productivity effects among non-movers would be much larger than those of movers. Given the nature of our data, it is impossible to rule out this possibility, and it clearly merits further investigation, although the lack of measured individual productivity in both sectors for non-movers naturally complicates the rigorous identification of these relationships.

However, several factors lean against this interpretation in our view, at least in the short-run. First, it is natural to think of the migration decision in terms of a Roy (1951) model, as we do above, in which those with the largest net utility benefits are most likely to move. This could lead our estimates to overstate gaps between sectors overall. While it is possible that those individuals who remain in the rural agricultural sector might receive large positive earnings gains from moving, their choice not to do so might simply reflect high financial or non-financial costs to migration. For instance, the bundle of amenities found in a large city is quite different than those in rural areas, and individuals may have strong and heterogeneous preferences for them, leading to large reductions in utility for some migrants even if wages rise. Poor individuals may also face credit constraints or other financial frictions that prevent them from exploiting wage gaps, and easing these constraints could boost migration rates, as argued for Indonesia by Bazzi (2016) and India by Munshi and Rosenzweig (2016).

A promising approach to estimating the returns to migration in low-income countries among those who are typically “non-movers” and may face such constraints is the Bryan et al. (2014) study in Bangladesh. They find that a moderate subsidy did induce a moderate share of recipients (roughly one fifth) to move to towns and cities for temporary work during the agricultural low season; the relatively low rate of migration may indicate that the utility costs of migration are non-trivial. Among movers, there is an estimated increase in per capita consumption among the sending household (excluding the migrant) of roughly 30% over two years (although effects appear more modest when the migrant is included) and 25% average gain in earnings (not statistically significant) among those assigned to the subsidy. Overall, the study provides some indication that there are positive returns to temporary seasonal migration among rural workers who are typically non-movers. Nonetheless, the earnings gains are fairly modest in size and note that they are closer in magnitude to the small gaps we estimate in this paper than to those found in many other recent contributions.

The case of urban-born non-movers is less well understood and raises some intriguing possibilities. Recall from Table 2 that individuals raised in urban areas have much higher cognitive scores (on a test of fluid intelligence) than those raised in rural areas. It is difficult to definitively determine the causes of this gap, but there are several plausible channels. One is simply that wave after wave of rural to urban (urban to rural) migration by positively (negatively) selected individuals over many decades, combined with partial heritability of cognitive ability, have reshaped the underlying ability distributions in these two sectors. This would simply be an inter-generational extension of the patterns of individual selection across urban and rural areas that we and Young (2013) document, and would not necessarily change the interpretation of our main results.

Another explanation, which is not mutually exclusive, is that there is a lower cost to skill acquisition in urban areas, either due to improved provision of schooling for children there or

something else about the nature of social interactions (e.g., the density of such interactions or other forms of intellectual stimulation in childhood). In other words, given the importance of early childhood circumstances for lifetime cognitive development (e.g., Gertler et al. 2014), growing up in a city might generate higher average adult skill levels. This would properly be understood as a causal effect of urban residence on individual labor productivity, albeit in the very long-run and on the movers' children rather than for themselves. Anecdotally, many migrants do claim to move in order to improve their children's economic wellbeing more than their own. These effects would not be captured even in the five-year follow-up period that we consider in this study (in Figure 4), but could be contributing to large, persistent and real causal urban-rural productivity gaps overall.²⁵

In our view, the study of sectoral productivity gaps remains a research area ripe for further empirical analysis. Some natural next steps include extending our long-run panel data analysis to new countries, settings and time periods (as appropriate data becomes available); conducting further experiments along the lines of Bryan et al. (2014) and McKenzie et al. (2010) to induce at least partially random selection in migration, thus generating "local" estimates in new sub-populations and improving understanding of the nature of costs and constraints facing potential migrants; and exploration of very long-run and even inter-generational effects of parent sectoral and residential choice on child ability and productivity.

²⁵ Indeed, Nakamura, Sigurdsson, and Steinsson (2016) study migration induced by a volcanic explosion in Iceland, and show that adult movers gain little from moving out of a rural area but their children earn far more in the long-run.

References

- Alvarez, J. A. (2015) “The Agricultural Wage Gap: Evidence from Brazilian Micro-data,” Working paper.
- Au, C.-C. and Henderson, J. V. (2006) “How migration restrictions limit agglomeration and productivity in China,” *Journal of Development Economics*, Vol. 80, pp. 350–388.
- Baird, S., Hamory, J., and Miguel, E. (2008) “Tracking, Attrition and Data Quality in The Kenyan Life Panel Survey Round 1 (KLPS-1),” *University of California CIDER*, Working Paper.
- Bazzi, S. (2016) “Wealth Heterogeneity and the Income Elasticity of Migration,” *American Economic Journal: Applied Economics*, Forthcoming.
- Bazzi, S., Gaduh, A., Rothenberg, A. D., and Wong, M. (2016) “Skill Transferability, Migration, and Development: Evidence from Population Resettlement in Indonesia,” *American Economic Review*, Vol. 106, pp. 2658–2698.
- Beegle, K., Weerdt, J. D., and Dercon, S. (2011) “Migration and Economic Mobility in Tanzania: Evidence from a Tracking Survey,” *Review of Economics and Statistics*, Vol. 93, pp. 1010–1033.
- Bryan, G., Chowdhury, S., and Mobarak, A. (2014) “Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh,” *Econometrica*, Vol. 82, pp. 1671–1748.
- Card, D. (1996) “The Effects of Union on the Structure of Wages: A Longitudinal Analysis,” *Econometrica*, Vol. 64, pp. 957–979.
- Caselli, F. (2005) “Accounting for Cross-Country Income Differences,” in P. Aghion and S. Durlauf eds. *Handbook of Economic Growth*, Vol. 1 of Handbook of Economic Growth: Elsevier, Chap. 9, pp. 679–741.
- Chiquiar, D. and Hanson, G. H. (2005) “International Migration, Self-Selection, and The Distribution of Wages: Evidence from Mexico and the United States,” *Journal of Political Economy*, Vol. 113, pp. 239–281.
- de la Roca, J. and Puga, D. (2016) “Learning by Working in Big Cities,” *Review of Economic Studies*, Forthcoming.
- de Mel, S., McKenzie, D., and Woodruff, C. (2009) “Measuring microenterprise profits: Must we ask how the sausage is made?” *Journal of Development Economics*, Vol. 88.
- Gertler, P., Heckman, J., Pinto, R., Zanolini, A., Vermeerch, C., Walker, S., Chang, S. M., and Sally, G.-M. (2014) “Labor Market Returns to an Early Childhood Stimulation Intervention in Jamaica,” *Science*, Vol. 344, pp. 998–1001.
- Gollin, D., Lagakos, D., and Waugh, M. E. (2014) “The Agricultural Productivity Gap,” *Quarterly Journal of Economics*, Vol. 129, pp. 939–993.


- Gollin, D., Parente, S., and Rogerson, R. (2002) “The Role of Agriculture in Development,” *American Economic Review*, Vol. 92, pp. 160–164.
- Graham, B. and Temple, J. (2006) “Rich nations, poor nations: how much can multiple equilibria explain?” *Journal of Economic Growth*, Vol. 11, pp. 5–41.
- Hendricks, L. and Schoellman, T. (2016) “Human Capital and Development Accounting: New Evidence from Wage Gains at Migration,” 2016 Meeting Papers 159, Society for Economic Dynamics.
- Herrendorf, B., Rogerson, R., and Valentinyi, Á. (2014) “Chapter 6 - Growth and Structural Transformation,” in P. Aghion and S. N. Durlauf eds. *Handbook of Economic Growth*, Vol. 2 of Handbook of Economic Growth: Elsevier, pp. 855–941.
- Herrendorf, B. and Schoellman, T. (2016) “Wages, Human Capital, and Structural Transformation,” *Working Paper*.
- Hsieh, C.-T. and Klenow, P. (2009) “Misallocation and Manufacturing TFP in China and India,” *Quarterly Journal of Economics*, Vol. 124, pp. 1403–1448.
- Johnston, B. and Kilby, P. (1978) *Agriculture and Structural Transformation: Economic Strategies in Late-developing Countries*: Oxford University Press.
- Kleemans, M. (2016) “Migration Choice under Risk and Liquidity Constraints,” Working Paper.
- Kleemans, M. and Magruder, J. (2016) “Labor Market Changes In Response To Immigration: Evidence From Internal Migration Driven By Weather Shocks,” *Economic Journal*, accepted.
- Kremer, M., Miguel, E., and Thornton, R. (2009) “Incentives to Learn,” *Review of Economics and Statistics*, Vol. 91, pp. 437–456.
- Kuznets, S. (1973) “Modern Economic Growth: Findings and Reflections.,” *American Economic Review*, Vol. 63, pp. 247–258.
- Lagakos, D. and Waugh, M. E. (2013) “Selection, Agriculture, and Cross-Country Productivity Differences,” *American Economic Review*, Vol. 103, pp. 948–980.
- Lee, K., Brewer, E., Christiano, C., Meyo, F., Miguel, E., Podolsky, M., Rosa, J., and Wolfram, C. (2015) “Barriers to Electrification for ”Under Grid” Households in Rural Kenya,” *Development Engineering*, Vol. 1, pp. 26–35.
- Lemieux, T. (1998) “Estimating The Effects of Unions on Wage Inequality in a Panel Data Model with Comparative Advantage and Nonrandom Selection,” *Journal of Labor Economics*, Vol. 16, pp. 261–291.
- Lewis, W. (1955) *The Theory of Economic Growth*, Homewood: R. D. Irwin.

- McKenzie, D., Gibson, J., and Stillman, S. (2010) “How Important Is Selection? Experimental vs. Non-Experimental Measures of The Income Gains from Migration,” *Journal of the European Economic Association*, Vol. 8, pp. 913–945.
- Miguel, E. and Kremer, M. (2004) “Worms: Identifying Impacts on Education and Health in The Presence of Treatment Externalities,” *Econometrica*, Vol. 72, pp. 159–217.
- Munshi, K. and Rosenzweig, M. (2016) “Networks and Misallocation: Insurance, Migration, and The Rural-Urban Wage Gap,” *American Economic Review*, Vol. 106, pp. 46–98.
- Nakamura, E., Sigurdsson, J., and Steinsson, J. (2016) “The Gift of Moving: Intergenerational Consequences of a Mobility Shock,” *Working Paper*.
- Pack, H. (1972) *Employment and productivity in Kenyan manufacturing*: Institute for Development Studies, University of Nairobi.
- Porzio, T. (2016) “Cross-Country Differences in The Optimal Allocation of Talent and Technology,” *Working Paper*.
- Preobrazhensky, E. (1921) *The crisis of Soviet industrialization*(1980 ed.), London: MacMillan.
- Restuccia, D. and Rogerson, R. (2008) “Policy distortions and aggregate productivity with heterogeneous plants,” *Review of Economic Dynamics*, Vol. 11, pp. 707–720.
- Restuccia, D., Yang, D. T., and Zhu, X. (2008) “Agriculture and aggregate productivity: A quantitative cross-country analysis,” *Journal of Monetary Economics*, Vol. 55, pp. 234–250.
- Rosenstein-Rodan, P. N. (1943) “Problems of Industrialisation of Eastern and South-Eastern Europe,” *Economic Journal*, Vol. 53, pp. 202–211.
- Rostow, W. W. (1960) *The Stages of Economic Growth: A Non-Communist Manifesto*, Cambridge: Cambridge University Press.
- Roy, A. D. (1951) “Some Thoughts on the Distribution of Earnings,” *Oxford Economic Papers*, Vol. 3, pp. 135–146.
- Rubalcava, L., Teruel, G., Thomas, D., and Goldman, N. (2008) “The Healthy Migrant Effect: New Findings from the Mexican Family Life Survey,” *American Journal of Public Health*, Vol. 98, pp. 78–84.
- Schultz, T. P. (1988) *Handbook of Development Economics*, Vol. 1, Chap. 13 Education Investments and Returns, pp. 543–630: Elsevier.
- Strauss, J., Witoelar, F., Sikoki, B., and Wattie, A. (2009) “The Fourth Wave of the Indonesian Family Life Survey (IFLS4): Overview and Field Report,” *WR-675/1-NIA/NICHD*.

- Stren, R., Halfani, M., and Malombe, J. (1994) *Beyond Capitalism vs. Socialism in Kenya and Tanzania*, Chap. Coping with Urbanization and Urban Policy, Colorado: Lynne Rienner.
- Suri, T. (2011) "Selection and Comparative Advantage in Technology Adoption," *Econometrica*, Vol. 79, pp. 159–209.
- Thomas, D., Witoelar, F., Frankenberg, E., Sikoki, B., Strauss, J., Sumantri, C., and Suriastini, W. (2012) "Cutting the costs of attrition: Results from the Indonesia Family Life Survey," *Journal of Development Economics*, Vol. 98, pp. 108–123.
- Yang, D. (2006) "Why Do Migrants Return to Poor Countries? Evidence from Philippine Migrants' Responses to Exchange Rate Shocks," *Review of Economics and Statistics*, Vol. 88, pp. 715–735.
- Young, A. (2013) "Inequality, The Urban-Rural Gap, And Migration," *Quarterly Journal of Economics*, Vol. 128, pp. 1727–1785.

Tables and Figures

Figure 1: Productivity Gap in Total Earnings



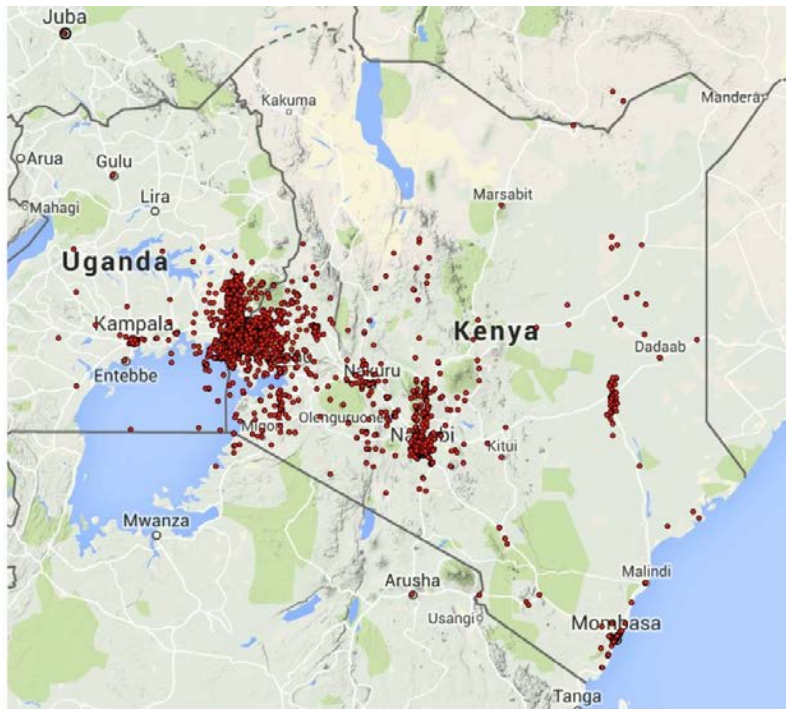
Notes: GLW refers to estimates from Gollin, Lagakos, and Waugh (2014), Online Appendix Table 4. For comparability, the figure reports log transformed numbers from their columns 4 and 5 for Indonesia and Kenya, respectively. Symbols here represent point estimates, and vertical lines represent 95% confidence intervals. Panel A estimates from the IFLS come from Table 5, panel A: “Raw” is the mean difference estimate from column (1), “Adjusted” is the regression adjusted mean difference estimate from column (3), and “FE, wages” is the fixed effects regression estimate of wages on an urban indicator and squared-age from column (7). Corresponding estimates from the KLPS come from Table 5, panel B. Estimates in panels C and D come from the same columns in Table 6, panels A and B, respectively. Note that the confidence intervals for the estimates from the IFLS are smaller than the size of the symbols and are therefore not visible.

Figure 2: Sample Areas

(A) Indonesia Family Life Survey




(B) Kenya Life Panel Survey




Notes: Panel A shows the residential locations of individuals during the 1988–2008 sample period of rounds 1–4 of the IFLS. For the Kenyan sample, Panel B shows individuals’ residential locations during the 1998–2014 sample period that was collected during rounds 2 and 3 of the KLPS. Individuals living outside of Kenya are dropped from the analysis. The location information of both datasets are described in more detail in Section 3.

Figure 3: Joint Distribution of Rural and Urban Productivities

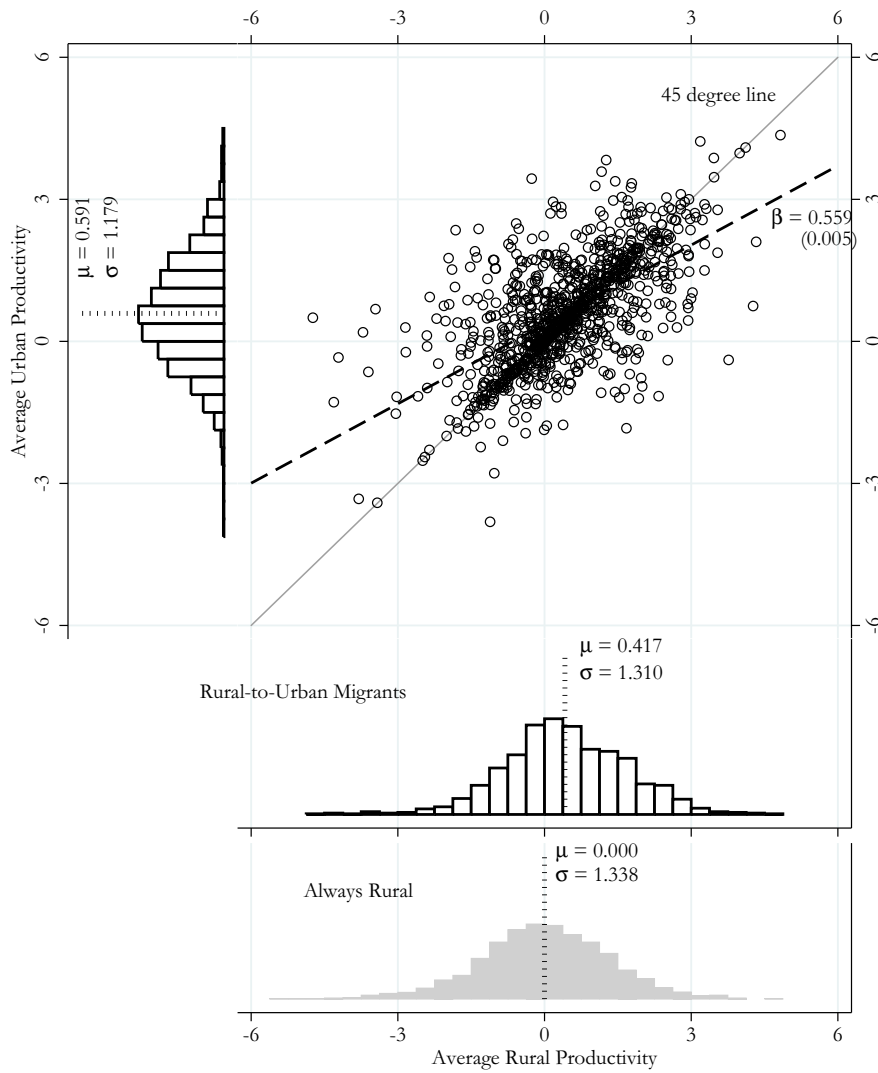
(A) Indonesia (Born Rural)



(B) Indonesia (Born Urban)




(C) Kenya (Born Rural)




Notes: Productivities are recovered individual-urban status effects from a fixed effects regression of log wages on squared age and indicators for time period on the same sample used in Table 5. Productivities are normalized such that the average productivity of rural non-migrants has zero mean. Histograms on the bottom of Panel A represent marginal distributions of rural productivities for “Always Rural” non-migrants (grey) and migrants (hollow). Marginal distribution of estimated urban productivities for migrants reported on the left (hollow). Means and standard deviations reported in log points. Scatterplot presents joint distribution for migrants with best fit line. Bootstrapped standard error of the slope reported in parentheses from 1,000 iterations of block sampling of individuals with replacement. Panel B presents a histogram of “Always Urban” urban productivities of non-migrants (grey) at the top left, an adjacent histogram of migrant urban productivities (hollow), and migrant rural productivities (grey) below. Joint distribution of urban and rural productivities and corresponding best fit line presented similar to panel A. Panel C mimics the format of Panel A except uses data from the KLPS.

Figure 4: Event Study of Urban Migration

(A) Indonesia (Born Rural)



(B) Kenya (Born Rural)



Notes: Panel A uses data from individuals in the IFLS who are born in rural areas, and Panel B uses data from the KLPS. Please refer to Section 3 for further details on the data. The top half of each panel reports event study coefficients β_{τ} from a regression of log wages described in equation 8 (in section 4D). The solid line represents the point estimate, and the dashed lines represent the 95% confidence interval. Estimates represent the difference in mean wages between movers and non-movers net of the difference that existed in the period prior to the move. Regressions include individual fixed effects, time fixed effects, squared-age, and indicator variables that pool observations exceeding a five year window of the move. The lower half of each panel reports the fraction of people having no rural observations from period zero to the period of interest. (The estimated fraction of survivors can in principle increase due to sample composition changes as can be seen in the lower half of panel B.) In the IFLS, there are 317 individuals who have observed wages in the year of the move and the year prior; 77 of these individuals report wages 5 years later. In the KLPS, these numbers are 320 and 51, respectively.

Table 1: Non-Agriculture/Agriculture and Urban/Rural

(A) Indonesia (Main Analysis Sample)

	Rural	Urban	Total
Agriculture	42.1%	3.4%	34.0%
Non-Agriculture	57.9%	96.6%	66.0%
Number of Observations	91,765	24,132	115,897

(B) Kenya (Main Analysis Sample)

	Rural	Urban	Total
Agriculture	25.9%	5.3%	15.1%
Non-Agriculture	74.1%	94.7%	84.9%
Number of Observations	61,139	67,300	128,439

(C) Kenya (12 Months with Subsistence Agricultural Module)

	Rural	Urban	Total
Agriculture	59.0%	9.2%	40.6%
Non-Agriculture	41.0%	90.8%	59.4%
Number of Observations	27,171	15,928	43,099

Notes: Panel A reports summary statistics from the Indonesia Family Life Survey (IFLS), and Panels B and C present data from the Kenya Life Panel Survey (KLPS); both are described in more detail in Section 3. Panel A shows the main Indonesian analysis sample of 115,897 individual-year observations, for individuals aged 16 and above for whom earnings measures are available. As described in Section 3, all reported wage employment, self employment, and subsistence agricultural employment is included. Panel B shows the main Kenyan analysis sample of 128,439 individual-month observations of individuals aged 16 and above for whom earnings measures are available. As described in Section 3, wage and self-employment are included for all years, including agricultural wage labor and agricultural self-employment if annual revenues exceeded 40,000 Ksh, and subsistence agricultural data is included from the 12 months preceding the KLPS 3 interview date if the respondent is both the main decision maker and the only supplier of labor in agriculture. Panel C shows data from the 12 months where subsistence agriculture data is available and counts all agricultural activities: including when the person is not the main decision maker and when others work on the agricultural activity; in the case of the latter, the agricultural productivity is weighted by the share of hours that the individual supplies. Each cell reports the percentage of observations by agricultural and non-agricultural sector, and by rural and urban area. In both the IFLS and KLPS, individuals are characterized by the sector of their main employment. The urban indicator from the IFLS is obtained from survey responses to the question: “Is the area you live in a village, a town or a city?” If the person reports living in a town or city, the urban indicator variable equals 1. For the KLPS, the urban indicator equals 1 if the person reports living in a large town or city. Please see the text in section 3B for further details on this classification. The list of Kenyan urban areas and frequency of occurrence in the panel are given in Table A4.

Table 2: Summary Statistics

(A) Indonesia

	All N=18211	Always Rural N=10592	Rural-to-Urban Migrants N=1985	Always Urban N=2606	Urban-to-Rural Migrants N=3028	Obs
Primary Ed.	0.800 [0.400]	0.722 [0.448]	0.902 [0.297]	0.960 [0.196]	0.872 [0.334]	18211
Secondary Ed.	0.285 [0.451]	0.172 [0.377]	0.357 [0.479]	0.575 [0.494]	0.383 [0.486]	18211
College	0.071 [0.258]	0.033 [0.178]	0.087 [0.282]	0.190 [0.392]	0.096 [0.294]	18211
Female	0.421 [0.494]	0.412 [0.492]	0.424 [0.494]	0.439 [0.496]	0.436 [0.496]	18211
Raven's Z-score	0.003 [0.961]	-0.117 [0.957]	0.121 [0.939]	0.315 [0.886]	0.119 [0.967]	11220

(B) Kenya

	All N=4537	Always Rural N=1546	Rural-to-Urban Migrants N=2991	Always Urban	Urban-to-Rural Migrants	Obs
Primary Ed.	0.742 [0.438]	0.648 [0.478]	0.791 [0.407]			4537
Secondary Ed.	0.361 [0.480]	0.245 [0.430]	0.420 [0.494]			4537
College	0.037 [0.189]	0.012 [0.110]	0.050 [0.218]			4537
Female	0.530 [0.499]	0.537 [0.499]	0.527 [0.499]			4537
Raven's Z-score	0.049 [0.986]	-0.145 [0.981]	0.149 [0.974]			4537

Notes: Panel A reports summary statistics from the IFLS and panel B reports summary statistics from the KLPS. Sample standard deviations reported in brackets below sample means. The sample is limited to respondents who report age, gender, and years of education and have at least one person-time observation that has income, hours, location of residence, and sector of occupation. In panel A (Indonesia), “Rural-to-Urban Migrants” are individuals born in rural areas and are observed in urban areas in our sample with data on wages, hours, and sector. “Urban-to-Rural Migrants” are defined similarly. In panel B (Kenya), all individuals are born rural; migrants are those who have subsequent observations with information on income, hours, and sector in urban areas. Rows correspond to the fraction within each column who have completed primary education, secondary education, and college; the fraction female; and the average score from a Raven’s matrices exam, normalized to be mean zero and standard deviation one.

Table 3: Correlates of Urban Migration

(A) Indonesia (Born Rural)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Primary Ed.	0.128*** (0.006)					0.0928*** (0.008)	0.101*** (0.006)
Secondary Ed.		0.153*** (0.010)				0.0858*** (0.012)	0.111*** (0.011)
College			0.184*** (0.021)			0.0155 (0.027)	0.0677** (0.023)
Female				0.00683 (0.007)		0.0300*** (0.008)	0.0306*** (0.007)
Raven's Z-score					0.0315*** (0.004)	0.0150*** (0.004)	
Constant	0.0617*** (0.004)	0.127*** (0.003)	0.150*** (0.003)	0.155*** (0.004)	0.145*** (0.004)	0.0344*** (0.007)	0.0444*** (0.005)
Observations	12577	12577	12577	12577	7861	7861	12577

(B) Kenya (Born Rural)

	(1)	(2)	(3)	(4)	(5)	(6)
Primary Ed.	0.167*** (0.017)					0.0880*** (0.019)
Secondary Ed.		0.171*** (0.014)				0.0979*** (0.017)
College			0.236*** (0.025)			0.121*** (0.027)
Female				-0.00898 (0.014)		0.0149 (0.014)
Raven's Z-score					0.0678*** (0.007)	0.0312*** (0.008)
Constant	0.535*** (0.015)	0.598*** (0.009)	0.650*** (0.007)	0.664*** (0.010)	0.656*** (0.007)	0.545*** (0.017)
Observations	4537	4537	4537	4537	4537	4537

Notes: Please see Table 2 for sample restrictions and row variable definitions. Each cell reports a regression coefficient with an indicator for being an urban migrant as the dependent variable. Both panels are estimated on individuals who are born rural. Columns 6 and 7 report coefficients from multiple regressions with corresponding rows as included covariates. Column 7 for Indonesia omits the Raven's matrix exam to preserve sample size. Robust standard errors reported below in parentheses, *** p<0.01, ** p<0.05, * p<0.1.

Table 4: Correlates of Employment in Non-Agriculture

(A) Indonesia (Born Rural)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Primary Ed.	0.270*** (0.010)					0.213*** (0.015)	0.240*** (0.010)
Secondary Ed.		0.268*** (0.008)				0.161*** (0.011)	0.190*** (0.009)
College			0.283*** (0.010)			0.0621*** (0.015)	0.0730*** (0.011)
Female				0.0742*** (0.008)		0.121*** (0.010)	0.125*** (0.008)
Raven's Z-score					0.0608*** (0.005)	0.0302*** (0.006)	
Constant	0.484*** (0.009)	0.632*** (0.005)	0.674*** (0.004)	0.655*** (0.006)	0.717*** (0.005)	0.449*** (0.015)	0.413*** (0.010)
Observations	12577	12577	12577	12577	7861	7861	12577

(B) Kenya (Born Rural)

	(1)	(2)	(3)	(4)	(5)	(6)
Primary Ed.	0.145*** (0.012)					0.106*** (0.014)
Secondary Ed.		0.106*** (0.008)				0.0464*** (0.009)
College			0.0943*** (0.007)			0.0188* (0.008)
Female				0.0104 (0.009)		0.0274** (0.009)
Raven's Z-score					0.0457*** (0.004)	0.0207*** (0.005)
Constant	0.796*** (0.012)	0.865*** (0.006)	0.900*** (0.005)	0.898*** (0.007)	0.901*** (0.004)	0.792*** (0.013)
Observations	4537	4537	4537	4537	4537	4537

Notes: Please see Table 2 for sample restrictions and row variable definitions. Each cell reports a regression coefficient with an indicator for being ever being employed in non-agriculture as the dependent variable. Panel A (Indonesia) is estimated on individuals who are born in rural areas, whereas panel B (Kenya) includes the full sample subject to previously defined sample restrictions. Please see Appendix Table A1 for analogous regressions of individuals born urban in Indonesia. Columns 6 and 7 report coefficients from a multiple regression with corresponding rows as included covariates. Column 7 for Indonesia omits the Raven's matrix exam to preserve sample size. Robust standard errors reported below in parentheses, *** p<0.01, ** p<0.05, * p<0.1.

Table 5: Non-Agricultural/Agricultural Gap in Earnings

(A) Indonesia

	Dependent variable: Log Earnings (in IDR)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7) Log Wage	(8) Log Real Wage
Non-agricultural employment	0.622*** (0.016)	0.514*** (0.016)	0.323*** (0.015)	0.309*** (0.019)	0.158*** (0.027)	0.171*** (0.025)	0.047 (0.031)	0.045 (0.031)
Log hours		0.531*** (0.025)	0.358*** (0.023)	0.395*** (0.027)	0.228*** (0.053)	0.323*** (0.034)		
Log hours squared		-0.021*** (0.005)	-0.005 (0.004)	-0.011** (0.005)	0.015 (0.010)	-0.014** (0.006)		
Female			-0.460*** (0.014)	-0.457*** (0.017)	-0.615*** (0.042)			
Age			0.061*** (0.002)	0.061*** (0.004)	0.052*** (0.006)			
Age squared			-0.001*** (0.000)	-0.001*** (0.000)	-0.001*** (0.000)	-0.001*** (0.000)	-0.000*** (0.000)	-0.000*** (0.000)
Years of education			0.030*** (0.005)	0.019*** (0.006)	0.055*** (0.012)			
Years of education squared			0.003*** (0.000)	0.003*** (0.000)	0.001 (0.001)			
Normalized Ravens				0.059*** (0.009)				
Normalized Ravens squared				0.021*** (0.007)				
Individual fixed effects	N	N	N	N	N	Y	Y	Y
Time fixed effects	N	Y	Y	Y	Y	Y	Y	Y
Switchers only					Y			
Number of observations	115897	115897	115897	68700	15648	115897	115897	115897
Number of individuals	18211	18211	18211	11220	1939	18211	18211	18211

(B) Kenya

	Dependent variable: Log Earnings (in KSh)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7) Log Wage	(8) Log Real Wage
Non-agricultural employment	0.791*** (0.076)	0.593*** (0.078)	0.569*** (0.072)	0.559*** (0.072)	0.331*** (0.105)	0.326*** (0.099)	0.134 (0.117)	0.121 (0.117)
Log hours		0.391** (0.169)	0.309** (0.154)	0.299** (0.152)	0.247 (0.231)	0.399** (0.192)		
Log hours squared		0.007 (0.028)	0.011 (0.026)	0.013 (0.026)	0.014 (0.041)	-0.014 (0.032)		
Female			-0.568*** (0.050)	-0.543*** (0.050)	-0.656*** (0.134)			
Age			0.174** (0.072)	0.170** (0.073)	0.111 (0.177)			
Age squared			-0.003* (0.001)	-0.003* (0.001)	-0.002 (0.004)	-0.001 (0.002)	-0.002 (0.002)	-0.002 (0.002)
Years of education			-0.015 (0.049)	-0.046 (0.048)	-0.073 (0.113)			
Years of education squared			0.006** (0.002)	0.007*** (0.002)	0.006 (0.007)			
Normalized Ravens				0.088*** (0.029)	0.095 (0.065)			
Normalized Ravens squared				-0.037 (0.025)	-0.178*** (0.064)			
Individual fixed effects	N	N	N	N	N	Y	Y	Y
Time fixed effects	N	Y	Y	Y	Y	Y	Y	Y
Switchers only					Y			
Number of observations	128439	128439	128439	128439	15657	128439	128439	128439
Number of individuals	4537	4537	4537	4537	364	4537	4537	4537

Notes: Panel A uses data from rounds 1–4 of the Indonesia Family Life Survey (IFLS), described in Section 3. Panel B uses data from rounds 2–3 of the Kenya Life Panel Survey (KLPS), also described in Section 3. The dependent variable in columns 1 to 6 is log earnings, which are the combined earnings from wage and self-employment, reported in Indonesian Rupiah. If an individual has multiple jobs in the same time period, earnings from all employment are included. The dependent variable in column 7 is log wage, which is obtained by dividing log earnings by total hours worked. The dependent variable in column 8 is log wage adjusted for differences in prices between urban and rural areas. The covariate “Non-agricultural employment” is an indicator variable which equals 1 if the person main employment is in the non-agricultural sector. The covariate log hours sums up hours worked in all employment. The sample size in column 4 is smaller in Panel A because the Raven’s test was administered only for a subset of the sample. The sample size in column 5 is smaller because it only includes “switchers” who have at least one observation in both the non-agricultural and agricultural sector. All regressions are clustered at the individual level. Robust standard errors are in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 6: Urban/Rural Gap in Earnings

(A) Indonesia

	Dependent variable: Log Earnings (in IDR)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7) Log Wage	(8) Log Real Wage
Urban	0.628*** (0.017)	0.534*** (0.016)	0.276*** (0.014)	0.245*** (0.018)	0.084*** (0.023)	0.018 (0.016)	-0.007 (0.019)	-0.103*** (0.019)
Log hours		0.512*** (0.025)	0.347*** (0.023)	0.387*** (0.027)	0.446*** (0.121)	0.322*** (0.034)		
Log hours squared		-0.012*** (0.005)	0.001 (0.004)	-0.006 (0.005)	-0.016 (0.020)	-0.013** (0.006)		
Female			-0.418*** (0.014)	-0.415*** (0.017)	-0.313*** (0.040)			
Age			0.061*** (0.002)	0.064*** (0.004)	0.084*** (0.007)			
Age squared			-0.001*** (0.000)	-0.001*** (0.000)	-0.001*** (0.000)	-0.001*** (0.000)	-0.000*** (0.000)	-0.000*** (0.000)
Years of education			0.039*** (0.005)	0.027*** (0.006)	0.020 (0.017)			
Years of education squared			0.003*** (0.000)	0.003*** (0.000)	0.003*** (0.001)			
Normalized Ravens				0.064*** (0.009)				
Normalized Ravens squared				0.020*** (0.007)				
Individual fixed effects	N	N	N	N	N	Y	Y	Y
Time fixed effects	N	Y	Y	Y	Y	Y	Y	Y
Switchers only					Y			
Number of observations	115897	115897	115897	68700	10888	115897	115897	115897
Number of individuals	18211	18211	18211	11220	1482	18211	18211	18211

(B) Kenya

	Dependent variable: Log Earnings (in KSh)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7) Log Wage	(8) Log Real Wage
Urban	0.851*** (0.051)	0.750*** (0.050)	0.676*** (0.047)	0.666*** (0.046)	0.334*** (0.067)	0.247*** (0.070)	0.132 (0.082)	0.054 (0.082)
Log hours		0.351** (0.164)	0.283* (0.150)	0.276* (0.149)	0.625** (0.276)	0.405** (0.198)		
Log hours squared		0.007 (0.027)	0.011 (0.025)	0.012 (0.025)	-0.036 (0.043)	-0.016 (0.033)		
Female			-0.517*** (0.048)	-0.495*** (0.049)	-0.452*** (0.094)			
Age			0.152** (0.073)	0.149** (0.073)	0.075 (0.115)			
Age squared			-0.002 (0.001)	-0.002 (0.001)	-0.001 (0.002)	-0.001 (0.002)	-0.002 (0.002)	-0.002 (0.002)
Years of education			-0.029 (0.045)	-0.054 (0.045)	-0.128* (0.070)			
Years of education squared			0.006*** (0.002)	0.007*** (0.002)	0.011*** (0.004)			
Normalized Ravens				0.079*** (0.028)	0.013 (0.049)			
Normalized Ravens squared				-0.018 (0.024)	-0.024 (0.041)			
Individual fixed effects	N	N	N	N	N	Y	Y	Y
Time fixed effects	N	Y	Y	Y	Y	Y	Y	Y
Switchers only					Y			
Number of observations	128439	128439	128439	128439	39417	128439	128439	128439
Number of individuals	4537	4537	4537	4537	1039	4537	4537	4537

Notes: Panel A uses data from the IFLS and Panel B uses data from the KLPS. Please refer to Section 3 for further details on the data and to the notes of Table 5 for additional information on the variables. For the IFLS, the urban indicator is obtained from survey responses to the question: “Is the area you live in a village, a town or a city?” If the person reports living in a town or city, the urban indicator variable equals 1. For the KLPS, the urban indicator equals 1 if a person lives in a large town or a city. Please see the text in section 3B for further details on this classification. Column 5 only includes switchers, who are defined as individuals with at least one observation in both an urban and rural area. All regressions are clustered at the individual level. Robust standard errors are in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 7: Gap in Earnings for those Born in Rural and Urban Areas, Indonesia

(A) Indonesian individuals born in rural areas

	Dependent variable: Log Earnings (in IDR)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Log Wage				Log Wage			
Non-agricultural employment	0.553*** (0.018)	0.314*** (0.017)	0.178*** (0.028)	0.059* (0.034)				
Urban					0.646*** (0.026)	0.320*** (0.021)	0.025 (0.021)	-0.004 (0.026)
Individual fixed effects	N	N	Y	Y	N	N	Y	Y
Control variables and time FE	N	Y	Y	Y	N	Y	Y	Y
Number of observations	83810	83810	83810	83810	83810	83810	83810	83810
Number of individuals	12577	12577	12577	12577	12577	12577	12577	12577

(B) Indonesian individuals born in urban areas

	Dependent variable: Log Earnings (in IDR)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Log Wage				Log Wage			
Non-agricultural employment	0.573*** (0.047)	0.274*** (0.039)	0.136** (0.054)	-0.008 (0.065)				
Urban					0.416*** (0.026)	0.197*** (0.021)	0.008 (0.025)	-0.013 (0.030)
Individual fixed effects	N	N	Y	Y	N	N	Y	Y
Control variables and time FE	N	Y	Y	Y	N	Y	Y	Y
Number of observations	32087	32087	32087	32087	32087	32087	32087	32087
Number of individuals	5634	5634	5634	5634	5634	5634	5634	5634

Notes: Panels A and B repeat the analyses of Table 5A and Table 6A for those born in rural and urban areas, respectively. Please refer to Section 3 for further details on the data and to the notes of Table 5 and 6 for additional information on the variables. Control variables include log hours worked, log hours worked squared, age, age squared, years of education, years of education squared and an indicator for being female. When also including individual fixed effects in columns 3 and 7, the control variables are reduced to only log hours worked, log hours worked squared, and age squared because the others are absorbed by the individual fixed effects. When using log wage as the dependent variable in columns 4 and 8, only age squared is included as a control variable. All regressions are clustered at the individual level. Robust standard errors are in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 8: Gaps in Consumption

(A) Indonesia

	Dependent variable: Log Consumption (in IDR)					
	(1)	(2)	(3)	(4)	(5)	(6)
Non-agricultural employment	0.543*** (0.011)	0.290*** (0.010)	0.030 (0.018)			
Urban				0.526*** (0.013)	0.296*** (0.012)	0.030 (0.022)
Individual fixed effects	N	N	Y	N	N	Y
Control variables and time FE	N	Y	Y	N	Y	Y
Number of observations	38280	38280	38280	38280	38280	38280
Number of individuals	19695	19695	19695	19695	19695	19695

(B) Kenya

	Dependent variable: Log Meals Eaten					
	(1)	(2)	(3)	(4)	(5)	(6)
Non-agricultural employment	0.043** (0.019)	0.024 (0.020)	-0.078* (0.044)			
Urban				0.035** (0.014)	0.026* (0.015)	-0.023 (0.040)
Individual fixed effects	N	N	Y	N	N	Y
Control variables and time FE	N	Y	Y	N	Y	Y
Number of observations	4308	4308	4308	4308	4308	4308
Number of individuals	3690	3690	3690	3690	3690	3690

Notes: Panel A uses data on total consumption from the IFLS, and Panel B uses data on meals eaten in the last day from the KLPS. Unlike previous tables, the sample includes individuals with and without earnings measures. Consumption data in the IFLS are obtained by adding up the value of food and non-food consumption at the household level and dividing this by the number of household members. The data was collected for each of the four waves so each household has four observations at most. Analyses by food and non-food consumption separately can be found in Appendix Table A6, and Appendix Tables A7 and A8 provide consumption analyses when using the sample with positive earnings measures. Data on meals eaten in Kenya are available from KLPS rounds 2 and 3 and refer to the day prior to the survey date. In the analysis sample, 0.6% of individual-time observations ate no meals in the prior day, 11.0% ate one meal, 53.0% ate two meals, 33.9% ate three meals, and 1.5% ate four or more. Control variables in both panels include age, age squared, years of education, years of education squared, and an indicator for being female. When also including individual fixed effects in columns 3 and 6, the control variables are reduced to only age squared because the others are absorbed by the individual fixed effects. All regressions are clustered at the individual level. Robust standard errors are in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 9: Urban/Rural Gap in Wages for Top 5 Cities

(A) Indonesia

	Dependent variable: Log Wages (in IDR)			
	(1)	(2)	(3)	(4)
Urban	0.428*** (0.017)	0.309*** (0.023)	0.061*** (0.021)	-0.020 (0.024)
Jakarta (population 10 million)		0.257*** (0.032)	0.252*** (0.028)	0.009 (0.040)
Surabaya (population 2.8 million)		-0.126* (0.071)	-0.092 (0.063)	0.058 (0.089)
Bandung (population 2.6 million)		0.383*** (0.060)	0.227*** (0.051)	0.233** (0.095)
Medan (population 2.5 million)		0.151** (0.071)	0.218*** (0.069)	-0.038 (0.117)
Bekasi (population 2.5 million)		0.454*** (0.063)	0.283*** (0.052)	-0.061 (0.070)
Individual fixed effects	N	N	N	Y
Control variables and time FE	N	N	Y	Y
Number of observations	110980	110980	110980	110980
Number of individuals	17933	17933	17933	17933


(B) Kenya

	Dependent variable: Log Wages (in KSh)			
	(1)	(2)	(3)	(4)
Urban	0.589*** (0.055)	0.376*** (0.075)	0.339*** (0.068)	0.048 (0.096)
Nairobi (population 3.4 million)		0.362*** (0.076)	0.300*** (0.066)	0.133* (0.078)
Mombasa (population 1.2 million)		0.257** (0.101)	0.238*** (0.091)	0.356*** (0.108)
Kisumu (population 0.4 million)		0.034 (0.188)	0.109 (0.157)	0.104 (0.484)
Nakuru (population 0.3 million)		0.248** (0.123)	0.136 (0.100)	-0.024 (0.169)
Eldoret (population 0.3 million)		0.120 (0.153)	0.002 (0.157)	-0.123 (0.184)
Individual fixed effects	N	N	N	Y
Control variables and time FE	N	N	Y	Y
Number of observations	128439	128439	128439	128439
Number of individuals	4537	4537	4537	4537

Notes: Panel A uses data from the IFLS and Panel B uses data from the KLPS. Please refer to Section 3 for further details on the data and to the notes of Table 5 for additional information on the variables. The covariate “Urban” is an indicator variable that equals 1 if the person lives in an urban area. Similarly, five city indicators are included for the five most populous cities in Indonesia and Kenya, respectively. Control variables include age, age squared, years of education, years of education squared and an indicator for being female. When also including individual fixed effects in columns 4, the control variables are reduced to only age squared. All regressions are clustered at the individual level. Robust standard errors are in parentheses, *** p<0.01, ** p<0.05, * p<0.1.


Appendix (for online publication)

Figure A1: Log GDP per Capita and Agricultural Share



Notes: Table source data is from Gollin, Lagakos, and Waugh (2014), Online Appendix Table 4. Kenya (KEN) and Indonesia (IDN) are highlighted.


Figure A2: Agricultural Share and Agricultural Productivity Gap




Notes: Table source data is from Gollin, Lagakos, and Waugh (2014), Online Appendix Table 4. Kenya (KEN) and Indonesia (IDN) are highlighted.

Figure A3: Event Study of Urban Migration for Urban Survivors


(A) Indonesia






Notes: Event study coefficients reported in top half of figure separately for “survivors” and “not-survivors.” “Survivor” status is defined as having no rural observations from period zero (when the individual moved an urban area) to the period of interest, corresponding exactly to the survivor rate graph on the lower half of the figure. Survivor coefficients (black line in the top half) obtained by interacting a survivor indicator with post-event time indicators described in Section IV.D; “not-survivor” coefficients (grey line in the top half) is the event time indicator interacted with one minus the survivor indicator. Panel A reports results for Indonesia, and Panel B reports results for Kenya. Please refer to Figure 4 notes for additional details on included control variables and computation of survivor rates.

Figure A4: Event Study of Rural Migration



Notes: Figure uses data on individuals in the IFLS who are born in urban areas. Event time indicator variables defined analogously to Figure 4 except with respect to individuals' first observed rural move. Coefficients multiplied by negative 1 to interpret difference in earnings as an urban premium. Sample includes 710 movers with wage observations at the time of move and one period prior; 118 individuals report wages five years later. Please refer to Figure 4 notes for additional details on included control variables and computation of survivor rates.


Figure A5: Event Study of Rural Migration for Survivors




Notes: Figure uses data on individuals in the IFLS who are born in urban. Event study coefficients reported in top half of figure separately for “survivors” and “not-survivors.” “Survivor” status is defined as having no urban observations from period zero (when the individual moved a rural area) to the period of interest, corresponding exactly to the survivor rate graph on the lower half of the figure. Survivor coefficients (black line in the top half) obtained by interacting a survivor indicator with post-event time indicators described in Section IV.D; “not-survivor” coefficients (grey line in the top half) is the event time indicator interacted with one minus the survivor indicator. Panel A reports results for Indonesia, and Panel B reports results for Kenya. Please refer to Figure 4 notes for additional details on included control variables and computation of survivor rates.

Figure A6: Event Study of Migration to Capital City

(A) Indonesia





Notes: Event study coefficients reported in top half of panels A and B separately for moves to the Jakarta and Nairobi, respectively (black line) or other urban areas (grey line). Estimates obtained by interacting a capital city indicator with all event time indicators described in Section IV.D; “Not Jakarta” and “Not Nairobi” are event time indicators interacted with one minus the respective capital city indicator. In the sample, 53 initial moves were to Jakarta, and 133 initial moves were to Nairobi. Please refer to Figure 4 notes for additional details on included control variables and computation of survivor rates.

Table A1: Correlates of Rural Migration—Indonesia (Born Urban)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Primary Ed.	-0.275*** (0.020)					-0.216*** (0.028)	-0.195*** (0.021)
Secondary Ed.		-0.192*** (0.013)				-0.127*** (0.020)	-0.132*** (0.015)
College			-0.195*** (0.019)			-0.0498 (0.029)	-0.0943*** (0.021)
Female				-0.00275 (0.013)		-0.0170 (0.017)	-0.0163 (0.013)
Raven's Z-score					-0.0549*** (0.009)	-0.0190* (0.009)	
Constant	0.789*** (0.018)	0.628*** (0.009)	0.565*** (0.007)	0.539*** (0.009)	0.576*** (0.009)	0.846*** (0.027)	0.798*** (0.020)
Observations	5634	5634	5634	5634	3359	3359	5634

Notes: This table is a rural migration analog of Table 3. Each cell represents a regression coefficient with an indicator for being a rural migrant as the dependent variable. The sample is restricted to individuals born in urban areas. Please see notes from Table 3.

Table A2: Correlates of Employment in Non-Agriculture—Indonesia (Born Urban)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Primary Ed.	0.188*** (0.020)					0.138*** (0.030)	0.155*** (0.020)
Secondary Ed.		0.101*** (0.007)				0.0728*** (0.009)	0.0801*** (0.007)
College			0.0529*** (0.007)			-0.0190 (0.011)	-0.0106 (0.007)
Female				0.0409*** (0.007)		0.0592*** (0.009)	0.0498*** (0.007)
Raven's Z-score					0.0280*** (0.005)	0.0130* (0.005)	
Constant	0.752*** (0.019)	0.876*** (0.006)	0.916*** (0.004)	0.906*** (0.005)	0.919*** (0.005)	0.731*** (0.029)	0.724*** (0.020)
Observations	5634	5634	5634	5634	3359	3359	5634

Notes: This table is a analogous to Table 4 but is estimated on individuals born in urban areas. Please see notes from Table 4.

Table A3: Correlates of Meals Eaten—Kenya

	(1)	(2)	(3)
	Log Consumption	Log Earnings	Log Wage
Log(Meals)	0.194* (0.090)	0.278*** (0.065)	0.221*** (0.066)
Number of observations	1062	4693	4333

Notes: Each cell reports a regression coefficient with the log of meals as the independent variable; dependent variables listed in the header of the table. These regressions do not have the sample restrictions found in Table 2. Log of household per capita consumption in column 1 available only for a subset of individuals from KLPS 3. Robust standard errors clustered by individual reported below in parentheses, *** p<0.01, ** p<0.05, * p<0.1.

Table A4: Kenya Urban Towns

	Population	Percentage of Urban Individual-Months
Nairobi	3,133,518	43.4
Mombasa	938,131	14.4
Busia	61,715	6.6
Nakuru	307,990	4.7
Kisumu	409,928	4.3
Eldoret	289,380	2.6
Kakamega	91,768	1.3
Kitale	106,187	1.2
Bungoma	81,151	1.1
Naivasha	181,966	1.0
Gilgil	35,293	0.5
Thika	139,853	0.3
Other	.	18.6

Notes: This table presents a list of reported towns from urban individual-month observations. Urban status is defined based on respondent answering that they live in a large town or city. Column 3 lists the fraction of individual months in analysis from a particular town. The source for town populations is the 2009 Kenya Census.

Table A5: Non-Agricultural/Agricultural Gap in Earnings using Alternative Definition of Agriculture

(A) Indonesia

	Dependent variable: Log Earnings (in IDR)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
				Log Wage				Log Wage
Only non-agricultural employment	0.536*** (0.016)	0.277*** (0.014)	0.048** (0.023)	0.000 (0.028)				
Any non-agricultural employment					0.627*** (0.017)	0.308*** (0.016)	0.209*** (0.024)	0.064** (0.030)
Individual fixed effects	N	N	Y	Y	N	N	Y	Y
Control variables and time FE	N	Y	Y	Y	N	Y	Y	Y
Number of observations	115897	115897	115897	115897	115897	115897	115897	115897
Number of individuals	18211	18211	18211	18211	18211	18211	18211	18211

(B) Kenya

	Dependent variable: Log Earnings (in KSh)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
				Log Wage				Log Wage
Only non-agricultural employment	0.764*** (0.070)	0.561*** (0.066)	0.226*** (0.077)	0.171* (0.090)				
Any non-agricultural employment					0.816*** (0.077)	0.592*** (0.073)	0.394*** (0.098)	0.074 (0.120)
Individual fixed effects	N	N	Y	Y	N	N	Y	Y
Control variables and time FE	N	Y	Y	Y	N	Y	Y	Y
Number of observations	128439	128439	128439	128439	128439	128439	128439	128439
Number of individuals	4537	4537	4537	4537	4537	4537	4537	4537

Notes: Panel A uses data from the IFLS, and Panel B uses data from the KLPS. The table repeats some of the analyses shown in Tables 5 and 6 with alternate definitions of non-agriculture. In the first “Only non-agricultural employment,” an individual-time is considered *agricultural* if any of their jobs are *agricultural*, and non-agricultural otherwise. In the second, “Any non-agricultural employment,” an individual-time is considered *non-agricultural* if any of their jobs are *non-agricultural*, and agricultural otherwise. For columns 4 and 8, the dependent variable is the log of earnings divided by hours worked. Control variables include log hours, log hours squared, age, age squared, years of education, years of education squared and an indicator for being female. When also including individual fixed effects in columns 3, 4, 7 and 8, the control variables are reduced to only age squared. All regressions are clustered at the individual level. Robust standard errors are in parentheses, *** p<0.01, ** p<0.05, * p<0.1.

Table A6: Non-Agricultural/Agricultural Gap in Earnings Within Rural Areas

(A) Indonesia

	Dependent variable: Log Earnings (in IDR)			
	(1)	(2)	(3)	(4) Log Wage
Non-agricultural employment	0.509*** (0.017)	0.278*** (0.016)	0.178*** (0.026)	0.054* (0.033)
Individual fixed effects	N	N	Y	Y
Control variables and time FE	N	Y	Y	Y
Number of observations	91765	91765	91765	91765
Number of individuals	14495	14495	14495	14495

(B) Kenya

	Dependent variable: Log Earnings (in KSh)			
	(1)	(2)	(3)	(4) Log Wage
Non-agricultural employment	0.353*** (0.091)	0.248*** (0.087)	-0.049 (0.115)	-0.199 (0.143)
Individual fixed effects	N	N	Y	Y
Control variables and time FE	N	Y	Y	Y
Number of observations	61139	61139	61139	61139
Number of individuals	2827	2827	2827	2827

Notes: Panel A uses data from the IFLS, and Panel B uses data from the KLPS. The table repeats some of the analyses shown in Table 5, but restricts the sample to observations where the individual resides in rural areas. For column 4, the dependent variable is the log of earnings divided by hours worked. Control variables include log hours, log hours squared, age, age squared, years of education, years of education squared and an indicator for being female. When also including individual fixed effects in columns 3 and 4, the control variables are reduced to only age squared. All regressions are clustered at the individual level. Robust standard errors are in parentheses, *** p<0.01, ** p<0.05, * p<0.1.

Table A7: Gap in Earnings for those Aged 30 or Younger, Indonesia

	Dependent variable: Log Earnings (in IDR)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
				Log Wage				Log Wage
Non-agricultural employment	0.443*** (0.024)	0.246*** (0.022)	0.067* (0.040)	-0.084* (0.050)				
Urban					0.458*** (0.019)	0.249*** (0.017)	0.037 (0.023)	-0.022 (0.029)
Individual fixed effects	N	N	Y	Y	N	N	Y	Y
Control variables and time FE	N	Y	Y	Y	N	Y	Y	Y
Number of observations	37263	37263	37263	37263	37263	37263	37263	37263
Number of individuals	9932	9932	9932	9932	9932	9932	9932	9932

Notes: This table uses data from the IFLS. The table repeats some of the analyses shown in Tables 5 and 6 but restricts the sample to observations where the individual is aged 30 years or fewer to allow better comparability to the KLPS sample. For columns 4 and 8, the dependent variable is the log of earnings divided by hours worked. Control variables include log hours, log hours squared, age, age squared, years of education, years of education squared and an indicator for being female. When also including individual fixed effects in columns 3, 4, 7 and 8, the control variables are reduced to only age squared. All regressions are clustered at the individual level. Robust standard errors are in parentheses, *** p<0.01, ** p<0.05, * p<0.1.

Table A8: Gap in Wage Earnings

(A) Indonesia

	Dependent variable: Log Wage Earnings (in IDR)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
				Log Wage				Log Wage
Non-agricultural employment	0.929*** (0.021)	0.451*** (0.020)	0.127*** (0.028)	-0.074* (0.038)				
Urban					0.655*** (0.019)	0.297*** (0.016)	-0.002 (0.016)	-0.044** (0.019)
Individual fixed effects	N	N	Y	Y	N	N	Y	Y
Control variables and time FE	N	Y	Y	Y	N	Y	Y	Y
Number of observations	65412	65412	65412	65412	65412	65412	65412	65412
Number of individuals	12408	12408	12408	12408	12408	12408	12408	12408

(B) Kenya

	Dependent variable: Log Wage Earnings (in KSh)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
				Log Wage				Log Wage
Non-agricultural employment	0.912*** (0.090)	0.648*** (0.080)	0.369*** (0.097)	0.097 (0.137)				
Urban					0.741*** (0.057)	0.617*** (0.051)	0.227*** (0.081)	0.163* (0.089)
Individual fixed effects	N	N	Y	Y	N	N	Y	Y
Control variables and time FE	N	Y	Y	Y	N	Y	Y	Y
Number of observations	92800	92800	92800	92800	92800	92800	92800	92800
Number of individuals	3844	3844	3844	3844	3844	3844	3844	3844

Notes: Panel A uses data from the IFLS, and Panel B uses data from the KLPS. The table repeats some of the analyses shown in Tables 5 and 6, but instead of using all available earnings as the dependent variable, this table only includes earnings from wage employment. For columns 4 and 8, the dependent variable is earnings from wage employment divided by hours worked. Control variables include log hours, log hours squared, age, age squared, years of education, years of education squared and an indicator for being female. When also including individual fixed effects in columns 3, 4, 7 and 8, the control variables are reduced to only age squared. All regressions are clustered at the individual level. Robust standard errors are in parentheses, *** p<0.01, ** p<0.05, * p<0.1.

Table A9: Gap in Self-Employment Earnings

(A) Indonesia

	Dependent variable: Log Self-Employment Earnings (in IDR)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
				Log Wage				Log Wage
Non-agricultural employment	0.370*** (0.021)	0.269*** (0.021)	0.092** (0.040)	0.099* (0.050)				
Urban					0.604*** (0.028)	0.412*** (0.026)	0.019 (0.033)	0.024 (0.042)
Individual fixed effects	N	N	Y	Y	N	N	Y	Y
Control variables and time FE	N	Y	Y	Y	N	Y	Y	Y
Number of observations	66412	66412	66412	66412	66412	66412	66412	66412
Number of individuals	10252	10252	10252	10252	10252	10252	10252	10252

(B) Kenya

	Dependent variable: Log Self-Employment Earnings (in KSh)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
				Log Wage				Log Wage
Non-agricultural employment	0.016 (0.175)	-0.022 (0.161)	0.158 (0.151)	0.147 (0.148)				
Urban					0.680*** (0.116)	0.457*** (0.109)	-0.077 (0.132)	-0.217 (0.289)
Individual fixed effects	N	N	Y	Y	N	N	Y	Y
Control variables and time FE	N	Y	Y	Y	N	Y	Y	Y
Number of observations	35316	35316	35316	35316	35316	35316	35316	35316
Number of individuals	1217	1217	1217	1217	1217	1217	1217	1217

Notes: Panel A uses data from the IFLS, and Panel B uses data from the KLPS. The table repeats some of the analyses shown in Tables 5 and 6, but instead of using all available earnings as the dependent variable, this table only includes earnings from self-employment. For columns 4 and 8, the dependent variable is earnings from self-employment divided by hours worked. Control variables include log hours, log hours squared, age, age squared, years of education, years of education squared and an indicator for being female. When also including individual fixed effects in columns 3, 4, 7 and 8, the control variables are reduced to only age squared. All regressions are clustered at the individual level. Robust standard errors are in parentheses, *** p<0.01, ** p<0.05, * p<0.1.

Table A10: Alternative Samples Kenya

(A) Subsistence agriculture included also if not main decision maker

	Dependent variable: Log Wage Earnings (in KSh)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
				Log Wage				Log Wage
Non-agricultural employment	0.792*** (0.075)	0.574*** (0.072)	0.321*** (0.098)	0.161 (0.116)				
Urban					0.847*** (0.051)	0.674*** (0.047)	0.247*** (0.070)	0.130 (0.082)
Individual fixed effects	N	N	Y	Y	N	N	Y	Y
Control variables and time FE	N	Y	Y	Y	N	Y	Y	Y
Number of observations	128428	128428	128428	128428	128428	128428	128428	128428
Number of individuals	4536	4536	4536	4536	4536	4536	4536	4536

(B) Subsistence agriculture not included

	Dependent variable: Log Wage Earnings (in KSh)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
				Log Wage				Log Wage
Non-agricultural employment	0.650*** (0.083)	0.454*** (0.079)	0.171 (0.105)	0.007 (0.129)				
Urban					0.827*** (0.052)	0.659*** (0.048)	0.230*** (0.070)	0.113 (0.084)
Individual fixed effects	N	N	Y	Y	N	N	Y	Y
Control variables and time FE	N	Y	Y	Y	N	Y	Y	Y
Number of observations	126091	126091	126091	126091	126466	126102	126102	126102
Number of individuals	4410	4410	4410	4410	4438	4413	4413	4413

Notes: Panels A and B use data from the KLPS, described in Section 3. Panel A also includes productivity from subsistence agriculture if the individual is not the main decision maker for the agricultural activity. Panel B excludes all data from subsistence agriculture. For columns 4 and 8, the dependent variable is total earnings divided by hours worked. Control variables include log hours, log hours squared, age, age squared, years of education, years of education squared and an indicator for being female. When also including individual fixed effects in columns 3, 4, 7 and 8, the control variables are reduced to only age squared. All regressions are clustered at the individual level. Robust standard errors are in parentheses, *** p<0.01, ** p<0.05, * p<0.1.

Table A11: Gap in Food and Non-Food Consumption, Indonesia

(A) Food Consumption						
	Dependent variable: Log Food Consumption (in IDR)					
	(1)	(2)	(3)	(4)	(5)	(6)
Non-agricultural employment	0.353*** (0.010)	0.174*** (0.009)	0.022 (0.018)			
Urban				0.342*** (0.011)	0.199*** (0.011)	0.033 (0.021)
Individual fixed effects	N	N	Y	N	N	Y
Control variables and time FE	N	Y	Y	N	Y	Y
Number of observations	38280	38280	38280	38280	38280	38280
Number of individuals	19695	19695	19695	19695	19695	19695
(B) Non-Food Consumption						
	Dependent variable: Log Non-Food Consumption (in IDR)					
	(1)	(2)	(3)	(4)	(5)	(6)
Non-agricultural employment	0.877*** (0.015)	0.511*** (0.014)	0.041 (0.026)			
Urban				0.836*** (0.017)	0.479*** (0.016)	0.036 (0.030)
Individual fixed effects	N	N	Y	N	N	Y
Control variables and time FE	N	Y	Y	N	Y	Y
Number of observations	38280	38280	38280	38280	38280	38280
Number of individuals	19695	19695	19695	19695	19695	19695

Notes: Both panels use data from the IFLS. Panels A and B repeat the consumption analyses shown in Table 8, broken down by food and non-food consumption respectively. Please refer to Table 8 for further details. Control variables include age, age squared, years of education, years of education squared and an indicator for being female. When also including individual fixed effects in columns 3 and 6, the control variables are reduced to only age squared. All regressions are clustered at the individual level. Robust standard errors are in parentheses, *** p<0.01, ** p<0.05, * p<0.1.

Table A12: Gap in Consumption (Main Analysis Sample), Indonesia

(A) Total Consumption

	Dependent variable: Log Consumption (in IDR)					
	(1)	(2)	(3)	(4)	(5)	(6)
Non-agricultural employment	0.536*** (0.012)	0.267*** (0.011)	-0.004 (0.020)			
Urban				0.495*** (0.013)	0.273*** (0.012)	0.032 (0.023)
Individual fixed effects	N	N	Y	N	N	Y
Control variables and time FE	N	Y	Y	N	Y	Y
Number of observations	32544	32544	32544	32544	32544	32544
Number of individuals	17157	17157	17157	17157	17157	17157

(B) Food Consumption

	Dependent variable: Log Food Consumption (in IDR)					
	(1)	(2)	(3)	(4)	(5)	(6)
Non-agricultural employment	0.355*** (0.011)	0.159*** (0.010)	-0.003 (0.020)			
Urban				0.324*** (0.012)	0.183*** (0.011)	0.037 (0.023)
Individual fixed effects	N	N	Y	N	N	Y
Control variables and time FE	N	Y	Y	N	Y	Y
Number of observations	32544	32544	32544	32544	32544	32544
Number of individuals	17157	17157	17157	17157	17157	17157

(C) Non-Food Consumption

	Dependent variable: Log Non-Food Consumption (in IDR)					
	(1)	(2)	(3)	(4)	(5)	(6)
Non-agricultural employment	0.860*** (0.016)	0.477*** (0.016)	0.000 (0.030)			
Urban				0.787*** (0.017)	0.446*** (0.016)	0.029 (0.032)
Individual fixed effects	N	N	Y	N	N	Y
Control variables and time FE	N	Y	Y	N	Y	Y
Number of observations	32544	32544	32544	32544	32544	32544
Number of individuals	17157	17157	17157	17157	17157	17157

Notes: All regressions use data from the IFLS. This table repeats the analyses shown in Table 8 and A11 using the main analysis sample, which excludes individual-year observations without earnings measures. Thus, the sample size is smaller than in Table 8. Control variables include age, age squared, years of education, years of education squared and an indicator for being female. When also including individual fixed effects in columns 3 and 6, the control variables are reduced to only age squared. All regressions are clustered at the individual level. Robust standard errors are in parentheses, *** p<0.01, ** p<0.05, * p<0.1.

Table A13: Unemployment and Job Search Behavior, Kenya

(A) Unemployment

	Dependent Variable: Unemployment or Subsistence Agriculture			Dependent Variable: Unemployment		
	(1)	(2)	(3)	(4)	(5)	(6)
Urban	-0.032*** (0.012)	-0.026** (0.012)	-0.001 (0.021)	0.156*** (0.008)	0.142*** (0.008)	0.116*** (0.012)
Individual fixed effects	N	N	Y	N	N	Y
Control variables and time FE	N	Y	Y	N	Y	Y
Mean dependent variable	0.297	0.297	0.297	0.079	0.079	0.079
Number of observations	10879	10879	10879	10879	10879	10879
Number of individuals	6757	6757	6757	6757	6757	6757

(B) Search Behavior

	Dependent variable: Total Hours Job Search		
	(1)	(2)	(3)
Urban	1.242*** (0.144)	1.221*** (0.150)	1.794*** (0.266)
Individual fixed effects	N	N	Y
Control variables and time FE	N	Y	Y
Mean dependent variable	1.887	1.887	1.887
Number of observations	10879	10879	10879
Number of individuals	6757	6757	6757

Notes: Panel A reports urban gaps in unemployment. The first three columns define an individual as being unemployed if they are searching for work and have no income from wage or salary employment. The second three columns define an individual as being unemployed if they are searching for work and have no income from wage, salary, or proceeds from subsistence agriculture reported in the agricultural module. Sample sizes differ from analysis of wage gaps because questions about job search are contemporaneous to the time of the survey and are not retrospective. The dependent variable in Panel B is the number of hours a person reports to be searching for work; this variable equals 0 if the person is not searching for work. Like Panel A, data was only collected on search behavior contemporaneous to the time of the survey and thus sample sizes are smaller. Control variables include age, age squared, years of education, years of education squared and an indicator for being female. When also including individual fixed effects in columns 3 and 6, the control variables are reduced to only age squared. All regressions are clustered at the individual level. Robust standard errors are in parentheses, *** p<0.01, ** p<0.05, * p<0.1.