

NBER WORKING PAPER SERIES

THE AGRICULTURAL PRODUCTIVITY GAP

Douglas Gollin
David Lagakos
Michael E. Waugh

Working Paper 19628
<http://www.nber.org/papers/w19628>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue
Cambridge, MA 02138
November 2013

We thank Lisa Starkman, Luisa Majnoni and Michael Jaskie for excellent research assistance. For helpful comments we thank Robert Barro and Larry Katz (the editors), four anonymous referees, Francesco Caselli, Matthias Doepke, Berthold Herrendorf, Larry Huffman, Remi Jedwab, Per Krusell, Kiminori Matsuyama, Maggie McMillan, Ed Prescott, Victor Rios Rull, Richard Rogerson, Mark Rosenzweig, Todd Schoellman, Rachel Soloveichik, Francis Teal, Aleh Tsyvinski, Dietz Vollrath and Alwyn Young, as well as seminar participants at Arizona State, Cambridge, Dartmouth, Oxford, Northwestern, Penn State, Toronto, UCSD and Yale, and conference participants at the NBER Summer Institute Economic Growth Meeting, the NBER Conference on Macroeconomics Across Time and Space, the AEA Annual Meetings (San Diego), the SED Annual Meetings (Ghent), the Conference on Economic Growth (Toronto), the Conference on Economic Growth and Cultural Change (Munich), and the World Bank's Annual Bank Conference on Development Economics. Portions of this paper were written while Gollin was on leave at the Yale School of Forestry and Environmental Studies. Financial support from the International Growth Centre is gratefully appreciated. All potential errors are our own. This version supersedes an earlier version with the title "The Agricultural Productivity Gap in Developing Countries". 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2013 by Douglas Gollin, David Lagakos, and Michael E. Waugh. 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.

The Agricultural Productivity Gap
Douglas Gollin, David Lagakos, and Michael E. Waugh
NBER Working Paper No. 19628
November 2013
JEL No. E01,E24,J61,O11,O13,O15,O18,O41,R11

ABSTRACT

According to national accounts data, value added per worker is much higher in the non-agricultural sector than in agriculture in the typical country, and particularly so in developing countries. Taken at face value, this “agricultural productivity gap” suggests that labor is greatly misallocated across sectors. In this paper, we draw on new micro evidence to ask to what extent the gap is still present when better measures of sector labor inputs and value added are taken into consideration. We find that even after considering sector differences in hours worked and human capital per worker, as well as alternative measures of sector output constructed from household survey data, a puzzlingly large gap remains.

Douglas Gollin
University of Oxford
Department of International Development
Queen Elizabeth House
3 Mansfield Road
Oxford OX1 3TB
United Kingdom
Douglas.Gollin@qeh.ox.ac.uk

Michael E. Waugh
New York University
Kaufman Management Center
44 West Fourth St. 7-160
New York, NY 10012-1126
mwaugh@stern.nyu.edu

David Lagakos
Department of Economics, 0508
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093
and NBER
lagakos@ucsd.edu

An online appendix is available at:
<http://www.nber.org/data-appendix/w19628>

1. Introduction

In almost every country in the world, the share of employment in agriculture is higher than the share of value added coming from agriculture. As a simple matter of arithmetic, this implies that value added per worker is higher in the non-agricultural sector than in agriculture. According to data from national income and product accounts, this “agricultural productivity gap” is around a factor three on average, and even higher in developing countries.

These large agricultural productivity gaps suggest important implications for understanding aggregate productivity. With minimal assumptions on production technologies, they suggest that labor is misallocated across sectors, and particularly so in developing countries. By re-allocating workers out of agriculture, where the value of their marginal product is low, and into other activities, aggregate output would increase even without increasing the amount of inputs employed in production. These gains could be particularly large in developing countries, where the agricultural productivity gaps and shares of employment in agriculture are largest.

In this paper, we take a step back and ask whether these gaps are still present when better measures of inputs and outputs are taken into consideration. In other words, we ask how much of the agricultural productivity gaps are artifacts of mismeasurement, as opposed to real differences in output per worker. Several existing studies have argued that these measurement issues may be first-order: [Caselli and Coleman \(2001\)](#), for example, argue that agriculture workers have relatively lower human capital than other workers; [Gollin, Parente, and Rogerson \(2004\)](#) suggest that agriculture output maybe underestimated due to home production; and [Herrendorf and Schoellman \(2013\)](#) claim that measurement error in agricultural value added data is prevalent even across U.S. States. Despite these concerns, the literature does not offer a clear answer to how important these measurement issues are in practice in developing countries.

To answer this question, we construct a new database from population censuses and household surveys and use it to provide better measures of agricultural productivity gaps in 151 countries of all income levels. We organize our analysis around possible biases that could affect value added per worker in the denominator (employment) and in the numerator (value added). We then use our new database to perform a sequence of adjustments to the data on agriculture’s shares of employment and value added. In the first set of adjustments, we construct measures of average hours worked and human capital per worker by sector. We consider several measures of human capital, including adjustments for sector differences in returns to schooling and experience. We find that taking into consideration sector differences in hours and human capital per worker jointly reduces the size of the average agricultural productivity gap by roughly one third overall and by one half in developing countries. We also find, however, that the remaining gaps are still large and particularly so in developing countries.

We then construct alternative measures of value added by sector using household income surveys from ten developing countries. Our data come from the World Bank's Living Standards Measurement Studies (LSMS), which are designed explicitly to obtain measures of household income and expenditure. They allow us to compute, among other things, the market value of all output—whether ultimately sold or consumed at home—produced by the households. We find that the gaps in value added per worker by sector implied by these household income surveys are similar in magnitude to those found in the national accounts. This suggests that mismeasurement of value added in national accounts is unlikely to account for the agricultural productivity gaps implied by national accounts data, at least in these developing countries.

We then consider a set of other potential explanations for the gaps, including sector differences in labor's share in production, potential discrepancies between income per worker and income per household, and urban-rural differences in amenities. We conclude that the agricultural productivity gap is unlikely to be completely explained by any of the measurement issues we address in the paper. What this suggests, we argue, is that a better understanding is needed of why so many workers remain in the agricultural sector, given the large residual productivity gaps with the rest of the economy. Understanding these gaps will help determine, in particular, whether policy makers in the developing world should pursue policies that encourage movement of the workforce out of agriculture.

As a first step in this direction, we present correlations of the adjusted agricultural productivity gaps with a set of country characteristics that we expect may be linked to the remaining gaps. The measures that we focus on relate to geographical features, measures of institutional quality, and measures of labor mobility. We find that the adjusted agricultural productivity gaps are correlated with all these variables, and in ways that suggest possible avenues for future research. For example, the adjusted gaps are larger, all else equal, when the quality of a country's institutions are low. This finding is consistent with the idea that the risk of expropriation might make workers less inclined to leave behind the security of existing living and working arrangements. While all we present are correlations, we feel that this evidence is suggestive of possible mechanisms that may help explain the measured productivity gaps.

We are not the first to point out the existence of large agricultural productivity gaps. [Lewis \(1955\)](#), for example, noted that in developing countries that "there is usually a marked difference between incomes per head in agriculture and in industry." The fact that the agriculture productivity gaps are most prevalent in poor countries was first shown by [Kuznets \(1971\)](#), and was later documented in richer detail by [Gollin, Parente, and Rogerson \(2002\)](#). These differences in sectoral productivity were viewed as critical by early development economists. [Rosenstein-Rodan \(1943\)](#), [Lewis \(1955\)](#), and [Rostow \(1960\)](#) viewed the development process as fundamentally linked to the reallocation of workers out of agriculture and into "modern" economic activities. Recently, [McMillan and Rodrik \(2011\)](#) argue that reallocations of workers to

the most productive sectors – taking the underlying productivity data at face value – would raise aggregate output per worker substantially in many countries.

Our contribution is to attempt to account for the gaps using richer data on labor and value added at the sector level than in any prior study. In particular, our paper is the first to make use of household survey-based measures of schooling attainment by sector and hours worked by sector in the cross-section of countries. Furthermore, we are the first to compare sector productivity levels computed from “macro” data, based on the national accounts, to those implied by “micro” data, based on household surveys of income. Our work is similar in this regard to that of [Young \(2012\)](#), who compares growth rates computed from national accounts data to those computed from household survey data in a set of developing countries.

Our work complements a related literature trying to understand why there is more cross-country variation in *physical* productivity in agriculture than in the non-agricultural sector ([Caselli, 2005](#); [Restuccia, Yang, and Zhu, 2008](#); [Chanda and Dalgaard, 2008](#); [Vollrath, 2009](#)). Proposed theories include distortions that reduce farm size ([Adamopoulos and Restuccia, Forthcoming](#)), selection of lower-ability workers into agriculture ([Lagakos and Waugh, 2013](#)), and low intermediate usage in agriculture ([Restuccia, Yang, and Zhu, 2008](#); [Donovan, 2013](#)). The current paper differs in that it compares the value of output produced by sector within each country.

The papers most closely related to ours are those of [Young \(2013\)](#) and [Herrendorf and Schoellman \(2013\)](#). [Young \(2013\)](#) uses cross-country, micro-level survey data to study urban-rural consumption gaps and connects these measures with migration outcomes. While he focuses on concepts that are distinct but related concepts to ours (real consumption vs. value added output per worker) and sectoral definitions (urban/rural vs. agricultural/non-agricultural), he finds gaps in consumption between urban and rural areas that are similar in magnitude to our findings. Moreover, by connecting these gaps with migration outcomes, he builds on the model of [Lagakos and Waugh \(2013\)](#) and provides compelling evidence suggesting that these gaps are well explained by selection on unobservable skill. We view [Young’s \(2013\)](#) results as providing an important path toward understanding the productivity gaps discussed in our paper.

[Herrendorf and Schoellman \(2013\)](#) study why agricultural productivity gaps are so large in most U.S. states. A key difference in the conclusions of our paper and theirs is that [Herrendorf and Schoellman \(2013\)](#) argue that systematic under-reporting of agriculture income (data that are collected in the U.S. from farmers’ tax returns) is a major factor in accounting for the low relative productivity of agriculture. Whether this type of mismeasurement is equally at play in countries outside the United States, and developing countries in particular, is an open question. The main similarity is that both studies find that sector differences in human capital per worker explain a substantial fraction of the gaps.

Finally, our work relates closely to the recent literature on misallocation and its role in explaining cross-country differences in total factor productivity and output per worker. Seminal examples of this line of research are Restuccia and Rogerson (2008) and Hsieh and Klenow (2009), who focus on the misallocation of capital across firms; or Caselli and Feyrer (2007) who study the misallocation of capital across countries. In contrast, we focus on the potential misallocation of workers *across sectors*. Our focus on the divide between the agricultural and non-agricultural sectors is important because the vast majority of workers in developing countries are in agriculture, suggesting that misallocation between these two sectors may be the most relevant source of sectoral misallocation.

2. The Agricultural Productivity Gap—Theory

We begin by asking what standard neoclassical theory implies for the agricultural productivity gap. Consider a neoclassical two-sector model featuring Cobb-Douglas production functions in the agricultural and non-agricultural sectors, free labor mobility across sectors and competitive labor markets. Assume further that the labor share in production is given by θ in each sector (an assumption we view as reasonable, but discuss later in Section 6.4), so that the production functions are:

$$Y_a = A_a L_a^\theta K_a^{1-\theta} \quad \text{and} \quad Y_n = A_n L_n^\theta K_n^{1-\theta} \quad (1)$$

where subscripts a and n denote agriculture and non-agriculture, and variables Y , L and K represent output, labor input and capital (and land) input.

The assumption of free labor mobility implies that the equilibrium wage for labor across the two sectors is the same. The assumption of competitive labor markets implies that workers are paid the value of their marginal product, and that firms hire labor up to the point where the marginal value product of labor equals the wage. Thus, marginal value products are also equalized. Combined with the production functions in (1), and letting p_a denote the relative price of Y_n , the ratio of marginal value products, and average value products, is

$$\frac{VA_n/L_n}{VA_a/L_a} = \frac{Y_n/L_n}{p_a Y_a/L_a} = 1. \quad (2)$$

We call this ratio of value added per worker in non-agriculture to agriculture the *agricultural productivity gap*, or APG. If the condition in (2) is not met, then this suggests that workers are misallocated relative to the competitive benchmark. For example, if the ratio of value added per worker between non-agriculture and agriculture is larger than one, there would seem to be an incentive for workers to move from agriculture to non-agriculture, simultaneously pushing up the marginal product of labor in agriculture and pushing down the marginal product of labor in non-agriculture. This process should tend to move the sectoral average products towards

equality.

An important point to note in condition (2) is that it does not depend on any assumptions about other factor markets. In particular, labor productivity should be equalized across sectors even in the presence of market imperfections that lead to misallocation of other factors of production. For example, capital markets could be severely distorted, but firm decisions and labor flows should nevertheless drive marginal value products—and hence value added per worker—to be equated. Thus, the model implies that if (2) does not hold in the data, the explanation must lie in either measurement problems related to labor inputs or value added, or in frictions of some kind in the labor market—nothing else.

Writing equation (2) in terms of agriculture’s share of employment and output gives:

$$\frac{(1 - y_a)/(1 - \ell_a)}{y_a/\ell_a} = 1, \quad (3)$$

where $y_a \equiv VA_a/(VA_a + VA_n)$ and $\ell_a \equiv L_a/(L_a + L_n)$. In other words, the ratio of each sector’s share in value added to its share in employment should be the same in the two sectors.

The relationship in (3) is the lens through which we look at the data. Under the (minimal) conditions outlined above, we first ask if the condition in (3) holds in cross-country data. One way to think about this exercise is along the lines of the work of [Restuccia and Rogerson \(2008\)](#) and [Hsieh and Klenow \(2009\)](#), who focus on the equality of marginal products of capital across firms; or the work of [Caselli and Feyrer \(2007\)](#), who study the equality of marginal products of capital across countries. Here, in contrast, we focus on the value of the marginal product of labor across sectors.

Later, in Section 6.4, we explore what happens when we allow labor shares to differ across sectors. We argue there that equality in labor shares turns out not to be a bad assumption given available estimates, and hence delay discussing this possibility until later.

3. The Agricultural Productivity Gap — Measurement and Data

In this section, we ask whether, in national accounts data, value added per worker is equated across sectors, as predicted by the theory. We begin with a detailed—perhaps tedious—description of how the national income and product accounts approach the measurement of agricultural value added and how national labor statistics quantify the labor force in agriculture. We conclude that while there are some inevitable difficulties in the implementation of these measures, there is no reason *ex ante* to believe that the data are flawed.

With these measurement issues clear, we then present the “raw,” or unadjusted, agricultural productivity gaps using aggregate value added and employment data. We show that the gap

is around a factor of three in the typical country, well above the prediction of the theory.

3.1. Conceptual Issues and Measurement: National Accounts Data

The statistical practices discussed below are standard for both rich and poor countries, but there are particular challenges posed in measuring inputs and outputs for the agricultural sector in developing countries. A major concern is that aggregate measures of economic activity and labor allocation in poor countries may be flawed—and may in fact be systematically biased by problems associated with household production, informality, and the large numbers of producers and consumers who operate outside formal market structures. Given these concerns, we focus on the conceptual definitions and measurement approaches used in the construction of national accounts data and aggregate labor measures.

To illustrate the potential problems, consider the example of Uganda, a country in which household surveys and agricultural census data show that as much as 80 percent of certain important food crops (cassava, beans, and cooking bananas) may be consumed within the farm households where they are grown. Most households are effectively in quasi-subsistence; the government reports that even in the most developed regions of the country, nearly 70 percent of households make their living from subsistence agriculture. In the more remote regions of the country, over 80 percent of households are reported as deriving their livelihoods from subsistence farming (Uganda Bureau of Statistics 2007b, p. 82).

Given these concerns, it is possible that value added measures will, by design or construction, omit large components of economic activity. As we discuss below this is not the case. Although value added may be measured with error, the conceptual basis for value added measurement is clear and well-defined.

3.2. Measurement of Value Added in Agriculture

Perhaps surprisingly, the small scale and informality of agricultural production in poor countries does not mean that their output goes largely or entirely unmeasured in national income and product accounts. At a conceptual level, home-consumed production of agricultural goods does fall within the production boundary of the United Nations (UN) System of National Accounts (SNA), which is the most widely used standard for national income and product accounts. The SNA specifically includes within the production boundary “the production of all agricultural goods for sale or own final use and their subsequent storage” (FAO (1996), p. 21), along with other forms of hunting, gathering, fishing, and certain types of processing. Within the SNA, there are further detailed instructions for the collection and management of data on the agricultural sector.

How is the measurement of these activities accomplished? Accepted practice in developing

countries is to measure the area planted and the yield of most crops, which can be surveyed at the national level, and to subtract off the value of important purchased intermediate inputs.¹ There are also detailed guidelines for estimating the value of output from animal agriculture and other activities, as well as for the consideration of inventory. Specific procedures also govern the allocation of output to different time periods.² Allowances are made for harvest losses, spoilage, and intermediate uses of the final product (e.g., crop output retained for use as seed). The final quantities estimated in this way are then valued at “basic prices,” which are defined as “the prices realized by [farmers] for that produce at the farm gate excluding any taxes payable on the products and including any subsidies.”

Although it is difficult to know how consistently these procedures are followed in different countries, the guidelines for constructing national income and product accounts are clear, and they apply equally to subsistence or quasi-subsistence agriculture as to commercial agriculture. Furthermore, there is no reason to believe that national income and product accounts for poor countries do an intrinsically poor job of estimating agricultural value added (as opposed to the value added in services or manufacturing, where informality is also widespread). Nor is there reason to believe that agricultural value added in poor countries is consistently underestimated, rather than overestimated.³

3.3. Measurement of Labor in Agriculture

Potential mismeasurement of labor in agriculture is another key concern. Because agriculture in poor countries falls largely into the informal sector, there are no detailed data on employment of the kind that might be found in the formal manufacturing sector. There are unlikely to be payroll records or human resources documentation. Most workers in the agricultural sector are unpaid family members and own-account workers, rather than employees. For example, in Ethiopia in 2005, 97.7 percent of the economically active population in agriculture consisted of “own-account workers” and “contributing family workers,” according to a national labor force survey made available through the International Labour Organization. A similar data set for Madagascar in 2003 put the same figure at 94.6 percent.

¹For some crops, only area is observed; for others, only production is observed. The guidelines provide detailed information on the estimation of output in each of these cases.

²The national accounting procedures also provide guidance on the estimation of intermediate input data. In the poorest countries, there are few intermediate inputs used in agriculture. But, conceptually, it is clear that purchased inputs of seed, fertilizer, diesel, etc., should be subtracted from the value of output. Data on these inputs can be collected from “cost of cultivation” or “farm management” surveys, where these are available, but the FAO recommends that these data “should be checked against information available from other sources,” such as aggregate fertilizer consumption data. Similar procedures pertain to animal products.

³Nevertheless, many development economists find it difficult to believe that national income accounts data for developing countries can offer an accurate picture of sectoral production. We revisit these concerns later in Section 5, where we construct alternative measures of value added by sector using household survey data from ten developing countries. Although these data have their own limitations, as we discuss later, we find that the large agricultural productivity gaps are present in these household survey data as well.

The informality of the agricultural sector may tend to lead to undercounting of agricultural labor. But a bigger concern is over-counting—which would lead to misleadingly low value added per worker in the sector. Over-counting might occur in at least two ways. First, some people might be mistakenly counted as active in agriculture simply because they live in rural areas. In principle, this should not happen; statistical guidelines call for people to be assigned to an industry based on the “main economic activity carried out where work is performed.” But in some cases, it is possible that enumerators might count individuals as farmers even though they spend more hours (or generate more income) in other activities. In rural areas in developing countries (as well as in rich countries), it is common for farmers to work part-time in other activities, thereby smoothing out seasonal fluctuations in agricultural labor demand. This might include market or non-market activities, such as bicycle repair or home construction.

A second way in which over-counting might occur is if hours worked are systematically different between agriculture and non-agriculture. In this situation, even if individuals are assigned correctly to an industry of employment, the hours worked may differ so much between industries that we end up with a misleadingly high understanding of the proportion of the economy’s labor that is allocated to agriculture. We explore this possibility directly in Section 4.1, below.

Note that this type of over-counting would affect sectoral productivity comparisons only if hours worked differ systematically across sectors—so that workers in non-agriculture supply more hours, on average, than workers in agriculture. At first glance, it might seem obvious that this is the case, but much of non-agricultural employment in poor countries is also informal. Many workers in services and even in manufacturing are effectively self-employed, and labor economists often argue that informal non-agricultural activities represent a form of disguised unemployment in poor countries, with low hours worked. To return to the Ethiopian data, in 2005, 88.4 percent of the *non-agricultural* labor force consisted of own-account workers and family labor. Thus, the predominance of self-employment and family business holds across sectors. If there are important differences in hours worked across sectors, we cannot simply assume that they result from differences in the structure of employment.

A final way in which over-counting of labor in agriculture might occur is if human capital per worker were higher in non-agriculture than in agriculture. If this were true, we would be overestimating the effective labor input in agriculture compared to non-agriculture. In this case, the underlying real differences in sectoral productivity would be smaller than the measured APGs. We address these possibilities directly in Section 4.4, to follow.

3.4. Raw Agricultural Productivity Gap Calculations

With these measurement issues clear, the following sections describe the sample of countries and our data sources, and then presents the “raw,” or unadjusted, agricultural productivity

gaps.

3.5. The Sample and Data Sources

Our sample includes all countries for which data on agriculture's shares of employment and value added are available. We restrict attention to countries with data from 1985 or later, though the vast majority of countries have data from 2000 or later. We end up with a set of 151 countries which have broad representation from all geographic regions and per-capita income levels. We exclude any countries in which employment shares are not based on nationally representative surveys, or in which agriculture's employment share is less than one percent.⁴

Our data on agriculture's share of employment come from the International Labour Organization (ILO). The underlying source for these data are censuses of population or labor force surveys conducted by the countries' statistical agencies. One advantage of using surveys based on samples of individuals or households is that they include workers in informal arrangements and the self-employed. Surveys of establishments or firms, in contrast, often exclude informal or self-employed producers from their sample.

Workers are defined as any individuals supplying labor for the production of goods within the boundary of the national income accounts (FAO (1996)). There is no minimum threshold for hours worked. This definition includes all workers who are involved in producing final or intermediate goods, including home-consumed agricultural goods. In general, employed workers are classified into sectors by their main industry of employment.

Our data on agriculture's share of value added come from the United Nations (UN) National Account Statistics. The underlying sources for these data are the national income and product accounts from each country. Industry classifications are made using the International Standard Industrial Classification System (ISIC). In all cases, these data are expressed in current-year local currency units.⁵

3.6. Raw Agricultural Productivity Gaps

Table 1 reports summary statistics for the raw APGs for all 151 countries and broken down by quartile of the income distribution. We refer to these as raw APGs because, unlike the calcu-

⁴For most countries, there are only a few years for which we can compute agricultural productivity gaps historically. The reason is that nationally representative household surveys are conducted infrequently, particularly in developing countries. Using the few countries for which data are available, the average APG in the years before 1985 is roughly similar in magnitude to the average APG in the period since 1985

⁵An alternative would be to use a single set of international comparison prices to value the agricultural output of each country. This would be relevant if we were making comparisons of real (i.e., physical) agricultural output per worker across countries, as in Caselli (2005), Restuccia, Yang, and Zhu (2008), Vollrath (2009) or Lagakos and Waugh (2013). In the current paper, however, we are interested in comparing the value of output produced per worker across sectors within each country.

Table 1: Raw Agricultural Productivity Gaps

	All Countries	Quartile of Income Distribution			
		Q1	Q2	Q3	Q4
10th Percentile	1.3	1.0	1.3	1.0	1.2
Median	2.6	1.7	2.7	2.8	4.3
Mean	3.5	2.0	3.2	3.4	5.6
90th Percentile	6.8	4.0	6.6	7.1	12.5
Number of Countries	151	38	38	38	37

Note: Income quartiles are determined using 2005 PPP GDP per capita. Q1 is the richest quartile and Q4 is the poorest quartile. The raw agricultural productivity gaps are defined as the ratio of value added per worker in the non-agricultural sector to value added per worker in the agricultural sector, without any adjustments to the underlying value added or employment data.

lations that follow, they do not incorporate any adjustments (e.g., for hours worked). The first data column describes the APG distribution for the entire sample of 151 countries. Across all countries, the mean APG is 3.5 and the median is 2.6, implying that value added per worker is roughly three times higher in non-agriculture than in agriculture. Moreover, the gaps are large for some countries: at the 90th percentile of the distribution, the gap is 6.8. These gaps are consistent with the results of [Young \(2013\)](#), who finds urban-rural consumption gaps of around four for a set of developing countries.

The remaining columns of Table 1 list the APGs by quartile of the distribution of 2005 PPP GDP per capita, with Q1 being the richest quartile. Looking first at the means, one can see that the APGs tend to be smaller in the richest countries than in the poorest; a similar pattern holds for the median. The average gap in the richest quartile of the income distribution (Q1) is 2.0, and the median is 1.7. In the poorest quartile (Q4), the average gap is 5.6, while the median is 4.3. Other metrics show that APGs are systematically larger in the poorest countries of the world than in the richest. A regression of log APGs on log GDP per capita yields a slope coefficient of -0.3 and is precisely estimated. Thus, while APGs are, on average, large everywhere, they are particularly large in the developing world.

Relative to the discussion in Section 2, it is abundantly clear that the data are not consistent with equation (3), which would give an APG of one. The raw data suggest very large departures from parity in sectoral productivity levels in all countries of the world, and particularly so in the poorest countries of the world.

Differences of this magnitude are striking. If we take these numbers literally, they raise the possibility of large misallocations between sectors within poor countries. Are such large disparities

plausible? Do these numbers reflect underlying gaps in real productivity levels and living standards? Or do they largely reflect flawed measurements of labor inputs and value added? In the following sections, we discuss the new data that we bring to bear on the question, and consider a number of ways in which mismeasurement may occur. We will also compare the magnitude of these possible mismeasurements with the observed gaps in productivity.

4. Improved Measures of Labor Inputs by Sector

In this section, we report the results of efforts to adjust the productivity gaps to account for potential differences in the quantity and quality of labor inputs across sectors. We base this analysis on a new database that we constructed, which contains sector-level data on average hours worked and average years of schooling for a large set of developing countries. We construct our data using nationally-representative censuses of population and household surveys, with underlying observations at the individual level.

One part of our data comes from International Integrated Public Use Microdata Series (I-IPUMS), from which we use micro-level census data from 56 countries around the world. We also get data on schooling attainment by sector from 51 countries from the Education Policy and Data Center (EPDC), which is a public-private partnership of the U.S. Agency for International Development (USAID) and the Academy for Educational Development. From a number of other countries we get schooling and hours worked from the World Bank's LSMS surveys of households. For a number of developed countries we get data from the Cross-National Equivalent File (Frick, Jenkins, Lillard, Lipps, and Wooden, 2007). The remainder of the data comes from individual survey data and published tables from censuses and labor force surveys conducted by national statistical agencies. Table 2 in the Online Appendix details the sources and data used in each of the countries.

4.1. Sector Differences in Hours Worked

We now ask whether the sectoral productivity gaps are explained by differences across sectors in hours worked. We find that in most of the countries for which we have data on hours worked, there are only modest differences in hours worked by sector; on average, workers in non-agriculture supply around 1.1 times more hours than workers in agriculture. Thus, hours-worked differences are unlikely to be the main cause of the large APGs we observe.

We measure hours worked for all employed workers, whether they are self-employed, employed in wage work, or a mix of both. The typical survey asks hours worked in a reference period of one to two weeks prior to the survey, although some report average hours worked in the previous year. We classify people as workers in either agriculture or non-agriculture, according to their main reported economic activity.

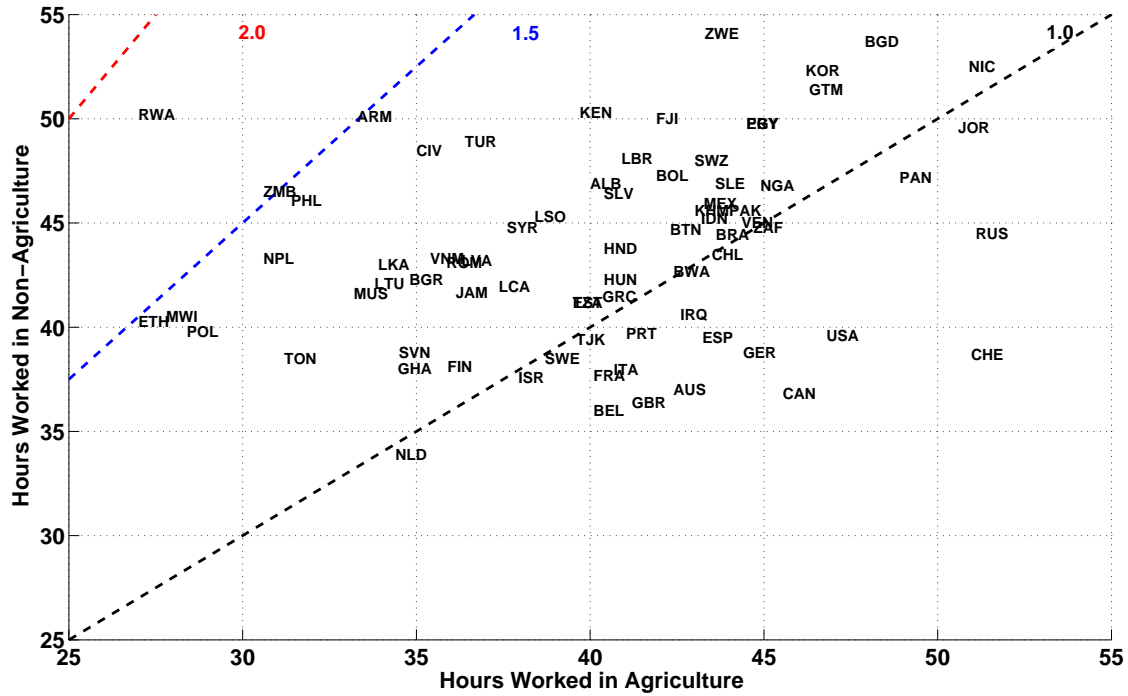


Figure 1: Average Hours Worked by Sector

For some countries, we cannot obtain measures of hours by agricultural or non-agricultural employment, but we are able, instead, to use hours worked by urban-rural status. Table 2 in the Online Appendix lists the countries for which we use urban-rural status to construct our hours measures.⁶ Using urban-rural status in some countries represents a potential limitation of our data, as the non-agricultural (agricultural) workforce and urban (rural) workforce do not correspond exactly to one another. However, in those countries for which we can measure average hours by both urban-rural status and agriculture-non-agriculture status, the two give similar average hours measures.

Figure 1 shows hours worked in non-agriculture, plotted against hours worked in agriculture, for each of the countries with available data. The 45-degree line, marked 1.0, corresponds to a situation in which average hours worked are identical in the two sectors. Similarly, the other two lines represent a factor of 1.5 and 2.0 differences in hours worked. Most of the observations are clustered closely around the 1.0 line, and the vast majority are below the 1.5 line, meaning that differences in hours worked across sectors are generally modest. An arithmetic average across all countries gives a factor 1.1 difference in hours worked in non-agriculture compared to agriculture.

⁶Our results change very little when using average hours among both employed and unemployed workers, classifying the latter according to their previous sector of employment or urban-rural status if previous sector is unavailable.

Hiding behind this result is a systematic relationship between the ratio of hours worked and the level of development. In Figure 1, rich countries lie mostly below the 45-degree line, whereas poor countries lie well above the 45-degree line. In countries in the top quartile of the income distribution, hours in agriculture are on par with those in non-agriculture, on average. For those in the bottom quartile of the income distribution, the ratio of hours worked is about 1.3. As shown in Table 1, these poor countries also have relatively large APGs. So while hours-worked differences overall do not seem to explain much of the APGs on average (as that would require an average ratio of around 3.5), in some countries lower hours worked in agriculture seem to be an important part of their large measured gaps.

4.2. Seasonality

One concern with these average-hours-worked statistics is potential bias due to seasonality. Seasonality in agricultural work is well known to make average hours higher during some months of the year (e.g., planting and harvesting) than others. If surveys systematically miss these busy periods, they may bias downward the average hours worked by agricultural workers. In contrast, if surveys are conducted systematically during harvest and planting times, they may bias hours upward.

One way to address possible bias from seasonality is to take a closer look at hours-worked data from the set of surveys for which survey dates are drawn representatively throughout the calendar months of the year. We can then compute average hours by sector by month, to assess the magnitude of seasonal fluctuations in hours worked. The challenge with this approach is that there are very few countries for which the data allow this comparison.

Figure 2 below plots hours worked by sector in each month of the calendar year for the five countries for which these calculations are feasible. Each subplot represents a country, and for each country, the dashed line shows average hours worked in non-agriculture, while the solid line represents average hours worked in agriculture. In most countries, one can see clear evidence of seasonality in agricultural hours worked. However, hours worked in non-agriculture are almost always higher than those in agriculture throughout the year, as in the conclusion of the previous section. Thus, for these countries at least, our results do not seem to be driven by seasonality.

For most other countries, the data on hours worked come from single-day snapshots—a standard approach used in censuses and labor force surveys. These snapshot studies typically ask people about hours worked during a reference period (often one or two weeks prior to the “census day”), but the census day is the same for everyone in the country. An advantage of surveys like this is that they minimize the potential for double-counting or under-counting people who change address during the course of the year. Snapshot surveys might seem, at first glance, to

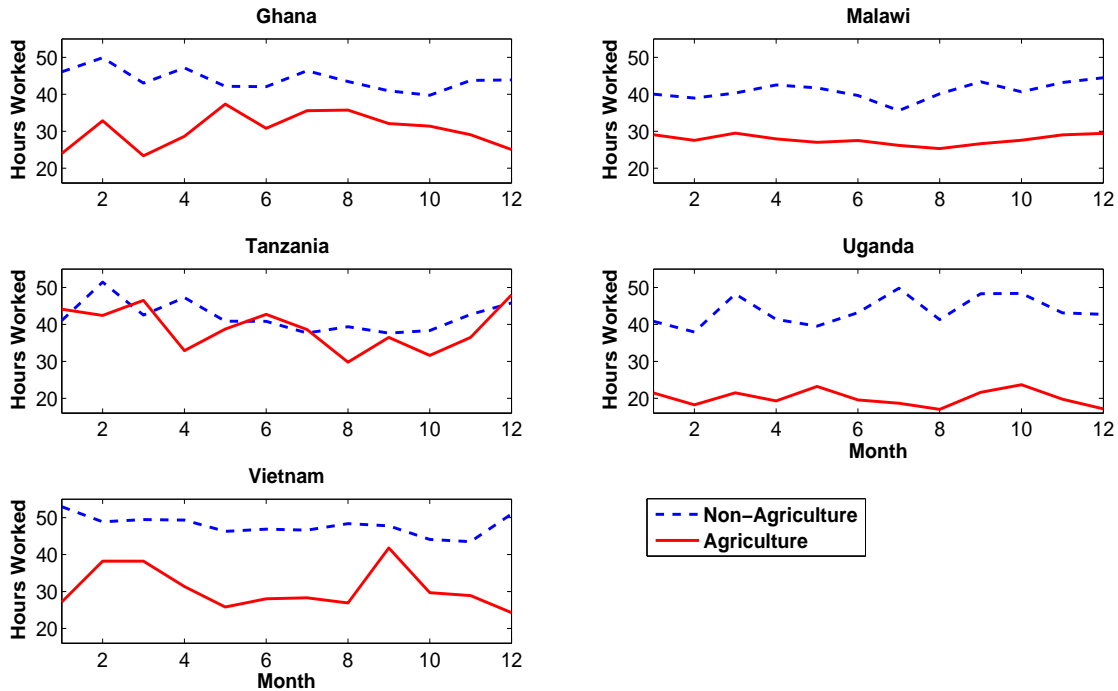


Figure 2: Average Hours Worked By Month by Sector

be particularly susceptible to a bias in hours worked. If the reference period falls during the harvest or planting period (or during a slow period in an agricultural off-season), we might expect these surveys to result in estimates of agricultural hours worked that are biased upwards (or downwards).

To address this concern, we compile data on the cropping seasons of each of the 47 countries for which we rely on snapshot surveys of hours worked (most of which are census data from IPUMS). For each country, we take the two most important staple foods produced—defined for our purposes as the vegetable-sourced foods that are the most important sources of calories consumed in the country, taken from FAO Food Balance Sheets. We then turn to a variety of sources to identify the key planting and harvesting dates for each staple food crop.⁷

We find that census days do not appear to be disproportionately located in seasons of peak labor demand or slack labor demand. To make this point formally, we calculate that for the two most important crops, the median country has either planting or harvest activities taking place for 9 months of the year; the mean value is 8.6 months. On average, then, nearly three quarters of the year is devoted to the planting or harvest of one of the two staple crops. After matching

⁷In identifying the two most important foods in each country, we ignore major staple foods that are largely imported. In general, we focus on starch staples, whether grains or roots and tubers; but for some countries, vegetable oils such as olive oil were important sources of calories, and in these cases, we treated these as food staples.

the census days to the crop calendars, we find that approximately three quarters of the census days fall within a planting or harvest period (36 of 47 countries, or 78 percent). This suggests that our findings regarding average hours worked are unlikely to be an artifact of seasonality.⁸

4.3. Hours Worked: A Further Breakdown

In the calculations above, we classify workers by their primary sector of employment and then attribute all their labor hours to that sector. A potential concern is that individuals classified as agricultural (non-agricultural) work a substantial fraction of their hours in non-agricultural (agricultural) activities. For example, suppose that individuals in agriculture in fact devote a large fraction of their hours to non-agricultural activities. In this case, we would be overcounting their hours worked in agriculture, leading to an underestimate of average labor productivity in agriculture. For this to be quantitatively important, it would need to be the case that a substantial fraction of hours is misallocated in this fashion.

To explore this possibility, we analyze individual-level data from LSMS household surveys for six countries with the appropriate data. Table 2 shows the results of this analysis. In this table, we show the hours worked in each sector by workers classified as agricultural or non-agricultural. As noted above, the classification of workers is based on their primary sector of employment. However, the LSMS data for these countries allow us to measure the hours worked by individuals across sectors in each of their economic activities.

These measures of hours worked show that to an overwhelming degree, individuals classified as working in agriculture do in fact allocate their time to agricultural activities; similarly, workers classified as non-agricultural allocate almost all of their time to non-agricultural activities. In all of these cases except that of the 1998 Ghana LSMS, we find that agricultural-classified workers devote 95 percent or more of their hours to agriculture; and in every case, we find that workers classified as non-agricultural devote at least 94 percent of their hours to non-agricultural activities.⁹

Although we have not carried out these painstaking calculations for all the countries with available micro data, we feel comfortable on the basis of the available evidence that the procedure we are using for calculating hours worked by sector accurately reflects the allocation of hours

⁸Of course, this does not imply that seasonality in agriculture is not important for individual households. Instead, our results suggest that at the aggregate (national) level, seasonality in agricultural hours is likely to be much less pronounced than seasonality at the household level due to the multiple planting and harvesting seasons for different crops in different regions on the census day.

⁹At first glance, these numbers appear to be inconsistent with the stylized fact that non-farm income represents an important source of earnings for rural households. In fact, our results are entirely consistent with that stylized fact. The reason is simply that “rural” and “agricultural” are different categories. In all of the micro data sets that we have examined, there are substantial fractions of rural households that are classified as non-agricultural. For example, in the 1998 Ghana LSMS data, 29.2 percent of rural workers are classified as non-agricultural, and 44.5 percent of rural income was non-agricultural.

Table 2: Hours Worked: A Further Breakdown

Country	Worker Classification	Sector of Hours Worked	
		Agriculture	Non-agriculture
Cote d'Ivoire (1988)	Agriculture	35.1	1.0
	Non-agriculture	0.7	49.2
Ghana (1998)	Agriculture	28.8	3.7
	Non-agriculture	2.0	30.6
Guatemala (2000)	Agriculture	47.6	1.3
	Non-agriculture	0.8	49.1
Malawi (2005)	Agriculture	26.4	1.4
	Non-agriculture	2.3	38.2
Tajikistan (2009)	Agriculture	39.5	0.1
	Non-agriculture	0.1	39.3
Uganda (2009)	Agriculture	18.7	2.1
	Non-agriculture	1.8	43.3

Note: This table reports average hours worked by sector for workers classified as agricultural workers and for workers classified as non-agricultural workers. Workers are classified by sector according to their primary sector of employment. Hours are classified by sector of job for each of the workers' jobs.

at the individual level.

4.4. Sector Differences in Human Capital

We next ask to what extent sectoral differences in human capital per worker can explain the observed APGs. We show that while schooling is lower, on average, among agricultural workers, the differences are not large enough to fully explain the measured gaps.

Our calculations in this section are related to those of [Vollrath \(2009\)](#), who also attempts to measure differences in average human capital between workers in agriculture and non-agriculture based on school enrollment data. While both sets of calculations have their limitations, ours improves on those of [Vollrath \(2009\)](#) in several dimensions. Most important, our calculations come from nationally representative censuses or surveys with direct information on educational attainment by individual. We also end up with estimates for a much larger set of countries. Finally, we attempt to adjust for quality differences in schooling across sectors.

As before, we compute average years of schooling by sector using household survey and census data. We measure hours worked for all individuals employed in either the agricultural or non-agricultural sector, when the data allow; otherwise, we use all economically active workers or classify workers by urban-rural status. In almost every country, we impute years of schooling using data on educational attainment (such as "primary school completed"), which is what


is available to us. Table 2 in the Online Appendix details the data source for each country. These imputations are, of course, likely to yield noisy measures of years of schooling, since a category such as “some secondary schooling completed” (for example) could correspond to several values for years of schooling. However, in all countries where we impute schooling, we do so in exactly the same way for non-agricultural and agricultural workers. Thus, the noise should in principle not systematically bias our measures of average years of schooling by sector.

As a further check of the data, we compared our aggregate average years of schooling versus Barro and Lee’s (2013) measure of educational attainment. The two measures are very similar, with a correlation of 0.89. Moreover, the average level is quite similar; in the cross-section of countries, average years of schooling is 8.25 in the Barro and Lee (2013) data set and 8.10 in our data. Overall, we view this result as reassuring in that our measures match up with Barro and Lee’s (2013) work.


Figure 3(a) shows our results for the 98 countries for which we constructed average years of schooling by sector. Again, the 45-degree line, marked 1.0, indicates equality in schooling levels, and the lines 1.5 and 2.0 represent those factor differences in years of schooling. As can be seen in the figure, in literally every country, average schooling is lower in agriculture than non-agriculture. Countries with the highest levels of schooling for individuals classified as working in agriculture tend to be closest to parity between the sectors. For example, the former Soviet block countries of Armenia, Kazakhstan, Uzbekistan, Georgia, and Ukraine have the highest schooling in agriculture and among the lowest ratios of non-agricultural to agricultural schooling. The ratios are generally higher for countries with less schooling among agriculture workers, with the lowest generally coming in francophone African countries. Mali, Guinea, Senegal, Chad and Burkina Faso have the lowest schooling for agricultural workers and among the highest ratios.

We are interested in the differences in human capital per worker that can be attributed to these differences in schooling. To turn years of schooling into human capital, we consider several different approaches. All of them assume that average human capital in sector j of country i can be expressed as $h_{j,i} = \exp(f(s_{j,i}; i))$, where $f(s_{j,i}; i) \equiv r_i \cdot s_{j,i}$, $s_{j,i}$ is average years of schooling in sector j of country i , and r_i is the return to each year of schooling in country i . Many macro studies simply assign a constant value to r_i across countries—assuming, for example, that each year of schooling increases wages by around ten percent (see, e.g., Klenow and Rodríguez-Clare (1997)). This assumption is based on the observation that Mincer returns to a year of schooling are roughly ten percent in countries of all income levels.¹⁰

¹⁰See Banerjee and Duflo (2005). Another approach, used by Hall and Jones (1999) and Caselli (2005), is to assume that there is some concavity in years of schooling, so that the first several years of schooling gives a higher return than subsequent years. When we use their approach we find results that are very close to those presented below.



(a) Years of Schooling by Sector



(b) Human Capital by Sector

Figure 3: Schooling and Human Capital by Sector

Figure 3(b) plots the results for human capital by sector using this approach. The resulting estimates of human capital by sector suggest that in virtually all countries, the average non-agricultural worker has between 1.0 and 1.5 times as much human capital as the average agricultural worker. The biggest ratios are still for the countries with the lowest human capital in both sectors, but the differences are less pronounced than those of schooling. This is simply because having (say) twice as many years of schooling implies having considerably less than twice as much human capital (see, e.g., the discussion of Mincer return estimates in [Banerjee and Duflo \(2005\)](#) and [Psacharopoulos and Patrinos \(2002\)](#)). The average across countries is a factor 1.3 difference in human capital of across the two sectors.

Similar to hours, this result masks some meaningful differences between rich and poor countries. For the countries in the top quartile of the income distribution, human capital is, on average, just 1.1 times larger in the non-agricultural sector than in agriculture. In contrast, for the countries in the bottom quartile, the difference between human capital is 1.4. Hence, in the countries where the APGs are the largest, we find that sector differences in human capital seems to be an important part of their large measured gaps.

4.5. Alternative Estimates of Returns to Schooling

One concern with the calculations in the previous section is that there are may be important differences across countries in the rates of return to schooling, and hence in the human capital accumulation of individuals with different years of schooling. To address this concern, we use country-specific estimates of the returns to schooling that have been compiled in three previous studies. Two of these sets of estimates can be traced to [Psacharopoulos and Patrinos \(2002\)](#), who generated a large list of country-specific rates of return, based on Mincer-type regressions. Based on these data, [Banerjee and Duflo \(2005\)](#) offered a modified set of estimates; an updated data set from the World Bank also provides estimates for some additional countries and some modifications to other numbers. Finally, a third set of country-specific estimates of returns to schooling comes from the work of [Schoellman \(2012\)](#). Unlike the other two data sources, [Schoellman \(2012\)](#) bases his estimates on the earnings of migrants to the United States, based on census data. Earnings are observed for migrants with different levels of education, allowing for estimates of country-specific rates of return to schooling.

We calculate sectoral differences in human capital per worker using all three sources of data on country-specific returns to education. Because these three data sets are incomplete in terms of country coverage, we can only calculate the sectoral differences for limited numbers of countries. The World Bank data and the [Banerjee and Duflo \(2005\)](#) data give essentially the same results; as a result, we report only the former. The data of [Schoellman \(2012\)](#) show lower returns to schooling in the poorest countries and thus generate different numbers for sectoral human capital levels.

Using the World Bank data, based on [Psacharopoulos and Patrinos \(2002\)](#), we find that sectoral differences in years of schooling translate into a level of human capital per worker that is 1.5 times higher in non-agriculture than in agriculture; in other words, each worker has 50 percent more human capital in non-agriculture. This compares to a figure of 1.3 when we use a constant ten percent rate of return to a year of schooling for all countries, as in our benchmark estimates. Using the estimates of [Schoellman \(2012\)](#), we find that the human capital per non-agricultural worker is, on average, 1.3 times higher than human capital per worker in agriculture.

To summarize our findings in this section, we find that there are substantial differences in human capital per worker across sectors arising through differences in schooling. Because education levels and educational attainment are almost universally lower in agriculture than in non-agriculture, we estimate that workers in the non-agricultural sector have 1.3 to 1.5 times as much human capital than those in agriculture, depending on the rate of return to schooling. This does appear to be an important source of differences in average labor productivity. However, these differences alone cannot account fully for the raw gaps observed in the data.

4.6. Adjusting Human Capital Using Literacy Rates

One limitation of the above analysis is that our procedure treats agricultural and non-agricultural workers with the same years of schooling as having the same levels of human capital. There are several good reasons to think that non-agricultural workers have more human capital conditional on the same amount of schooling. The quality of schooling in rural areas in many countries is below that of schooling in urban areas (see, e.g., [Williams \(2005\)](#) and [Zhang \(2006\)](#)). Children growing up in rural areas may have worse nutrition or worse health more generally, which leads them to miss school or absorb less while at school. They may also get less input from their parents, who may themselves be less educated, or devote fewer resources to education for their children. All of these factors imply that our estimates above may tend to overestimate the human capital level of agriculture workers by looking just at years of schooling completed.

To address this possible bias, we present a simple new method of measuring human capital differences by sector, using literacy data in addition to schooling data. The basic idea is that literacy, particularly in primary schools, is one of the main observable measures of human capital. Thus, literacy rates for workers by years of schooling completed in the two sectors are informative about how much human capital individuals have.

What we observe in our micro data are the literacy rates for non-agricultural and agricultural workers in country i conditional on having completed s years of schooling, which we denote $\ell_i^n(s)$ and $\ell_i^a(s)$ for $s = 0, 1, 2, \dots$. If the amount of human capital by year of schooling were the same for the two groups, then $\ell_i^n(s)$ and $\ell_i^a(s)$ would be the same (at least approximately) for




Figure 4: Literacy Rates by Years of Schooling, Uganda

each s . Instead, we find that in almost every country in our sample, $\ell_i^n(s) > \ell_i^a(s)$ for most or all values of s . In other words, literacy rates are higher for non-agricultural workers at most or all schooling levels.

Figure 4 illustrates the literacy data by sector for Uganda. The x -axis contains years of schooling completed and the y -axis shows the literacy rates $\ell_i^n(s)$ and $\ell_i^a(s)$ for the two sectors by years of schooling completed. Note that at each year of schooling completed, non-agricultural workers have literacy rates that are at least as high as those of agricultural workers, with the biggest difference coming for the lower years of schooling completed (particularly one year).

Our method is as follows. First we interpolate the literacy outcome data for agricultural workers and create a continuous literacy function of schooling: $\tilde{\ell}_i^a(s)$. This function, which for the case of Uganda is the dotted curve in Figure 4, allows us to evaluate literacy rates for agricultural workers for non-integer years of schooling. We then posit that, in country i , s years of schooling for agricultural workers are as effective as $s\gamma_i$ years of schooling for non-agricultural workers, and we set γ_i to the value that solves

$$\min_{\gamma} \sum_{s=1}^{\bar{s}} \left(\tilde{\ell}_i^n(\gamma s) - \tilde{\ell}_i^a(s) \right)^2 \quad (4)$$

where \bar{s} is the maximum years of schooling over which we minimize. In other words, we pick the value of γ that equates as closely as possible the literacy rates between agricultural workers with s years of schooling and non-agricultural workers with $s\gamma$ years of schooling. The limit \bar{s} is necessary because literacy rates in both sectors tend to be close to one hundred

percent for individuals with sufficiently high years of schooling, and thus literacy rates for these individuals are not informative about human capital differences by sector. We set the value of \bar{s} equal to five because literacy rates for all countries tend to be close to one hundred percent by then (as in year six and beyond in Figure 4).

In the example of Uganda, we find that $\gamma_{UGA} = 0.82$, meaning that agricultural workers with s years of schooling have the equivalent of 0.82 times as much human capital as non-agricultural workers with s years of schooling. Available data allow us to make similar calculations for 17 other developing countries.¹¹ The average estimate is 0.87. The range of all other estimates runs from a low of 0.62 in Guinea to a high of 0.95 in Bolivia. Mexico, Venezuela and Vietnam have other notably low estimates, all around 0.75. Only Tanzania has an estimate above one.

For the 17 countries, we make these adjustments to our human capital stocks, where $h_{a,i}^q = \exp(r_i \hat{\gamma}_i s_{a,i})$ for each country i . After the adjustments, we find that human capital is between 1.2 and 1.7 times higher in non-agriculture, with an average of 1.6. For comparisons, in our benchmark calculations, we find that on average human capital is 1.3 times higher in agriculture. Thus, relaxing the assumption that workers with the same years of schooling have the same amount of human capital in each sector makes a modest difference relative to the calculations of the previous section, at least using this methodology.

4.7. Human Capital from Experience

Human capital can be accumulated from experience, in addition to that arising from schooling. This may matter for our calculations if workers in agriculture tend to accumulate less human capital over their working years than workers in the non-agriculture sector. To address how important this may be in practice, we draw on estimates of wage returns to potential experience by sector from [Lagakos, Moll, Porzio, and Qian \(2013\)](#). They estimate the returns to potential experience by sector by running Mincer regressions for agricultural and non-agricultural workers in a set of twenty countries. They find that returns to potential experience are modestly lower among agricultural workers than others in most countries.

We define the human capital from schooling and experience among workers in country i and sector j as $h_{j,i} = \exp(f(s_{j,i}; i) + g(e_{j,i}; i, j))$ where $f(s_{j,i}; i)$ is (as before) the return to the average level of schooling in country i and sector j , and $g(e_{j,i}; i, j)$ is the return to the average level of experience in country i and sector j . The function $f(s_{j,i}; i)$ is computed for each country-sector pair exactly as before. The functions $g(\cdot; i, j)$ are computed using the quintic polynomials in potential experience estimated by [Lagakos, Moll, Porzio, and Qian \(2013\)](#), and the average levels of experience, $e_{j,i}$, that they compute.

¹¹These countries, and their estimated γ values, are Argentina (0.87), Bolivia (0.95), Brazil (0.95), Chile (0.92), Ghana (0.90), Guinea (0.62), Malaysia (0.93), Mali (0.89), Mexico (0.77), Panama (0.87), Philippines (0.80), Rwanda (0.88), Tanzania (1.25), Thailand (0.90), Uganda (0.82), and Venezuela (0.78).

What we find is that human capital stocks defined in this way are between 1.5 and 2.1 times as large in the non-agricultural sector as in agriculture. The average ratio is 1.9. These ratios are higher than those calculated before, reflecting the fact that returns to potential experience are, on average, somewhat lower among agricultural workers than among workers in other sectors. The average potential experience level is virtually identical in the two sectors and does play a significant role.

We conclude that human capital from experience is likely to be an important factor in understanding sector productivity gaps. Nevertheless, the magnitudes are still not large enough to explain the full size of the gaps. Our findings in this section parallel those of [Herrendorf and Schoellman \(2013\)](#), who find that across U.S. States and a smaller set of countries, human capital differences from experience, in addition to schooling, are important determinants of sector productivity gaps. The main caveat in our analysis is, of course, that we have evidence on sector returns to experience from only twenty countries. Hence, in our adjustments that follow, we stick to our human capital measures from schooling only.

4.8. Adjusted APGs

We now compute the “adjusted” agricultural productivity gaps, which take into consideration the sector differences in hours worked and human capital. We are able to do adjustments for both hours and human capital by sector for 72 countries. For each country, we construct the adjusted APG by dividing the raw APG by the ratio of hours worked and the ratio of human capital.

Figure 5 plots the adjusted APGs versus the raw APGs on the x -axis. The 45-degree line, marked 1.0, indicates that the APGs after adjustments are the same as the raw APGs; the line marked 0.50 indicates a 50 percent reduction in adjusted APGs relative to the raw APGs. For nearly all countries, the adjusted APGs lie below the 45-degree line and just above the 0.50 line. The first two columns of Table 3 provide more detail by reporting summary statistics of the raw and adjusted APG for countries. After all adjustments, the mean gap is 2.2 while the median is 1.9. The 10th percentile is 1.0 and the 90th percentile is 6.4. Comparing these adjustments to the raw APGs we see that the reductions reduced the gap by roughly one third, from around three down to around two.

The final four columns of Table 3 report the adjusted APGs by quartile of the PPP GDP per capita distribution. Comparing these results to those in Table 1 shows that our adjustments had a relatively larger impact for the least developed countries. For example, both the median and mean gap decreased by nearly 50 percent for those countries in the bottom quartile. In contrast, the median and mean gap decreased by about 15 percent for those countries in the top quartile. A result of the relatively larger reduction in APGs for poor countries is that there is




Figure 5: Raw and Adjusted Agricultural Productivity Gap

a substantially weaker relationship between the adjusted APGs and level of development. A regression of the log adjusted APGs on log GDP per capita yields a slope coefficient of -0.14 and is precisely estimated. For comparison, without the adjustments, the slope coefficient is nearly twice as large, at -0.26.

As can be seen in Table 3, even after adjustments for hours and human capital, the agricultural productivity gaps are still large in most countries, and particularly so in the developing countries. In the richest quartile of countries (Q1), the average gap is 1.7 while the median gap is 1.4. For the middle-income groups, mean and median gaps are in the ballpark of 2.0. For the poorest quartile (Q4), the median is 2.3 and the mean gap is 3.0.

We also compute a set of adjusted APGs for all 151 countries. For any country in which hours or schooling data are missing, we impute them as the average values of those of other countries in the same income quartile. We find very similar results to those in Table 3; after adjustments, the median and mean gap are 1.7 and 2.3. For the poorest quartile, their gaps are relatively larger, with a median and mean gap of 2.3 and 3.1.

Taken together, these numbers tell an important story: while our adjustments do help explain a sizable fraction the raw productivity differentials by sector, the remaining gaps are still quite large.

Table 3: Agricultural Productivity Gaps and All Adjustments

Measure	Raw APG	All Adjustments	All Adjustments by Quartile			
			Q1	Q2	Q3	Q4
10th Percentile	1.3	1.0	0.8	1.2	0.7	1.3
Median	3.1	1.9	1.4	2.0	2.1	2.3
Mean	3.5	2.2	1.7	2.1	1.9	3.0
90th Percentile	6.4	4.3	3.3	2.8	4.3	5.6
Number of Countries	72	72	18	16	18	20

Note: Income quartiles are determined using 2005 PPP GDP per capita. Q1 is the richest quartile and Q4 is the poorest quartile. The raw APG is defined as the ratio of value added per worker in the non-agricultural sector to value added per worker in the agricultural sector, without any adjustments to the underlying value added or employment data. The adjusted APG is defined as the the ratio of value added per worker in the non-agricultural sector to value added per worker in the agricultural sector after adjusting for average hours worked per worker and average human capital per worker.

5. Measures of Value Added by Sector from Micro Data

We now ask to what extent the agricultural productivity gaps implied by national accounts data are an artifact of mismeasurement of agricultural value added in national accounts data in practice. Mismeasurement may manifest itself in several ways. First, while national accounts data should, in theory, measure home production, agricultural output may in practice be underestimated due to home production, as argued by [Gollin, Parente, and Rogerson \(2004\)](#). Second, national accounts data may feature other types of bias disproportionately affecting agriculture. For example, [Herrendorf and Schoellman \(2013\)](#) argue that the agricultural productivity gaps present in the majority of U.S. states arise largely from mismeasurement due to the treatment of land and proprietors' income.

To answer this question, we use household survey data for a set of developing countries to construct new alternative measures of value added by sector. We focus on developing countries primarily because, as we've shown thus far, the gaps are largest there even after all of our adjustments. Moreover, concerns about mismeasurement due to home production are most relevant in developing countries, since home production is much more prevalent there. These micro-data allow us to compute income by economic activity for nationally representative samples of households, which we then aggregate to construct value added by agricultural and non-agricultural activity. A key feature of these data is that we observe a wide range of household enterprises that may or may not be accounted for properly in the national accounts.

What we find is that the shares of value added computed from the household data are similar to

those of the national accounts. As a result, the agricultural productivity gaps we compute using the household data are similar to those implied by the national accounts. While the household survey data are not without their own limitations, as we discuss below, these results suggest that the agricultural productivity gaps in developing countries are real, rather than artifacts of measurement problems with national accounts data.

5.1. Household Income Surveys

The household survey data we use come from the World Bank’s Living Standards Measurement Studies (LSMS). The LSMS surveys typically involve the collection of detailed data at the household (and individual) level on income, health, education, and other “outcome” measures; expenditure and consumption; labor allocation; asset ownership; and details on agricultural production, business operation, and other economic activities. The surveys undertaken in different countries do not always follow identical methodologies; nevertheless, substantial efforts have been made to allow for as much international comparability as possible, for example in the treatment of home production. In micro-development economics, data from these household surveys are generally seen as representing a high standard for data quality (Deaton (1997)).

We have ten developing countries for which we can measure value added by sector using household data. These are Armenia (1996), Bulgaria (2003), Cote d’Ivoire (1988), Guatemala (2000), Ghana (1998), Kyrgyz Republic (1998), Pakistan (2001), Panama (2003), South Africa (1993) and Tajikistan (2009). The Online Appendix provides more detail about each of the surveys. While small, our set of countries features a variety of geographic locales with four countries from Asia, two from the Americas, one from Europe, and three from Africa. It also features a wide variety of income levels, with three countries below \$2,000 PPP income per capita (Cote d’Ivoire, Ghana and Tajikistan), two between \$2,000 and \$5,000 (Kyrgyz Republic and Pakistan), two between \$5,000 and \$10,000 (Armenia and Guatemala), and three slightly above \$10,000 (Bulgaria, South Africa and Panama).

5.2. Measuring Value Added from Household Income Surveys

We construct value added in agriculture using the household survey data as follows. Letting i index a household, we define value added in agriculture to be:

$$VA_a = \sum_i y_{a,i}^{SE} + \sum_i y_{a,i}^L + \sum_i y_{a,i}^K \quad (5)$$

where $y_{a,i}^{SE}$, $y_{a,i}^L$ and $y_{a,i}^K$ represent self-employed agricultural income, agricultural labor income, and agricultural capital income of household i . Self-employed agricultural income of house-

hold i is defined as:

$$y_{a,i}^{SE} = \sum_{j=1}^J p_j (x_{i,j}^{home} + x_{i,j}^{market} + x_{i,j}^{invest}) - COSTS_{a,i}, \quad (6)$$

where j indexes all agricultural goods in the economy; p_j is the farm-gate price of good j ; and the three $x_{i,j}$ terms are the quantities of good j used for home consumption, market sales, and investment. In most cases, households with agricultural production report $x_{i,j}^{home}$, $x_{i,j}^{market}$ and $x_{i,j}^{invest}$ for each crop j in kilograms, and report p_j for all crops for which some sales were made. For other crops, the surveys report a local or regional average price. $COSTS_{a,i}$ is the cost of intermediate goods purchased, plus hired labor and rented capital (and land) used for production. Conceptually, $y_{a,i}^{SE}$ represents the value of all output produced by i net of any costs.

Agricultural labor income, $y_{a,i}^L$, is defined as all income paid in currency or in kind for labor services rendered by any member of the household in the agriculture sector. Wage income is measured at the individual level and then aggregated to the household level. Agricultural capital income, $y_{a,i}^K$, is defined as all income earned in currency or in kind for rental of housing, land or equipment, plus interest payments. Capital income is measured directly at the household level. Since it is virtually impossible to assign capital income to a particular sector, we assume that all capital income earned by agricultural households is agricultural, and all capital income earned by non-agricultural households is non-agricultural. We classify households as being either agricultural or non-agricultural based on which sector the majority of the household's workers are employed, and in the event of ties, which sector the majority of self-employment plus wage income comes from.

Value added in the non-agricultural sector is defined analogously as:

$$VA_n = \sum_i y_{n,i}^{SE} + \sum_i y_{n,i}^L + \sum_i y_{n,i}^K, \quad (7)$$

where $y_{n,i}^{SE}$, $y_{n,i}^L$ and $y_{n,i}^K$ represent self-employed non-agricultural income, non-agriculture employment income, and non-agricultural capital income of household i . Self-employed non-agricultural income is defined as:

$$y_{n,i}^{SE} = REV_{n,i} - COSTS_{n,i} \quad (8)$$

where $REV_{n,i}$ is self-reported revenues in non-agricultural businesses owned by household i , and $COSTS_{n,i}$ is any intermediate or factor cost incurred by these non-agricultural businesses. Non-agricultural labor income, $y_{n,i}^L$, and non-agricultural capital income $y_{n,i}^K$ are defined as above, only for the non-agricultural sector. For households with non-agricultural income, revenues and input costs are reported directly in all countries for which we have data.

Conceptually, our value added measures represent the total value of all payments made to factors of production applied to the production of output in each sector. Labor and capital income are unambiguous payment for labor and capital services used for production. The terms $y_{a,i}^{SE}$ and $y_{n,i}^{SE}$ represent payments made to entrepreneurs in the two sectors, and capture a mix of labor and capital income (see [Gollin \(2002\)](#)).

For each country, we compute value added by sector, as described above, and then compute agriculture's share of total value added. We then compute agriculture's employment share by classifying each workers by her primary industry of employment. Workers are defined to be all economically active adults aged 15 or older. Using these two shares for each country we can construct the ratio of value added per worker in non-agriculture to agriculture, which is essentially the "micro" analog of our raw APG measures.

Our calculations of value added from these micro data have several advantages relative to national accounts measures. First, they unambiguously include home production of agriculture. Second, they include "informal" income sources, such as small-scale self-employment or informal wage employment, which may not be included completely in national accounts data. Finally, they focus only on domestic households, and exclude (for example) large multinational national resource firms that may contribute a lot to domestic value added in non-agriculture without much effect on the income of domestic residents.

Our value added measures also have several limitations. First, given the relatively small sample sizes (usually several thousand), the surveys are unlikely to capture the income of the very highest income earners in the economy, who may be business owners or simply those with high wage income. Second, as is always true with surveys of income, self-employment income could be underreported. We worry particularly that non-agricultural self-employment income is under-reported, since non-agricultural business owners are typically asked directly to report their revenues (unlike in agriculture, where farm owners report physical quantities of output, crop by crop.) To the extent that this is the case, our non-agricultural value added measures may be biased downward, and, hence, so may our APG estimates. Finally, we were able to make these calculations for only ten countries, and each country's estimates are based on income data for only a few thousand people, which limits the generality of our findings.

5.3. Results: APGs from Household Income Surveys

Table 4 shows the results of the calculations of value added by sector for a set of ten developing countries. The first data column shows the share of workers in agriculture according to the micro survey data.¹² The second and third data columns show agriculture's share in value

¹²We do not report the "macro" employment shares in agriculture for two reasons. First, there is no conceptual difference between how we compute employment shares by sector and how sector employment shares are computed in aggregate statistics. Second, aggregate statistics on employment shares by are not available in most of the

Table 4: Micro and Macro Data and Agricultural Productivity Gaps

	Agriculture Share of				
	Employment	Value Added		APG	
		Micro	Macro	Micro	Macro
Armenia (1996)	34.2	36.8	32.8	0.9	1.1
Bulgaria (2003)	14.1	11.7	18.4	1.2	0.7
Cote d'Ivoire (1988)	74.3	32.0	42.1	4.7	4.0
Ghana (1998)	53.9	36.0	33.3	2.2	2.3
Guatemala (2000)	40.2	15.1	18.7	3.8	2.9
Kyrgyz Republic (1998)	56.9	39.5	39.3	2.0	2.0
Pakistan (2001)	46.9	25.8	22.6	2.5	3.0
Panama (2003)	27.0	7.8	11.8	4.4	2.7
South Africa (1993)	11.0	5.0	7.0	2.3	1.7
Tajikistan (2009)	41.0	24.7	30.1	2.1	1.6
Average	40.0	23.4	25.6	2.6	2.2

Note: “Micro” means calculated using LSMS household survey data. “Macro” means calculated using national accounts data. APGs are calculated using the shares of value added from micro and macro data, and the shares of employment from micro data.

added according to the macro data (the national accounts) and the micro data.

There are several points to take away from Table 4. First, the shares of value added in agriculture are fairly similar in both the macro and micro data with no apparent bias in either direction. For example, for some countries, such as Cote d’Ivoire, Tajikistan, and Bulgaria, the micro value added shares are larger than the macro shares. Whereas countries such as Armenia and Ghana have micro value added shares lower than the macro shares. The final row of Table 4 summarizes this result, showing that, on average, the micro value added share is nearly the same as the macro value added share.

Second—and most revealing—the micro employment shares (except for Bulgaria) are all larger than the micro value added shares, resulting in APGs greater than one. The final column summarizes this result by showing the micro APGs. The average micro APG is 2.2, with countries such as Cote d’Ivoire, Pakistan, Guatemala and Panama having the largest APGs. For comparison, the second to the last column reports the macro APGs, which use the macro value added shares and the micro employment shares. The average APG from the macro data is 2.6, and for the most part the same countries having the largest gaps in the micro data are those with the largest gaps in the macro data. Thus, both sets of data suggest large APGs, albeit with

years of our surveys.

somewhat smaller gaps computed from the micro data.

We conclude that, in spite of the differences in data and methodology between our calculations and those of the national accounts, the two measures provide surprisingly similar estimates of the size of the APGs in these developing countries. While countries may differ in the size of the employment and value added shares of agriculture, there are no countries for which micro and macro sources paint a substantially different picture of agriculture's share in aggregate value added.¹³ Thus, at least for these ten developing countries, substantial gaps in value added per worker by sector appear prominently in household survey data, and the magnitude of the gaps is similar to that found in the national accounts data.

6. Other Explanations

Thus far, we have argued that agricultural productivity gaps are still large even after adjusting for improved measures of labor input, and that measures of the shares of value added by sector are similar in household surveys and national accounts data. In this section, we discuss several possible explanations for these residual agriculture productivity gaps.¹⁴

6.1. Household Income and Expenditure by Sector

The theory of Section 2 assumes that workers supply labor to one particular sector and are indifferent between work in the two sectors. In reality, decisions are often made at the household level, and households often diversify income across different types of economic activity. Thus, it could be that households primarily involved in agriculture earn combined incomes from agricultural and non-agricultural activities that are, on average, equal to the total incomes of non-agricultural households. To address this question, we use our LSMS data to ask whether a gap exists between the average income per worker of agricultural households and non-agricultural households.

As above, we define households as being either agricultural or non-agricultural based on where the majority of their workers are employed (and in a tie, which sector is the source of the majority of their self-employment plus wage income.) We define income of a household i in sector

¹³One potential explanation for the similarity of the micro and macro numbers is that the underlying data sources are in fact the same or similar. In Tanzania, for example, the value added of agriculture is based largely on an extensive survey of rural households called the *Agricultural Sample Census*, combined with a second nationally representative household survey called the *Household Budget Survey*. Unfortunately, we do not know how much national statistical agencies in other developing countries base their value added estimates on household surveys.

¹⁴To be sure, a number of other measurement issues remain. As one example, the non-agricultural sector includes a number of industries—such as government services—in which output is valued at the cost of inputs and in which labor markets may not be fully competitive. If these sectors receive inflated wages, it will be misleadingly reflected in the data as high productivity.

s as

$$y_i = \sum_{j=a,n} y_{j,i}^{SE} + \sum_{j=a,n} y_{j,i}^L + y_{s,i}^K. \quad (9)$$

In other words, total income represents self-employed income from businesses in both sectors, wage income from both sectors, and capital income from the sector in which the household is classified. Note that for many households, at least one of the entries is zero. We define the total number of workers by household as the total number of economically active persons aged 15 or older.

In addition, we also measure the average expenditure per worker by sector. The rationale is that total household expenditure may provide a more accurate measure of income than direct measures of income, again due to underreporting of self-employed income (see, e.g., [Deaton \(1997\)](#) and [Ravallion \(2003\)](#).) The Online Appendix provides more detail about how expenditure is measured in each survey. To construct the measures of income per worker and expenditure per worker by sector, we use the same LSMS described above in Section 5.1.

Table 5 presents the results, with the last two columns reporting the ratio of income and expenditure per worker for non-agricultural households relative to agricultural households. For convenience, we also reproduce our measures of the APGs using the micro approach from Table 4. The results show that, for the most part, the gaps in income per worker are similar to the gaps in value added per worker. The average gap in income per worker is 2.1 relative to a 2.2 gap in value added per worker. The relative rankings are very similar as well. The countries with the largest APGs also have the largest gaps in income per worker.

Expenditures per worker data paint a similar picture—most countries exhibit gaps in expenditure per worker across sectors. The relative ranking across countries in expenditure gaps is similar to the ranking of gaps in income per worker and micro-APGs. However, the magnitudes of the expenditure gaps are often lower. For example, the average expenditure per worker gap is 1.7 relative to 2.1 in income per worker and 2.2 in micro-APGs.¹⁵

We conclude that household income per worker appears lower for agricultural households than non-agricultural households, with gaps similar in magnitude to those we observed in value added per worker in Section 5. Thus, at least for these ten countries, it appears unlikely that an explanation of the large residual APGs comes down to the distinction between agricultural workers and agricultural households. Put differently, income gaps are present between agri-

¹⁵Part of the reason that the expenditure gaps may be smaller than income gaps is the existence of certain types of insurance arrangements, which often involve transfers from richer households to poorer ones. One prominent example are remittances from non-agricultural workers (in the city, say) back to relatives in agricultural areas. These types of transfers could account for the differences between the value added per worker, income per worker and expenditure per worker metrics. In Pakistan, for example, where the largest discrepancy exists between the gap in income and expenditure per worker, [Ilahia and Jafareyc \(1999\)](#) argue that remittances to rural agricultural households provide for a substantial amount of the consumption of rural agricultural families.

Table 5: Micro APGs and Ratios of Income and Expenditure Per Worker

Country	APG Micro	Income per Worker Ratio	Expenditure per Worker Ratio
Armenia (1996)	1.1	0.7	0.9
Bulgaria (2003)	0.7	1.4	1.2
Cote d'Ivoire (1988)	4.0	3.5	3.2
Ghana (1998)	2.3	2.0	1.9
Guatemala (2000)	2.9	3.2	2.4
Kyrgyz Republic (1998)	2.0	1.3	1.8
Pakistan (2001)	3.0	3.2	1.4
Panama (2003)	2.7	2.8	2.1
South Africa (1993)	1.7	1.7	1.2
Tajikistan (2009)	1.6	1.2	1.1
Average	2.2	2.1	1.7

Note: APGs are calculated as in Table 4. Ratios of income per worker are calculated as income per worker in agricultural households divided by income per worker in non-agricultural households. Households are classified as agricultural if the majority of their workers report agriculture as their primary sector of employment, and in the event of equal numbers of workers in each sector, whether the majority of the household's income comes from agricultural activities. Ratios of expenditure per worker are calculated in the same way, but using total household expenditure.

cultural and non-agricultural households, not just between agricultural and non-agricultural workers.

6.2. Differences in Amenities and Cost of Living

One possible explanation of the residual gaps is that there are offsetting benefits to living and working in rural areas, as opposed to cities. One might imagine, for instance, that there is better access to public services or to health care in rural areas.

On measure after measure, however, access to key public goods are consistently lower in rural areas than in urban areas, in the developing world at least. A recent [World Bank \(2013\)](#) report notes that on almost all measures of health and human capital, rural areas of developing countries lag behind urban areas; often, the greatest rural-urban discrepancies are found in the poorest countries. For instance, infant mortality rates in 40 sub-Saharan African countries are substantially higher in rural areas than in urban areas (80 deaths per 1,000 live births, compared to 65). Child malnutrition (as measured by stunting) is higher in rural areas. Literacy rates are higher in urban areas, and educational quality is consistently higher in urban areas. Specific public services and infrastructure are also more common in urban areas than in rural, in part

due to higher per unit costs of service delivery in rural areas: babies are more likely to be delivered at health centers in urban areas than in rural areas (78 percent compared with 43 percent in a sample of 28 sub-Saharan countries). Access to sanitation and piped water are lower in rural areas than urban areas. Electrification rates in sub-Saharan Africa are far lower in rural areas than urban areas (see [Hewett and Montgomery \(2001\)](#)). In short, it is difficult to find measures of public goods provision that are not lower in rural areas of the developing world than urban areas.

One form of “amenity” that may help explain the residual gaps is the varying degree of informal social insurance provided for individuals who stay within their rural communities, as opposed to the risk-sharing available to those who migrate to urban centers. An abundant literature (well surveyed and analyzed both theoretically and empirically by [Morten \(2013\)](#)) demonstrates that risk-sharing within villages, while imperfect, is nevertheless extremely effective at providing protection from negative shocks. Individuals who migrate to cities may lose this form of protection.

Another possible explanation for the reluctance of people to move from agriculture to non-agriculture, even in the face of productivity gaps, is the potential increase in the cost of living in urban areas. In other words, even though incomes may be higher in urban areas, they may be offset by the higher cost of the same basket of goods and services.

Although few countries have detailed data that permit such calculations, the available data suggest that differences in urban-rural cost of living do not account for the remaining difference in incomes between rural and urban areas. For instance [Munshi and Rosenzweig \(2013\)](#) find that in India, daily wage rates for individuals with less than primary school are still substantially higher in urban areas even after controlling for cost-of-living differences. [Nguyen, Albrecht, Brogman, and Westbrook \(2007\)](#) reach a similar conclusion for Vietnam, as do [Sicular, Yue, Gustafsson, and Li \(2007\)](#) for China.¹⁶

6.3. Selection

Another potential explanation for the large residual sector productivity gaps is that workers select themselves by sector in such a way that the average worker in the non-agricultural sector earns a higher average wage. [Lagakos and Waugh \(2013\)](#) formalize one version of this story, in which workers are heterogeneous in ability in each sector and choose where to supply their labor. In a parameterized version of their model, the average wage is higher in the non-agricultural sector in equilibrium, even without barriers to moving out of agriculture.

¹⁶See also the work of [Ravallion, Chen, and Sangraula \(2007\)](#), who look at estimates of the cost of a basket of goods purchased by households living on \$ 1 per day in rural and urban areas in a set of developing countries. In general, they do find that their basket is cheaper in rural areas, but the differences are fairly modest compared to the agricultural productivity gaps in the current paper.

The reason is that the underlying distribution of non-agricultural abilities is more dispersed than the distribution of agricultural ability, and workers with the highest endowments of non-agricultural ability disproportionately enter the non-agricultural sector.

Young (2013) provides the most compelling evidence for selection. He uses cross-country, micro-level survey data to study urban-rural consumption gaps and connects these measures with migration outcomes. Similar to our results, he finds large gaps in real consumption between urban and rural areas. Different from us, he is able to observe migration between rural and urban areas and finds large flows of people moving from urban to rural areas despite the fact that average consumption is lower in the rural area. This migration pattern is inconsistent with a simple model of migration which would predict large losses in welfare for these migrants. However, Young (2013) argues that this pattern is consistent with selection. Building on the work of Lagakos and Waugh (2013), he finds strong empirical support for a model in which workers sort on the basis of unobservable skill.

The work of Beegle, De Weerd, and Dercon (2011) is suggestive of selection as well. They show that students in Tanzania who attend more years of school are more likely to move out of agricultural work and into the urban non-agricultural sector. One interpretation of this finding is that those individuals with greater cognitive abilities are the ones selecting more years of schooling and, subsequently, work in the non-agricultural sector. This is consistent with the evidence from Young (2013), who shows that rural-to-urban migrants are typically better educated than rural permanent residents, and urban-to-rural migrants are typically less educated than urban permanent residents. It is also consistent with our evidence in Section 4.4, which shows much higher schooling attainment rates among non-agricultural workers.

A final piece of evidence comes from the work of Miguel and Hamory (2009), who draw on a unique data set of cognitive ability scores from rural students in Kenya. The students were given the cognitive ability test in primary school and then later followed. The authors find that the scores are a very strong predictor of who later migrates out of agricultural areas to take non-agricultural employment. Their estimates are that individuals scoring one standard deviation higher on cognitive ability scores are roughly 17 percent more likely to migrate.

6.4. Sector Differences in Labor's Share in Production

Up to this point, we have maintained the assumption that labor shares in production are the same in agriculture and non-agriculture. Could sector differences in labor shares account for much of the remaining gap? To answer this question, we consider a modification of the Cobb-Douglas production function in (1), but now the importance of labor and other inputs in pro-

duction differs across sectors:

$$Y_a = A_a L_a^{\theta_a} K_a^{1-\theta_a} \quad \text{and} \quad Y_n = A_n L_n^{\theta_n} K_n^{1-\theta_n}. \quad (10)$$

Now, the firms' first-order conditions plus free labor mobility imply that sector differences in value added per worker are given by the ratio of the Cobb-Douglas factor shares for labor:

$$\frac{V A_n / L_n}{V A_a / L_a} = \frac{Y_n / L_n}{p_a Y_a / L_a} = \frac{\theta_a}{\theta_n}. \quad (11)$$

Equation (11) suggests that we can explain the remaining sectoral differences in average labor productivity if θ_n is approximately half as large as θ_a .

Is labor's share in non-agricultural production half as large as in agriculture? Several pieces of evidence suggest that it is not. First, estimates of agricultural labor shares computed using producer-level time series or cross-sectional data suggest that this is not the case. [Fuglie \(2010\)](#) provides a recent review of the estimates from around the world. His data imply that the average share of labor relative to land, equipment and structures is 0.58 for China, India, Indonesia, Brazil, Mexico, and sub-Saharan Africa, while the corresponding figures for the U.S. and U.K. are 0.51 and 0.52. In order to explain the residual APGs found in the paper, the non-agricultural labor share would have to be in the range of 0.25 to 0.29, which is highly implausible given that the aggregate labor shares are in the ballpark of two thirds (see, e.g., [Gollin \(2002\)](#).)

A second piece of evidence against θ_n being substantially lower than θ_a comes from the relationship between aggregate labor shares and income per capita across countries. As [Gollin \(2002\)](#) points out, labor shares—once adjusted for the mixed income of the self-employed—vary little across countries, and the variation is largely uncorrelated with income per capita. If this is the case, and if agriculture's share of GDP varies systematically with income per capita (as is widely understood), then labor shares cannot differ very much between agriculture and non-agriculture; otherwise, we would observe large and systematic variation in aggregate labor shares. We do not, which suggests that θ_n and θ_a are similar in size.

The final piece of evidence comes from observations from share tenancy arrangements. In much of the world, large areas of agricultural land are farmed by operators under share tenancy arrangements, in which the operators pay a fraction of gross or net output to land owners in lieu of a cash rent (see, e.g., [Otsuka \(2007\)](#).) These arrangements are informative about cost shares in production since the operator typically provides all the labor, while the land owner provides the land and buildings. In principle, then, the split of gross output between the operator and the land owner, along with the allocation of capital costs and intermediate input costs, will allow for the calculation of the (net) share of labor in production. In practice, it may be difficult to arrive at precise calculations, because relationships between land owners and operators may be

quite complicated (see, e.g., [Jacoby and Mansuri \(2009\)](#).) Nevertheless, the gross output share and the cost shares provide a useful—if crude—estimate of the factor shares.

A striking stylized fact in the share tenancy literature is that over time and across countries, most share contracts seem to involve 50-50 splits of both gross output and intermediate inputs. [Otsuka \(2007\)](#) refers to this as the “commonly observed rate,” and [Otsuka, Chuma, and Hayami \(1992\)](#) note that “the output sharing rate is almost universally 50 percent under share tenancy in many developing countries.”¹⁷ [Jacoby and Mansuri \(2009\)](#) note that in survey data for rural Pakistan, in 1993 and 2001, “nearly three-quarters of share-tenants . . . report a 50-50 output sharing rule.” The 50-50 split is also common in modern-day agriculture in the United States (see [Canjels \(1998\)](#)). The predominance of the 50-50 split would tend to suggest that θ_a is around one half. If this is true, then one would require θ_n of around 0.25 to explain the residual APGs, which is again highly implausible given that aggregate estimates are around two thirds.

7. Correlates of the Adjusted Gaps

In this section, we study the correlations of the adjusted APGs with some salient observable country characteristics that we expect may be linked to the size of the residual gaps. The idea is that these correlations can point toward possible explanations for the remaining productivity gaps. The measures that we focus on relate to geographical features, measures of institutional quality, and measures of labor mobility. We find that the residual gaps are correlated with all these variables, and in ways that suggest possible avenues for future research. An important caveat is that results should be interpreted as only correlations and not as causal relationships.

In what follows, we present the results of regressions of (the log of) our adjusted APG measures on observable country characteristics. To give ourselves the maximum number of countries in the regressions, we include all countries in our sample. For any country in which hours or schooling data are missing, we impute them as the average values of those of other countries in the same income quartile. Results using only countries with complete data present largely similar coefficient estimates and are available upon request.

The first set of variables we consider pertain to geographic features of the countries in question. Researchers such as [Bloom and Sachs \(1998\)](#) have argued that being located in the tropics leads to lower physical productivity in agriculture; having less fertile soil may have the same effect. We therefore include as independent variables the fraction of a country’s area that is in the tropics and the fraction of soil that is classified as fertile by the FAO. [Nunn and Puga \(2012\)](#) have explored the role of ruggedness, which basically measures the average variability in elevation within a country, and have argued that ruggedness leads to lower physical productivity in

¹⁷They further note that the 50-50 split was historically pervasive in many parts of the world, to the extent that the French and Italian words for share tenancy (*metayage* and *mezzadria*, respectively) mean “splitting in half.”

agriculture. We also suspect that rugged terrain may make it harder for workers to relocate internally to take advantage of better opportunities elsewhere, either because relocating is more difficult or because knowledge of those opportunities diffuses more slowly. We thus include their measure of ruggedness as well. The last geographic measure we include is the fraction of all imports that are fuel or metals. One might suspect that if nonagricultural value added per worker is relatively high, this could be because the non-agricultural sector is based heavily on natural resources such as oil, gas or precious metals.

The second set of variables pertains to institutions. A large literature has emphasized the negative economic consequences of poorly functioning institutions (see e.g. [Hall and Jones \(1999\)](#) and [Acemoglu, Johnson, and Robinson \(2001\)](#).) In our context, one might imagine that migrating internally might be much more difficult when property rights are poorly protected. For example, a lack of property rights in land markets makes it hard for rural residents to sell their land (in many cases their main form of savings) and move to urban areas. Similarly, the risk of expropriation may make workers less inclined to leave behind the security of tight-knit village or family living arrangements. There are many measures to proxy the quality of a country's institutions.¹⁸ For simplicity, we use the rule of law index from the the World Bank's Worldwide Governance Indicators database ([Kaufmann, Kraay, and Mastruzzi \(2008\)](#)). This index aggregates information from a variety of stakeholders pertaining to the extent to which agents have confidence in and abide by the rules of society.

The third group of variables relates to more specific factors that might inhibit the movement of workers within a country.¹⁹ In some countries there are either formal (such as the hukou system in China) or informal barriers to internal migration. These types of restrictions are an obvious explanation for large gaps in sector value added per worker. In our regressions, we include data from [Cingranelli and Richards \(2010\)](#), who provide a systematic characterization of restrictions on the ability of people to move within a country. They rely on US State Department country human rights reports to categorize a country as having restrictions on internal mobility if various actions in the country are recorded. Note that these actions might conflict with official regulations. For example, a country's laws may allow for the freedom of movement, but the government may use tactics to intimidate and/or prevent people from moving. This variable takes the value of one if there are moderate or severe restrictions on domestic movement and zero otherwise.

Another factor we consider is ethnic fractionalization. One could also imagine that a society that is ethnically fractionalized may inhibit the ability of workers to reallocate. For ex-

¹⁸We considered other measures of institutions, such as the political institutional quality index constructed by [Henisz \(2000\)](#). Using this measure in place of the rule of law index gives very similar results to those in Table 6.

¹⁹We also explored the correlation of employment laws on the adjusted APGs using the employment law index of [Botero, Djankov, La Porta, Lopez-de Silanes, and Shleifer \(2004\)](#). We found only a weak correlation and did not include the results in our main specification, as these measures are available only for a limited set of countries.

ample, those in agricultural areas may be ethnically or linguistically different than those in non-agricultural areas which would raise the effective cost of migration. To explore this possibility, we include the measure of ethnic fractionalization of [Alesina, Devleeschauwer, Easterly, Kurlat, and Wacziarg \(2003\)](#). Their measure is combination of both racial and linguistic characteristics, and it reflects the probability that two randomly selected individuals belong to a different ethnolinguistic group.

Table 6 presents the results. Column (1) focuses on our geography measures. Ruggedness has a positive and significant sign, and the coefficient implies that a one standard deviation increase in ruggedness is associated with an 18 percent larger residual gap. This is consistent with a story in which ruggedness leads to less internal migration and lower physical productivity in agriculture. The fraction of land that is in the tropics is negatively correlated with the residual gaps, while the fraction of land that is fertile is positively correlated. These findings are consistent with a world in which being in the tropics or having less-fertile soil reduces physical productivity in agriculture, and then for some reason workers do not move out, or relative agricultural prices rise, enough to equate the value of output per worker in the two sectors. Neither variable is statistically significant, however. The fraction of exports of fuel or minerals has a positive and significant sign, and the coefficient implies that a one standard deviation increase is associated with an 18 percent larger residual gap. Again, this finding is consistent with the idea that we might observe large residual gaps in countries where the non-agricultural sector is heavily based on natural resources.

Columns (2) - (4) focus on our proxy for institutional quality, restrictions on domestic movement, and ethnic fractionalization. Individually, each of these variables is correlated with the residual gaps, is statistically different than zero, and has a sign consistent with our discussion above. Countries that lack the rule of law have larger adjusted APGs, with a one standard deviation decrease in the rule of law associated with a 16 percent higher residual gap. This finding is very much consistent with the idea that poor protection of property rights would provide workers with a disincentive to migrate. Countries that have restrictions on domestic movement have larger residual gaps, and these restrictions lead to, on average, a 24 percent increase in the residual gaps. Ethnic fractionalization is also associated with larger residual gaps, with a one standard deviation increase in fractionalization associated with a 13 percent higher residual gaps. Again, this evidence is consistent with the idea that fractionalization may impede the ability of workers to move.

Column (5) includes all these measures (geography, institutions, mobility restrictions) together. The signs for all variables are unchanged relative to the previous regression. However, it is only ruggedness that remains statistically different from zero, and the magnitudes are muted in some cases. This appears to be because several of these measures are correlated with each other— i.e. countries with weak institutions also happen to be large fuel and mineral exporters

Table 6: Correlations of Adjusted Productivity Gaps and Country Characteristics

	Dependent Variable: Log Adjusted APG				
	(1)	(2)	(3)	(4)	(5)
Ruggedness	0.902 (0.318)***				0.850 (0.355)**
Fertile soil	-0.153 (0.237)				-0.081 (0.236)
Tropical climate	0.178 (0.112)				0.051 (0.124)
Fuel and Mineral Exports	0.630 (0.256)**				0.378 (0.276)
Rule of Law Index		-0.781 (0.222)***			-0.440 (0.287)
Restrictions on Domestic Movement			0.235 (0.112)**		0.052 (0.121)
Ethnic Fractionalization				0.495 (0.200)**	0.328 (0.236)
Constant	0.310 (0.159)*	1.023 (0.128)***	0.547 (0.060)***	0.423 (0.089)***	0.474 (0.316)
R^2	0.14	0.08	0.03	0.04	0.18
N	140	146	146	142	136

Note: OLS regression. Coefficients are reported with robust standard errors in brackets. *, **, *** indicates significance at the 10, 5, and 1 percent level. Ruggedness is taken from [Nunn and Puga \(2012\)](#) and is normalized to lie between zero and one. Fertile soil and tropical climate are the fraction of land that is fertile or lies in the tropics. Fuel and mineral exports is the fraction of total merchandise exports. Rule of law index is from [Kaufmann, Kraay, and Mastruzzi \(2008\)](#) and lies between zero and one. Restrictions on movement is a dummy variable that takes the value one if a country is classified by [Cingranelli and Richards \(2010\)](#) to restrict internal movement. Ethnic Fractionalization is from [Alesina, Devleeschauwer, Easterly, Kurlat, and Wacziarg \(2003\)](#) and lies between zero and one.

and are ethnically fractionalized. The inability to separately identify these effects is a limitation of these results. However, we do believe that Table 6 provides evidence warranting further investigation of these relationships, and understanding the exact mechanisms by which, say, weak institutions could lead to measured productivity gaps, is an important topic for future research.

8. Conclusion

According to national accounts data, value added per worker is much higher in the non-agricultural sector than in agriculture in most countries. This agricultural productivity gap, when taken at face value, suggests that labor is greatly misallocated across sectors. In this paper, we ask to what extent the gap is still present when better measures of sector inputs and outputs are taken into consideration. To answer this question, we construct a new data set for

a large number of countries, with measures of hours worked and human capital per worker by sector. We also consider alternative measures of value added per worker constructed from household income surveys for a smaller set of developing countries.

We find that even after taking all these measurement issues into consideration, a puzzlingly large agricultural productivity gap remains. The value of output per worker in non-agriculture still appears to be roughly twice as high as in agriculture in the typical country, and even higher in the typical developing countries. The implication is that there should be large income gains from workers moving out of agriculture and into other economic activities.

A recent randomized controlled trial by [Bryan, Chowdhury, and Mobarak \(2013\)](#) in rural Bangladesh finds just this. In their experiment, a treatment group of households were given small subsidies to out-migrate from rural areas to nearby cities during the “lean season.” The authors find large and persistent benefits accruing to treatment households from this intervention, including evidence that the treatment households are more likely to send seasonal migrants to urban areas (without subsidy) for at least several years after the initial intervention. On average, these migrating workers and their families experienced sizable increases in income relative to workers that stayed behind.

References

- ACEMOGLU, D., S. JOHNSON, AND J. A. ROBINSON (2001): "The Colonial Origins of Comparative Development: An Empirical Investigation," *American Economic Review*, 91(5), 1369–1401.
- ADAMOPOULOS, T., AND D. RESTUCCIA (Forthcoming): "The Size Distribution of Farms and International Productivity Differences," *American Economic Review*.
- ALESINA, A., A. DEVLEESCHAUWER, W. EASTERLY, S. KURLAT, AND R. WACZIARG (2003): "Fractionalization," *Journal of Economic Growth*, 8(2), 155–194.
- BANERJEE, A., AND E. DUFLO (2005): "Growth Theory Through the Lens of Development Economics," in *Handbook of Economic Growth*, ed. by P. Aghion, and S. Durlauf.
- BARRO, R. J., AND J.-W. LEE (2013): "A New Data Set of Educational Attainment in the World, 1950-2010," *Journal of Development Economics*, 104, 184–198.
- BEEGLE, K., J. DE WEERDT, AND S. DERCON (2011): "Migration and Economic Mobility in Tanzania," *Review of Economics and Statistics*, 93, 1010–1033.
- BLOOM, D. E., AND J. D. SACHS (1998): "Geography, Demography, and Economic Growth in Africa," *Brookings Papers on Economic Activity*, 1998(2), 207–295.
- BOTERO, J. C., S. DJANKOV, R. LA PORTA, F. LOPEZ-DE SILANES, AND A. SHLEIFER (2004): "The Regulation of Labor," *Quarterly Journal of Economics*, 119(4), 1339–1382.
- BRYAN, G., S. CHOWDHURY, AND M. MOBARAK (2013): "Escaping Famine through Seasonal Migration," Unpublished Manuscript, Yale University.
- CANJELS, E. (1998): "Risk and Incentives in Sharecropping: Evidence from Modern U.S. Agriculture," The New School Center for Economic Policy Analysis Working Paper Number IV.
- CASELLI, F. (2005): "Accounting for Cross-Country Income Differences," in *Handbook of Economic Growth*, ed. by P. Aghion, and S. Durlauf.
- CASELLI, F., AND W. J. COLEMAN (2001): "The U.S. Structural Transformation and Regional Convergence: A Reinterpretation," *Journal of Political Economy*, 109(3), 584–616.
- CASELLI, F., AND J. FEYRER (2007): "The Marginal Product of Capital," *Quarterly Journal of Economics*, 122(2), 535–568.
- CHANDA, A., AND C.-J. DALGAARD (2008): "Dual Economies and International Total Factor Productivity Differences: Channelling the Impact from Institutions, Trade and Geography," *Economica*, 75, 629–661.

- CINGRANELLI, D. L., AND D. L. RICHARDS (2010): "The Cingranelli-Richards Human Rights Dataset Dataset Version 2010.05.17," URL: <http://www.humanrightsdata.org>.
- DEATON, A. (1997): *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*. The World Bank.
- DONOVAN, K. (2013): "Agricultural Risk, Intermediate Inputs, and Cross-Country Productivity Differences," Unpublished Manuscript, University of Notre Dame.
- FAO (1996): *A System of Economic Accounts for Food and Agriculture*. Food and Agriculture Organization.
- FRICK, J. R., S. P. JENKINS, D. R. LILLARD, O. LIPPS, AND M. WOODEN (2007): "The Cross-National Equivalent File (CNEF) and its Member Country Household Panel Studies," *Schmollers Jahrbuch*, 127, 627–654.
- FUGLIE, K. (2010): *The Shifting Patterns of Agricultural Production and Productivity Worldwide* chap. 4. Iowa State University.
- GOLLIN, D. (2002): "Getting Income Shares Right," *Journal of Political Economy*, 110(2), 458–474.
- GOLLIN, D., S. L. PARENTE, AND R. ROGERSON (2002): "The Role of Agriculture in Development," *American Economic Review Papers and Proceedings*, 92(2).
- (2004): "Farm Work, Home Work and International Productivity Differences," *Review of Economic Dynamics*, 7.
- HALL, R. E., AND C. I. JONES (1999): "Why Do Some Countries Produce So Much More Output per Worker than Others?," *Quarterly Journal of Economics*, 114(1), 83–116.
- HENISZ, W. (2000): "The Institutional Environment for Economic Growth," *Economics and Politics*, 12(1), 1–31.
- HERRENDORF, B., AND T. SCHOELLMAN (2013): "Why is Agricultural Productivity So Low in Agriculture?," Unpublished Manuscript, Arizona State University.
- HEWETT, P. C., AND M. MONTGOMERY (2001): *Poverty and Public Services in Developing-Country Cities*, vol. 154. Population Council New York.
- HSIEH, C.-T., AND P. J. KLENOW (2009): "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 124, 1403–1448.
- ILAHIA, N., AND S. JAFAREYC (1999): "Guestworker Migration, Remittances and the Extended Family: Evidence From Pakistan," *Journal of Development Economics*, 58(2), 285–512.

- JACOBY, H., AND G. MANSURI (2009): "Incentives, Supervision, and Sharecropper Productivity," *Journal of Development Economics*, 88(2), 232–241.
- KAUFMANN, D., A. KRAAY, AND M. MASTRUZZI (2008): "Governance Matters VII : Aggregate and Individual Governance Indicators 1996-2007," Policy Research Working Paper Series 4654, The World Bank.
- KLENOW, P. J., AND A. RODRÍGUEZ-CLARE (1997): "The Neoclassical Revival in Growth Economics: Has it Gone Too Far?," in *NBER Macroeconomics Annual 1997*, ed. by B. S. Bernanke, and J. Rotemberg. MIT Press, Cambridge, MA.
- KUZNETS, S. (1971): *Economic Growth of Nations: Total Output and Production Structure*. Harvard University Press.
- LAGAKOS, D., B. MOLL, T. PORZIO, AND N. QIAN (2013): "Experience Matters: Human Capital and Development Accounting," Unpublished Manuscript, Princeton University.
- LAGAKOS, D., AND M. E. WAUGH (2013): "Selection, Agriculture and Cross-Country Productivity Differences," *American Economic Review*, 103(2), 948–80.
- LEWIS, W. (1955): *The Theory of Economic Growth*. Irwin.
- MCMILLAN, M. S., AND D. RODRIK (2011): "Globalization, Structural Change and Productivity Growth," NBER Working Paper No. 17143.
- MIGUEL, E., AND J. HAMORY (2009): "Individual Ability and Selection into Migration in Kenya," Unpublished Manuscript, University of California, Berkeley.
- MINNESOTA POPULATION CENTER (2013): "Integrated Public Use Microdata Series, International: Version 6.2," Minneapolis: University of Minnesota.
- MORTEN, M. (2013): "Temporary Migration and Endogenous Risk Sharing in Village India," Unpublished Manuscript, Yale University.
- MUNSHI, K., AND M. ROSENZWEIG (2013): "Networks and Misallocation: Insurance, Migration, and the Rural-Urban Wage Gap," Unpublished Manuscript, Brown University.
- NGUYEN, B. T., J. W. ALBRECHT, S. B. BROGMAN, AND D. M. WESTBROOK (2007): "A Quantile Regression Decomposition of Urban-Rural Inequality in Vietnam," *Journal of Development Economics*, 83(2), 466–490.
- NUNN, N., AND D. PUGA (2012): "Ruggedness: The Blessing of Bad Geography in Africa," *Review of Economics and Statistics*, 94(1), 20–36.

- OTSUKA, K. (2007): "Efficiency and Equity Effects of Land Markets," *Handbook of Agricultural Economics*, 3, 2671–2703.
- OTSUKA, K., H. CHUMA, AND Y. HAYAMI (1992): "Land and Labor Contracts in Agrarian Economies: Theories and Facts," *Journal of Economic Literature*, 30(4), 1965–2018.
- PSACHAROPOULOS, G., AND H. A. PATRINOS (2002): "Returns to Investment in Education: A Further Update," World Bank Policy Research Working Paper 2881.
- RAVALLION, M. (2003): "Measuring Aggregate Welfare in Developing Countries: How Well Do National Accounts and Surveys Agree?," *Review of Economics and Statistics*, 85(3), 645–652.
- RAVALLION, M., S. CHEN, AND P. SANGRAULA (2007): "New Evidence on the Urbanization of Global Poverty," *Population and Development Review*, 33(4), 667–701.
- RESTUCCIA, D., AND R. ROGERSON (2008): "Policy Distortions and Aggregate Productivity with Heterogeneous Establishments," *Review of Economic Dynamics*, 11(4), 707–720.
- RESTUCCIA, D., D. T. YANG, AND X. ZHU (2008): "Agriculture and Aggregate Productivity: A Quantitative Cross-Country Analysis," *Journal of Monetary Economics*, 55, 234–250.
- ROSENSTEIN-RODAN, P. (1943): "Problems of Industrialisation of Eastern and South-Eastern Europe," *The Economic Journal*, 53(210), 202–211.
- ROSTOW, W. (1960): *The Stages of Economic Growth: A Non-Communist Manifesto*. Cambridge, Cambridge University Press.
- SCHOELLMAN, T. (2012): "Education Quality and Development Accounting," *Review of Economic Studies*, 79(1), 388–417.
- SICULAR, T., X. YUE, B. GUSTAFSSON, AND S. LI (2007): "The Urban-Rural Income Gap and Inequality in China," *Review of Income and Wealth*, 53(1), 93–126.
- VOLLRATH, D. (2009): "How Important are Dual Economy Effects for Aggregate Productivity?," *Journal of Development Economics*, 88(2), 325–334.
- WILLIAMS, J. H. (2005): "Cross-National Variations in Rural Mathematics Achievement: A Descriptive Overview," *Journal of Research in Rural Education*, 20(5).
- WORLD BANK (2013): *Global Monitoring Report 2013: Rural-Urban Dynamics and the Millennium Development Goals*. The World Bank.
- YOUNG, A. (2012): "The African Growth Miracle," *Journal of Political Economy*, 120(4), 696–739.

——— (2013): “Inequality, the Urban-Rural Gap and Migration,” *Quarterly Journal of Economics*, 128(4), 1727–1785.

ZHANG, Y. (2006): “Urban-Rural Literacy Gaps in Sub-Saharan Africa: The Roles of Socioeconomic Status and School Quality,” *Comparative Education Review*, 50(4).