The Agricultural Productivity Gap in Developing Countries

Douglas Gollin

Williams College

David Lagakos

Arizona State University

Michael E. Waugh

New York University

July 2011

ABSTRACT

According to national accounts data for developing countries, value added per worker is on average four times higher in the non-agriculture sector than in agriculture. Taken at face value this "agricultural productivity gap" suggests that labor is greatly misallocated across sectors in the developing world. In this paper we draw on new micro evidence to ask to what extent the gap is still present when better measures of inputs and outputs are taken into consideration. We find that even after considering sector differences in hours worked and human capital per worker, urban-rural cost-of-living differences, and alternative measures of sector income from household survey data, a puzzlingly large agricultural productivity gap remains.

Email: douglas.gollin@williams.edu, lagakos@asu.edu, mwaugh@stern.nyu.edu. Preliminary and incomplete – comments welcome. All potential errors are our own. We thank Lisa Starkman for excellent research assistance. For helpful comments we thank Francesco Caselli, Matthias Doepke, Berthold Herrendorf, Maggie McMillan, Richard Rogerson, Todd Schoellman, and Francis Teal, as well as from seminar participants at Northwestern University, the conference on Economic Growth and Cultural Change at the University of Munich, and the World Bank's Annual Bank Conference on Development Economics. This paper was 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.

1. Introduction

The agriculture sector accounts for large fractions of employment and value added in developing countries. Almost always, agriculture's share of employment is higher than its share of value added. As a simple matter of arithmetic, this implies that value added per worker is higher in the non-agriculture sector than in agriculture. According to data from national income and product accounts, this "agricultural productivity gap" (APG) is around a factor of four in developing countries, on average. In many poor countries the gaps are even higher, with a number of countries having gaps above ten.

These large agricultural productivity gaps have several important implications for developing countries. First, with minimal assumptions on production technologies, they imply that labor is misallocated across sectors. Second, they imply that developing countries trail the developed world by a much larger margin in agriculture than in non-agriculture. Put together, these two implications suggest that the problem of economic development is closely linked to an apparent "misallocation" of workers across sectors, with too many workers in the less-productive agriculture sector.

In this paper we draw on new micro evidence to ask to what extent these gaps are still present when better measures of inputs and outputs are taken into consideration. Our analysis addresses a basic, yet unanswered question: how much of the agricultural productivity gaps are due to problems of omitted factors and mis-measurement, as opposed to real differences in income per worker? To answer this question, we consider a sequence of adjustments to the data on agriculture's shares of value added and employment. These adjustments attempt to control for differences across sectors in hours worked per worker, human capital per worker, and cost-of-living differences between rural and urban areas. We then ask to what extent value added in the national accounts differs from value added measures constructed using data from household income surveys.

Our analysis draws on a new database that we constructed from population censuses and labor force surveys for a large set of developing countries. We use these data to construct measures of hours worked by sector for 46 developing countries and measures of human capital by sector for 97 countries. We complement these data with evidence on urban-rural differences in the cost of living for 78 developing countries, constructed by the World Bank. We find that taking these differences into consideration jointly reduces the size of the average agricultural productivity gap to around two.

We then ask whether the gaps are consistent with sector value added measures computed from household income survey data. We construct these measures from the World Bank's Living Standards Measurement Studies (LSMS), which are designed explicitly to obtain measures of household income and expenditure. These surveys allow us to compute, among other things, the market value of all output produced by agricultural households, whether they are ultimately sold or consumed at home. They also allow to construct measures of average income and expenditure by agriculture and non-agricultural households.

Our analysis suggests that the agricultural productivity gaps in most developing countries are unlikely to be completely explained by any of the measurement issues we address. We conclude that a better understanding is needed of why so many workers remain in the agriculture sector, given the large residual productivity gaps that we find in most developing countries. Understanding these gaps will help determine, in particular, whether policy makers in the developing world should pursue polices that encourage movement of the workforce out of agriculture.

To be sure, we are not the first to point out the existence of large agricultural productivity gaps in some countries. As Lewis (1955) noted (pp. 349-40):[1]

> There is usually a marked difference between incomes per head in agriculture and in industry. Some of the difference in money income is illusory; rural workers get some income in kind, pay less for many things they buy (especially food and living accommodation) and do not have to spend so much as the urban population on some other costs of living and enjoying (e.g., transportation). Nevertheless, when account is taken of this . . . real income per head is lower in agriculture than it is in manufacturing.

These differences in sectoral productivity were viewed as critical by early development economists, who saw the development process as fundamentally linked to the reallocation of workers across sectors through the expansion of modern industry, oriented at least in part towards export markets. Thus, Rosenstein-Rodan (1943), Lewis (1955), and Rostow (1960) viewed development as essentially identical, with the movement of people out of agriculture and into "modern" economic activities.

More recently, the work of Caselli (2005), Restuccia, Yang, and Zhu (2008), Chanda and Dalgaard (2008), and Vollrath (2009) has shown that the apparent misallocation of workers across agriculture and non-agriculture can account for the bulk of international income and productivity differences. The reason is that most poor countries have very unproductive agricultural sectors, yet employ most of their workers in agriculture.[2]

Our paper builds on these studies by bringing in richer data on inputs and outputs at the sector

[1] The fact that the agriculture productivity gaps are most prevalent in poor countries was first shown by Kuznets (1971), and later documented in richer detail by Gollin, Parente, and Rogerson (2002). Interestingly, Gollin, Parente, and Rogerson (2002) note that the disparities were fairly small in today's rich countries at moments in the historical past when their incomes were substantially lower than at present.

[2] In related work, McMillan and Rodrik (2011) argue that reallocations of workers to the most productive sectors would raise income dramatically in many developing countries.

level. In particular, our paper is the first to make use of census-based measures of schooling attainment by sector, hours worked by sector, and cost-of-living differences in urban and rural areas. 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.

The paper most closely related to ours is by Herrendorf and Schoellman (2011), who ask why agricultural productivity gaps are so large in most U.S. states. While both studies are ultimately motivated by the large sectoral productivity differences in developing countries, theirs makes use of U.S. data which is richer than ours in several dimensions. One similarity is that we both find a large role for differences in human capital across sectors in explaining the sector productivity gaps.

2. Agricultural Productivity Gap — Theory

In this section, we discuss some implications of standard neoclassical theory for data. Consider the standard neoclassical two-sector model featuring constant returns to scale in the production of agriculture and non-agriculture, along with free labor mobility across sectors and competitive labor markets.[3] 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 firms hire labor up to the point where the marginal value product of labor equals the wage. Since wages are equalized across sectors, this implies that marginal value products are also equalized:

$$p_a \frac{\partial F_a(\mathbf{X})}{\partial L} = \frac{\partial F_n(\mathbf{X})}{\partial L} = w, \tag{1}$$

where subscripts a and n denote agriculture and non-agriculture. Units are chosen here such that the non-agricultural good is the numeraire, p_a is the relative price of the agricultural good, and $\mathbf{X}$ is a vector of inputs (including labor) used in production.

If the production function displays constant returns to scale, then marginal products are proportional to average products with the degree of proportionality depending on that factors share in production. Defining $1-\alpha_a$ and $1-\alpha_n$ as the shares of labor in production, the constant-return production functions imply:

$$(1-\alpha_a) \times \frac{p_a Y_a}{L_a} = (1-\alpha_n) \times \frac{Y_n}{L_n}. \tag{2}$$

Noting that $p_a Y_a$ and Y_n equals value added in the agriculture and non-agriculture sector, equation

[3] Parametric examples in the literature include Gollin, Parente, and Rogerson (2004), Gollin, Parente, and Rogerson (2007) and Restuccia, Yang, and Zhu (2008).

(2) says that value added per worker across the two sectors should be equated (modulo differences in labor shares). Assuming that labor shares are the same across sectors implies that

$$\frac{Y_n/L_n}{p_a Y_a/L_a} \equiv \frac{VA_n/L_n}{VA_a/L_a} = 1. \tag{3}$$

If the condition in (3) 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, we should see workers 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 (3) is that it says nothing about misallocation in other factor markets. 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 (3) does not hold in the data, the explanation must lie either in either measurement problems related to labor inputs or in frictions of some kind in the labor market – nothing else.

Writing equation (3) in terms of agriculture's share of employment and output gives:

$$\frac{(1-y_a)/(1-\ell_a)}{y_a/\ell_a} = 1. \tag{4}$$

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 (4) is the lens through which we look at the data. Under the (minimal) conditions outlined above, we first ask if the condition in (4) holds in cross-country data. One way to think about this exercise is along the lines of Restuccia and Rogerson (2008) and Hsieh and Klenow (2009) who focus on the the equality of marginal products of capital across firms; or 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.

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 above. 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 inevitably some 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. We show that the gap is around a factor of four on average in developing countries, 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 poor — and may in fact be 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 Uganda for example. In Uganda 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. To begin with, home-consumed production of agricultural goods does fall within the production boundary of the UN System of National Accounts, 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 is to measure the area planted and yield of most crops, which can be surveyed at the national level, and to subtract off the value of purchased intermediate inputs.[4] 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. Detailed procedures also govern the allocation of output to different time periods.[5] 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 to be "the prices realized by them 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.[6]

3.3. Measurement of Labor in Agriculture

Mis-measurement of labor in agriculture is another key issue. Because agriculture in poor countries falls largely into the informal sector, there are not detailed data on employment of the kind that might be found in the formal manufacturing sector. There are unlikely to be payroll records or hu-

[4] For 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.

[5] The 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 for animal products.

[6] Nevertheless, many development economists find it difficult to believe that national income accounts data for developing countries can offer an accurate picture of sectoral production. To revisit these concerns later in Section 5, where we construct, for a number of countries, alternative measures of value added in agriculture, using household survey data. Although a number of the same methodological challenges arise (e.g., with respect to the valuation of home-consumed agricultural outputs), we find that the large agricultural productivity gaps remain.

man 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 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.

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 also 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 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.[7] 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 this results from differences in the structure of employment.

[7]This is an issue studied in some detail by Vollrath (2010) recently, and dates back to the dual economy theory of Lewis (1955), in which he posited a surplus of labor in agriculture.

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. In this were true, we would be overestimating the labor input in agriculture compared to non-agriculture, as the productivity of agriculture workers would be lower on average than for other workers. We address these possibilities directly in Section 4.2, to follow.

3.4. Raw Agricultural Productivity Gap Calculations

With these measurement issues clear, this section describes the sample of countries, our data sources, and it presents the "raw," or unadjusted, agricultural productivity gaps.

The Sample and Data Sources

Our sample of countries includes all *developing countries* for which data on the shares of employment and value added in agriculture is available. By developing countries, we mean countries for which income per capita, in US Dollars expressed at exchange rates, is below the mean of the world income distribution.[8] We restrict attention to countries with data from 1985 or later, and the majority of countries have data from 1995 or later. We end up with a set of 112 countries which have broad representation from all geographic regions and per-capita income levels within the set of developing countries. In each country we focus our attention on the most recent year in which data is available.

Our main source of data on agriculture's share of employment is the World Bank's World Development Indicators (WDI). We supplement these with employment data by sector compiled by the International Labor Organization (ILO). The underlying source for all these data are nationally representative censuses of population or labor force surveys conducted by the countries' statistical agencies.[9] One advantage of using surveys based on of 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 to be the "economically active population" in each sector. The economically active population refers to all persons who are unemployed or employed and supply any labor in 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.

[8] This cutoff is arbitrary; however the results of the analysis do not differ meaningfully if we use the classifications of the World Bank or other international organizations.

[9] We exclude a small number of countries in which employment shares in agriculture are based on non-nationally representative surveys, such as urban-only samples, or surveys of hired workers.

In general, employed workers are classified into sectors by their reported main economic activity, and unemployed workers are classified according to their previous main economic activity.

Our data on agriculture's share of value added come from the WDI. The underlying sources for these data are the national income and product accounts from each country. In all cases these data are expressed at current-year local currency units.[10] Industry classifications are made in the majority of cases using the International Standard Industrial Classification System (ISIC).

Raw Agricultural Productivity Gaps

Table 2 reports summary statistics for the raw APGs for our set of developing countries. We refer to these as raw APGs because they are before any adjustments (e.g. for hours worked), unlike the calculations that follow. The first data column describes the APG distribution for the entire sample of 112 countries when weighting by population, our preferred method. Across all countries, the mean value of the gap is 4.0, implying that value added per worker is approximately four times higher in non-agriculture than in agriculture. The median is slightly lower, at 3.7. Even at the 5th percentile of the distribution, the gap is greater than unity (1.7), implying that in almost all countries for which we have data, the simple prediction of (3) is inconsistent with the data. At the 95th percentile of the distribution, the gap is 5.4.

The second data column of Table 2 presents the same statistics when not weighting. The results are largely similar, with the unweighted mean APG at 3.6 and the median at 3.0. When not weighting, the range of gaps is larger across countries. The 5th percentile is 1.1, and the 95th percentile is now 8.8. Still, the majority of countries have gaps above unity, unlike as predicted by the simple model.

Figure 6 shows histograms of the APG by region. Africa has the highest average APG, and all countries with gaps above ten (Burkina Faso, Chad, Guinea, Madagascar and Rwanda) are in Africa. Still, in all regions – Africa, Asia, the Americas and Europe – the average country is well above unity, and each region has a number of countries with gaps above four. These data suggest that the large gaps are not confined to developing countries in one area of the world.

Relative to the discussion in Section 2, it is abundantly clear that the data are not consistent with (4), which would give an APG of one. The raw data suggest very large departures from parity in sectoral productivity levels among these developing countries.

Differences of this magnitude are striking. If we take these numbers literally, they raise the pos-

[10] An alternative would be to use a single set of international comparison prices to value the agricultural output of each country. This might be more relevant if we were making comparisons of agricultural productivity across countries as in Caselli (2005), Restuccia, Yang, and Zhu (2008) or Vollrath (2009); in the current paper, however, we are only interested in comparing sector value added per worker within each country.

sibility of very 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 input and output? In the following sections, we discuss the underlying data and consider a number of ways in which mis-measurement may occur. We will also compare the magnitude of these possible mis-measurements with the observed gaps in productivity.

4. Accounting for Agricultural Productivity Gaps

In this section, we report the results of efforts to adjust the productivity gaps to account for some obvious differences in the quantity and quality of labor inputs. We base this analysis on a new database that compiles country-level data on schooling, labor, and other variables. All of the data used in this section originate in nationally-representative censuses of population and labor force surveys, with underlying observations at the individual level.

Our data comes in part from International Integrated Public Use Microdata Series (I-IPUMS), from which we use micro-level census data from 44 developing countries around the world (Minnesota Population Center (2011)). 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 with 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. 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 1 details the sources and data used in each of the 112 developing countries in our data.

4.1. Sector Differences in Hours Worked

In this section we ask whether the sectoral productivity gaps can be explained by differences across sectors in hours worked. We find that in most of the countries for which we have hours data, there are only modest differences in hours worked by sector; on average, workers in non-agriculture supply around 1.2 times more hours than workers in agriculture. We conclude that hours worked differences are unlikely to be the main cause of the large APGs we observe.

Specifically, we measure hours worked for all workers in the labor force, including those unemployed during the survey, for whom we count zero hours worked. The typical survey asks hours worked in the week or two weeks prior to the survey, although some report average hours worked in the previous year.[11] We classify people as workers in either agriculture or non-agriculture, ac-

[11] One potential limitation of using hours in the previous week or two weeks is if the survey was conducted during

cording to their main reported economic activity. For unemployed workers not reporting a main economic activity, we classify them as agricultural if they live in rural areas, and as non-agricultural if they live in urban areas.

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 1 lists the countries for which we use urban-rural status to construct our hours measures. In these countries, as in the others, we count unemployed workers as having worked zero hours.[12] 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. One consolation is that, in the developing world, most workers in urban areas work in non-agricultural activities, and most rural workers work in agriculture. Furthermore, 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 6 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 where average hours worked are identical in the two sectors. Similarly, the other two lines represent factor of 1.5 and 2.0 differences in hours worked. Most of the observations are clustered closely around the 1.0 line, and all but a few are below the 1.5 line, meaning that hours worked differences across sectors are generally modest. An arithmetic average across countries gives a factor 1.2 difference in hours worked in non-agricultural compared to agriculture.

This pattern does not vary much across regions, with average ratios of 1.2 for developing countries in Africa, Europe, and Asia (Table 4), and an average ratio of 1.0 in the Americas. Uganda and Rwanda have the most pronounced differences in hours worked, with roughly 1.7 times as many hours worked in non-agriculture as agriculture in these countries. Notably, these countries also have large APGs.[13] So while hours worked differences overall do not seem to explain much of the large APGs, in some countries lower hours worked in agriculture seems to be an important part of their large measured gaps.

4.2. 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

intense work periods, such as harvest or planting periods, or off periods, such as right after the harvest. In general, the surveys are conducted over many months or even years, however. In all of the countries for which we can make the calculations, roughly similar numbers of households were surveyed in each month of the survey period.

[12] Our results change very little when using average hours among only employed workers.

[13] Jordan and Armenia are also outliers, although neither has a large APG or agricultural employment share.

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 agriculture and non-agriculture workers. While both sets of calculations have their limitations, ours improve on those of Vollrath (2009) in several dimensions. Most importantly, our calculations come from nationally representative censuses or surveys with direct information on educational attainment by individual.[14] We also end up with estimates for a much larger set of countries, and 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. As for our hours measures, we use all employed or unemployed people in the agricultural and non-agricultural sectors when possible, and otherwise we use urban-rural status. When direct measures of years of schooling completed are available, we use those. When they are not, we impute years of schooling using educational attainment data. Table 1 details which countries use years of schooling directly and which use educational attainment data. These imputations are likely to yield noisy measures of years of schooling of course, as "some primary 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 noisiness should in principle not systematically bias our measures of average years of schooling by sector.

The first panel of Figure 3 shows our results for the 97 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 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.

Table 4 shows that average years of schooling in non-agriculture, for all the countries with available data, is 2.0 times as high as in agriculture. This ratio varies across region: in developing countries in Europe, the difference is a factor of just 1.4, driven by the former Soviet block countries; in Asia

[14] Those used by Vollrath (2009) are imputed using school enrollment data.

and the Americas the ratio is around 2.0, and in Africa, schooling levels are 2.8 times higher in non-agriculture than in agriculture.[15] Thus, in many countries, human capital differences have the potential to explain some substantial fraction of the APGs.

To turn years of schooling into human capital, we heed the findings of Banerjee and Duflo (2005) and Psacharopoulos and Patrinos (2002), who conclude that each year of schooling increases wages by around 10%. They arrive at their estimates using a large number of Mincer return estimates from countries around the world.[16] In particular, we assume that average human capital in sector j in country i is $h_{j,i} = \exp(0.10 \cdot s_{j,i})$ where $s_{j,i}$ is average years of schooling in sector j, country i.

The bottom panel of Figure 3 shows the results of our calculations of average human capital by sector. 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, according to the Mincer return estimates discussed by Banerjee and Duflo (2005) and Psacharopoulos and Patrinos (2002), having (say) twice as many years of schooling implies having considerably less than twice as much human capital. The weighted average across countries is a factor 1.4 difference in human capital of across the two sectors. As can be seen in Table 4, the average is a little higher in the Americas at 1.5, and lower in Europe at 1.3.[17]

4.3. Adjusting for Education Quality using Literacy Rates

One limitation of the analysis above is that our procedure treats years of schooling among agriculture workers as equally valuable as those among non-agriculture workers. There is evidence, however, that the quality of schooling in rural areas in many developing countries is below that of schooling in urban areas. For example Williams (2005) and Zhang (2006) provide evidence that literacy rates and test scores in mathematics and reading are most often lower in rural schools than urban ones. Thus, our estimates above may tend to overestimate the human capital level of agriculture workers, who in general received their schooling from lower-quality rural schools.

[15] It is worth noting again that our data set is limited to countries that have income per capital less than half the level in the U.S. Thus, when we refer to countries in Europe or the Americas, we are explicitly excluding advanced economies.

[16] Our results change very little when using the concave human capital function of schooling used by Caselli (2005), Hall and Jones (1999), and Herrendorf and Valentinyi (Forthcoming) in their accounting exercises.

[17] By comparison, Vollrath (2009) finds that human capital in non-agriculture is higher by a factor of only around 1.2, averaging across countries. In other words, we suggest that more of the agricultural productivity gaps can be explained by human capital differences. The proximate reason for this is that our measures yield higher levels of schooling in both sectors than Vollrath's, but we find a substantially higher level of schooling in non-agriculture than he does, while our measures for the agricultural sector are only slightly higher.

In this section, we consider the effect of adjustments for education quality differences. Here we present a simple new method of adjusting for quality differences in schooling among agricultural and non-agricultural workers using literacy data. The basic idea is that literacy, particularly in primary schools, is one of the main components of the human capital that students receive through schooling. Thus literacy rates for workers by years of schooling completed in the two sectors are informative about quality differences in schooling received by workers in the two sectors.

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 $s = 0, 1, 2....$ If the quality of schooling received 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 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, and hence an average year of schooling received by the non-agricultural workers must have been more effective than an average year received by the agricultural workers.

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 1 year.) The differences in literacy are largely absent by about 6 years of schooling completed, with virtually all workers literate by then, hence we cut the graph off then.

To pin down how much more effective a year of urban education is than a rural year in country i, our method is the following. 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-agriculture workers, and set γ_i to the value that solves

$$\min_{\gamma} \sum_{s=1}^{\bar{s}} \left(\tilde{\ell}_i^n(\gamma s) - \tilde{\ell}_i^a(s) \right)^2 . \tag{5}$$

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, over a range of s values up to some value $\bar{s}$. Since primary school ends at 5 years in most countries, and since most workers are literate by then, setting $\bar{s}$=5 seems warranted. In the example of Uganda, we find that $\gamma_{UGA} = 0.82$, meaning that a each year of schooling for agriculture workers is worth 82% of a year of schooling for the typical non-agriculture worker in terms of acquiring literacy. We assume therefore that a year of schooling for agriculture workers is worth 82% of a year of schooling for non-agriculture workers in terms of acquiring human capital.

Table 3 shows the results of each of similar calculations that we made for the 17 countries with available data. The average estimate is 0.87, suggesting real but modest differences in schooling quality across countries. All but one country has an estimate of γ less than one. Only Tanzania has an estimate above one; why rural schools appear to fare better than urban ones is a question for which we do not yet have a clear answer. 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 are other notably low estimates, all around 0.75.

Figure 5 displays the ratios of human capital in non-agriculture to agriculture using the quality-adjusted agriculture human capital estimates, calculated as $h^q_{a,i} = \exp(\hat{\gamma}_i s_{a,i})$ for each country i, and the original unadjusted estimates. Countries above the 45-degree line are those that have higher ratios once the quality adjustments are made. As can seen from the figure, the differences in ratios are modest in general. Many of the adjusted ratios are virtually identical to the unadjusted ones, and the biggest adjustments are small, on the order of a factor 0.2 increase (for Vietnam) or smaller.[18]

We conclude that these education quality adjustments, while perhaps crude, suggest that quality differences in schooling do not substantively alter the our findings regarding human capital differences by sector. In the average developing country, human capital per worker is 1.4 times as high in the non-agriculture sector as the agriculture sector, and this ratio is basically unchanged when schooling quality using our method.

4.4. Cost-of-Living Differences

Next we turn to cost-of-living differences between rural and urban areas. The prediction of Equation 4 is that average productivities should be equalized across sectors. But this prediction is for real measures of average productivity, and it assumes that the nominal income earned by workers in each sector has the same purchasing power. In reality, there are many reasons to suspect that the cost of living is lower in rural areas, which have lower population density and easier access to food supplies.

Fortunately, proxies for the cost of living in rural and urban areas are available for a large number of developing countries. Ravallion, Chen, and Sangraula (2009) use the World Bank's country

[18] We find that even when assuming a counterfactually low ratio of one year of rural schooling to 0.5 years of urban schooling, the quality adjustments lead to fairly modest differences in human capital ratios. Under this assumption, the average ratio among these 17 countries rises from 1.4 to just 1.6.

studies from a set of 78 developing countries to compute the cost of the basket of goods consumed by households living on $1 per day in rural and urban areas. While this basket is not necessarily the same as the basket of the average household in the countries studied, Ravallion, Chen, and Sangraula (2009) argue that most poor households (who consume mostly food) have a basket that is quite similar, and hence a cost of living that is similar. For example, they found very similar urban-rural cost of living differentials when re-computing the cost of a basket consumed by households living on $2 per day.[19]

Figure 6 shows a histogram of the ratio of cost of living in urban areas to rural areas. As can be seen in the figure, virtually no countries have lower prices in rural areas, and the average developing country has an urban cost of living that is roughly 1.3 times that of rural areas. The median is slightly lower than the mean. Thus, part (but not all) of the APGs may reflect differences in sector costs of living.

4.5. Adjusted APGs

We now compute the "adjusted" agricultural productivity gaps, which take into consideration the sector differences in hours worked, human capital, and cost of living. We do not have all these data for all the countries in our sample, and hence we proceed in two ways. First, we compute the adjusted APG for each of the 35 countries for which we have complete data. Second, we compute the APG for every country in our sample by imputing any missing data. We do this by assigning any missing value to be the weighted average ratio across all other countries with data.[20] For each country, we construct the adjusted APG by dividing the raw APG by the ratio of hours worked, the ratio of human capital, and the urban-rural price ratios.

Table 5 shows summary statistics of the adjusted APG distributions for countries with complete data and then all countries in our data. For both sets of countries, the mean adjusted gap is 1.9. The means are 1.9 and 1.8, respectively. Thus the typical country has an APG around half as high once all our adjustments are made. The 5th percentiles are 1.1 and 0.8, while the 95th percentiles are now down to 2.9 and 2.6 respectively. Thus, in virtually all countries, adjusted APGs are substantially lower than their raw counterparts. Figure 6 illustrates this decline in more detail by plotting the distributions of APGs before and after adjustments for the two sets of countries.

Figure 9 provides more detail on how the adjusted and raw APG values differ for the countries

[19] In principle, many countries collect consumer price data at many locations, in order to construct consumer price indices. These data should make it possible to calculate cost-of-living differences between rural and urban areas. In practice, however, the underlying data are seldom publicly available.

[20] Most of the imputed values are for ratios of hours worked, since hours measures were available for the fewest countries. Our results do not change substantially when using alternative imputation methods, such as projecting missing data using GDP per capita and regional dummies.

for which we have complete data. The top panel of Figure 9 shows all countries. Most notably, Rwanda and Zambia have big raw gaps, of 14 and 9.5 respectively, and much smaller gaps after our adjustments, with both countries below 4. The bottom panel provides a “close up” of the same countries minus those with raw APG values of over 7. Now one can see that Lesotho and Uganda have initial gaps of around 7, and adjusted gaps of around 2 and 3 respectively. Interestingly, the remainder of the countries tend cluster along a ray of slope one-half from the origin, suggesting that our adjustments explain around one half of their raw gaps.

While, on the one hand, explaining roughly one half the raw APG measures represents success for our adjustments, the remaining gap of around two is puzzlingly large. The implication is that there should be large income gains from moving workers out of agriculture and into other economic activities. Thus, we conclude that our adjustments thus far take us part of the way – but only part of the way – towards explaining the differences in productivity between sectors. We now turn to several other potential explanations of the remaining gaps.[21]

4.6. Sector Differences in Labor's Share in Production

One maintained assumption of the simple model is that labor shares in production are the same in agriculture and non-agriculture. We now relax this assumption and ask whether sector differences in labor shares could account for much of the remaining gap. We argue that evidence suggests that it cannot, and that assuming equal labor shares in the two sectors does not change the nature of our analysis in any important way.

Consider a variant of the production function in (6) where the importance of labor and other inputs in production differs across sectors:

$$Y_a = L_a^{\theta_a} K_a^{\phi_a} X_a^{1-\theta-\phi} \quad \text{and} \quad Y_n = L_n^{\theta_n} K_n^{\phi_n} X_n^{1-\theta_n-\phi_n}. \tag{6}$$

One can show that the firms' first order conditions imply that sector differences in value added per worker are given by the ratio of the Cobb-Douglas elasticities:

$$\frac{VA_n/L_n}{VA_a/L_a} = \frac{Y_n/L_n}{p_a Y_a/L_a} = \frac{\theta_a}{\theta_n}. \tag{7}$$

Thus, we could explain the remaining sectoral differences in average labor productivity if θ_n is

[21] One candidate explanation is that that the gaps have arisen recently, and workers have simply not had sufficient time to reallocate across sectors in response. However, we find that for virtually all countries for which historical data is available from the WDI, the average APG in the period 1985 to the present is similar in magnitude to the average APG in the period 1960 to 1984.

approximately half as large as θ_a. Is this a plausible explanation?

Evidence from National Accounts data

One source of data on sectoral labor shares is the income side of the national accounts. In this account, GDP is divided into different types of income, with the principal categories being employee compensation, the operating surplus of firms, depreciation, and indirect taxes and subsidies. Not all countries report the income side of the national accounts by sector. As a result, there are few systematic studies of sectoral labor shares across countries.

Some data are available on the employee compensation shares in agriculture (as in Gollin (2002), Table 5), but these do not accurately reflect labor shares in sectors where much of the labor force consists of family labor. (In practice, employee compensation does not usually include the mixed income of the self-employed.) Moreover, as pointed out by Valentinyi and Herrendorf (2008), careful calculations of sectoral labor shares require adjustment for the cross-sector flows of intermediate goods, making these calculations relatively complicated. Nevertheless, there are strong reasons to believe that aggregate labor shares are basically the same in rich and poor countries.

As Gollin (2002) points out, labor shares – once adjusted for the mixed income of the self-employed – vary relatively little across countries, and the variation is not highly correlated 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. To quantify this, in many poor countries, agriculture accounts for 25-50% of GDP, while in rich countries it may be only 1% of GDP. If poor countries also had a labor share in agriculture that was half as large as the labor share in non-agriculture, then the aggregate labor share in poor countries would be noticeably lower. For instance, if the labor share in non-agriculture was θ_n= 0.67 , as suggested in Gollin (2002), then a poor country with agriculture producing 30% of GDP would have an aggregate labor share of 0.57, compared with an aggregate labor share of 0.67 in a rich country.

What limited evidence on sectoral shares is available suggests that labor shares in agriculture are, if anything, lower than labor shares in non-agriculture. Gollin (2002) reports employee compensation shares of output for a set of countries with available data. The agricultural sector has the lowest shares of all the sectors in these data – perhaps because these measures frequently appear to exclude the imputed labor income of unpaid family workers.

Econometric Estimates of Factor Shares

A large empirical literature attempts to estimate agricultural production functions using cross-section, time series, and panel data. A dual literature estimates parameters using cost data. This literature is problematic because input use is fundamentally endogenous, so there are frequently puzzling signs and magnitudes in the coefficient estimates. Typically, however, these estimates find labor shares for agriculture that are *lower* than those for non-agriculture, not the other way around.

Fuglie (2010) summarizes the estimated cost shares of labor from a number of recent studies. Leaving aside studies of transition economies that have unusually low labor shares, Fuglie finds in his analysis that the labor shares in agricultural production functions tend to be low. He reports results from micro studies showing cost shares or production elasticities for labor that range from 0.31 (for sub-Saharan Africa) to 0.46 (India and Indonesia). Estimates for rich countries include figures of 0.20 for the U.S., 0.30 for the U.K., 0.39 for Japan, and 0.23 for South Africa.

These estimates are not directly comparable to the macro concepts of labor shares, because the calculations are often based on analyses in which the gross value of agricultural output is regressed on inputs of labor, land, and a range of intermediate goods. These include energy, chemicals, and seed – all of which would be viewed as intermediates in the national income and product accounts. Accordingly, we can take the labor share as a fraction of the summed shares of labor, land, structures, and machines.

Working with the data of Fuglie (2010) for China, India, Indonesia, Brazil, Mexico, and sub-Saharan Africa, the average share of labor relative to labor plus land and structures is 0.62. When the denominator is expanded to include machinery (and energy, which is lumped together with machinery), the labor share falls to an average of 0.58. Obviously it falls even farther if the denominator is expanded to include livestock, which are in some countries viewed as a form of agricultural capital. For the U.S., the labor share relative to labor, land, and structures is 0.51, and for the U.K., the number is 0.52.

Observations from Share Tenancy

Another source of information – perhaps less obvious – on labor's share in agriculture comes from the abundant literature on share tenancy contracts. 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. Share tenancy is an ancient and ubiquitous practice, and in some parts of the world – particularly developing countries in

Asia – share tenancy accounts for the majority of all tenanted farm land (Otsuka (2007))[22]. Share contracts typically specify the fraction of gross output that will be received by the landowner. The contracts normally also specify a set of cost-sharing arrangements for variable inputs (e.g., fertilizer, seed, and chemicals) and for the usage or rental costs of equipment.

The share contracts offer some insight into the cost shares of different factors of production. The operator provides all the labor, and 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) shares of labor, land, and capital. In practice, it may be difficult to arrive at precise calculations, because relationships between land owners and operators may be quite complicated. These relationships are often based on long-term leases with indefinite time horizons, so that they constitute repeated games. There are also many accounts of interlinked contracts, so that operators may also borrow from land owners for consumption purposes; land owners may provide some informal insurance; and cost sharing formulas may be modified informally. As a result, Jacoby and Mansuri (2009) write that "[o]utput and cost shares alone thus do not fully characterize the terms of the share-contract." 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% under share tenancy in many developing countries." They 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." Contemporary studies continue to find the 50-50 split to be comon in developing countries. 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. In Iowa, in 2007, 54% of all farm land is leased, and about 20-25% of the leased land is covered by share contracts. Of the land under share contracts in Iowa, 93% involved a 50-50 split of output and in 80-90% of the cases, there was also a 50-50 split of the costs of fertilizer, seed, herbicides, insecticides, and drying (Iowa State University (2007)). The 50-50 split is also found commonly in other parts of the U.S., but Canjels (1998) draws on a broader 1988 survey of U.S. farmers and concludes that other splits, with the owner's share ranging from 0.40 to 0.75, depending on the state and crop.

[22]Most land, however, is farmed by owners, cash rentals and share cropping represent a minority of total farm land.

Calculating the actual *ex post* shares of labor, land, and capital would require extremely detailed data on the sharing of particular costs as well as the various state-contingent adjustments that are made in the allocation of cost and revenue. Nevertheless, the predominance of the 50-50 split of both output and inputs would tend to suggest that labor's share in agricultural production is around one-half. Moreover, the recurrence of the same cost shares across countries and over time would tend to support the notion that the labor share can be modeled as identical across countries.

Summarizing the Evidence on Labor's Share in Agriculture

Taking the various approaches into account, it seems difficult to produce estimates of labor shares that are *higher* for agriculture than for non-agriculture. Under the circumstances, differences in labor shares are unlikely to play any large role in explaining the sectoral productivity gaps; if anything, it looks as though they should exacerbate those gaps.

5. Alternative Measures of Sectoral Productivity

One possible concern with the sectoral productivity measures reported above is the quality of the data. Many development economists are outspoken in characterizing national accounts data for poor countries as fanciful if not fictitious. Even more sober appraisals, such as that of Ravallion (2003), find that measures of consumption per capita derived from national income accounts differ in levels and growth rates from the alternative measures of consumption or expenditure that are reported in nationally representative household survey data.

To address this concern over data quality, we construct alternative measures of sectoral productivity that rely on measures of productivity derived entirely from household survey data. Although the measures reported above, in Section 4, did make some use of micro data, those measures simply relied on the micro data for hours and schooling to make adjustments to productivity measures that were derived from national income accounts. By contrast, in this section, we construct entirely new measures of productivity from detailed household survey data.

The reason for doing this is to address the possibility that national accounts measures of agriculture's share of output may be inaccurate. In particular, if the national accounts underestimate agricultural value added (or overestimate value added in non-agriculture), then our measures of average productivity may be erroneous. It is also possible that the national labor force surveys and census data are mistakenly assigning people to the agricultural sector, creating a form of mismeasurement that our previous adjustments cannot remedy.

Household surveys typically collect enormously detailed data on the characteristics and decisions of respondent families. We use these household survey data to construct a number of measures of labor and output. We conclude that, even though the household survey and national account evidence are based on entirely different methodologies, both provide evidence that productivity is substantially lower among agricultural workers than other workers.

5.1. Household Income and Expenditure Surveys

Over approximately the past twenty-five years, international institutions, in collaboration with national statistical authorities, have carried out detailed household surveys in numerous countries, including many developing countries. These surveys have been widely used in the development economics literature, and they are widely seen as providing high-quality data that offer an independent check on many statistical aggregates.

One of the largest collections of these household surveys is available through the World Bank's Living Standards Measurement Study (LSMS). The LSMS surveys typically involve the collection of detailed data at the household level (and occasionally at the 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 follow identical methodologies, but many of them contain information on variables of interest for our analysis.

The data from these surveys provide information on economic activity at the household level. Because individuals are asked about all their economic activities, the household surveys typically do a good job of depicting non-market activities, including agricultural production by quasi-subsistence households. In micro development economics, data from these household surveys are generally seen as representing a high standard for data quality.

5.2. Measuring Value Added from Household Surveys

To begin, we construct measures of value added per worker from the household surveys. For households that operate a farm (or other agricultural business), we define agriculture value added as

$$VA_{a,i} = y_{a,i} - INT_{a,i} \tag{8}$$

where $y_{a,i}$ is "gross agricultural output," defined below, and INT_i is the cost of intermediate goods, such as fertilizer or seeds. Gross agricultural output is measured as as follows. Let j index farm goods, and let there be J total goods. For household i, let $x_{i,j}^{home}$, $x_{i,j}^{market}$ and $x_{i,j}^{invest}$ be the total quantity of good j that is produced and subsequently consumed as home, sold in the market place, or re-invested (as seed) for the following year. Agricultural revenue is defined as

$$y_{a,i} = \sum_{j=1}^{J} p_j \left(x_{i,j}^{home} + x_{i,j}^{market} + x_{i,j}^{invest} \right), \tag{9}$$

where p_j is the farm-gate price of good j. In general, the household surveys report $x_{i,j}^{home}$, $x_{i,j}^{market}$ and $x_{i,j}^{invest}$ directly (in kilograms) for each j. For prices, in the case of goods j that were sold on the marketplace, we set p_j to be the sale price. For goods j that were not sold, we let p_j be the self-reported price that the household would have fetched if they sold their output.[23]

For households that operate a non-agricultural business, we define value added as

$$VA_{n,i} = y_{n,i} - INT_{n,i} \tag{10}$$

where $y_{n,i}$ is non-farm business revenue and $INT_{n,i}$ is the cost of intermediates used up in production. In general the households report $y_{n,i}$ or their profits directly.

While the surveys vary somewhat from country to country, 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 World Bank uses a local or regional average price. For households with non-agricultural income, typically $y_{n,i}$ is reported directly. Intermediate usage in the two sectors are either reported directly or reported as part of total input costs, i.e. along with wage payments to hired workers and rental payments to rented land or equipment. When total input costs are reported, we compute the fraction that is due to intermediates using the ratio of aggregate payments to hired labor and capital to aggregate input costs.[24]

For each country, we aggregate value added by sector across households to compute the share of aggregate value added coming from agriculture. We then compute the employment share of agriculture in each country, classifying each worker by her primary industry of employment. 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.

[23]There are some complexities involved in figuring out the appropriate prices to use. For households that sell some goods to the market, the quantities of the same goods that are consumed at home must (through revealed preferences) have a shadow value higher than the market price. Using market prices will underestimate the value of production, in this sense. Another issue that arises is the adjustment of market prices for transport costs. In principle, home consumed output should be valued at the farm-gate price, which is generally lower than the market price for goods that the household sells (and is generally higher than the market price for goods that the household purchases).

[24]Specifically, let $COST_{s,i}$ be the total input costs of household i for production in sector s. Denote by α_s the ratio of aggregate wage earnings plus payments for rented land or equipment to aggregate input costs in sector s. Then $INT_{s,i} = y_{s,i} - (1 - \alpha_s)COST_{s,i}$.

5.3. Results: Value Added per Worker

Table 6 shows the results of the calculations for four countries: Cote d'Ivoire, Guatemala, Pakistan and South Africa.[25] The first data column shows the share of workers in agriculture according to the micro survey data. The second and third data columns show agriculture's share in value added according to the macro data (the national accounts) and the micro data. In spite of the differences in methodology, the macro and micro sources give similar answers as to the share of value added originating in agriculture. In Cote d'Ivoire for example, the macro data implies a value added share of 32.0 percent, while the micro data implies a share of 37.7 percent.

The final two columns of Table 6 show the raw APG measures implied by the macro and micro data. What varies in these two columns is whether the value added share comes from the national accounts data (macro) or household surveys (micro). In both cases the employment share of agriculture comes from the micro surveys; in none of the years was a separate survey of households conducted. Furthermore, the basic methodology for computing employment shares in agriculture is roughly the same in the LSMS surveys as the labor force surveys and population censuses used to construct the macro APG numbers.

In all four countries, the macro and micro raw APG measures are similar. Cote d'Ivoire and Pakistan have the biggest gaps in the macro data, of 4.3 and 4.2, and in the micro data they also show up as having the biggest gaps, at 4.0 and 4.3. Guatemala has a macro gap of 3.4 and a micro gap of 3.3. South Africa has a macro gap of 1.6 and a micro gap of 1.5. At least for these countries, we conclude that large gaps in value added per worker are found in household survey evidence, just as in the national accounts.

5.4. Income and Expenditure per Worker

One challenge to the use of the value added data described above is that it could be biased, due to households misreporting their income. Deaton (1997) claims that expenditure data are more accurate than the income data. In Table 7 we address this concern by computing the ratio of expenditure per worker by sector, in addition to value added per worker (as above) and income per worker.

We find that in all four countries, gaps in value added and income per worker are similar in magnitude, and in all but one country expenditure per worker gaps are of similar magnitude. In Pakistan, the exception, the expenditure per worker gap is substantially smaller than that of value added or income, at 1.9, compared to 4.3 for value added per worker. While Pakistan's gap is still sizable, its smaller magnitude suggests that large remittances from non-agriculture households to agriculture

[25] We are in the process of making similar calculations for roughly a dozen other countries.

areas may play an important role. In future work we can test this proposition (using remittance data) and expand our analysis to more countries. The conclusion thus far, however, is that the large agricultural productivity gaps found in the macro data are very much present in micro data as well.

6. Conclusion

According to national accounts data from developing countries, value added per worker is on average four times higher in the non-agricultural sector than in agriculture. This agricultural productivity gap, when taken at face value, suggests that labor is greatly misallocated in developing countries. In this paper we ask to what extent the gap is still present when better measures of inputs and outputs are taken into consideration. To do so we construct a new data set for a large number of developing countries, with measures of hours worked and human capital per worker by sector, urban-rural cost-of-living differences, and alternative measures of value added per worker constructed from household income survey evidence.

We find that even after taking all these measurement issues into consideration, a puzzlingly large agricultural productivity gap remains. Output per worker in non-agriculture still appears to be nearly twice as high as in agriculture.

A number of measurement issues remain. One is that 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. A second measurement problem is that the costs of living for rural semi-subsistence farmers may be overstated by price indices based on local market prices. Many households in poor countries may in fact face very low prices for a range of home-produced goods, so that their realized utility levels are higher than would be suggested by income and expenditure data.

Our results suggest that the typical resident of the developing world could roughly double her real income by moving out of the agriculture sector. Why don't workers in developing countries move out of agriculture to close this gap? Answering this question seems like an important step in understanding economic development.

References

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.

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 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.

DEATON, A. (1997): *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*. The World Bank.

FAO (1996): *A System of Economic Accounts for Food and Agriculture*. Food and Agriculture Organization.

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.

——— (2007): "The Food Problem and the Evolution of International Income Levels," *Journal of Monetary Economics*, 54, 1230–1255.

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.

HERRENDORF, B., AND T. SCHOELLMAN (2011): “Why is Agricultural Productivity so Low in the United States?,” Discussion paper, ASU.

HERRENDORF, B., AND A. VALENTINYI (Forthcoming): “Which Sectors Make the Poor Countries So Unproductive?,” *Journal of the European Economic Association*.

HSIEH, C.-T., AND P. J. KLENOW (2009): “Misallocation and Manufacturing TFP in China and India,” *Quarterly Journal of Economics*, 124, 1403–1448.

IOWA STATE UNIVERSITY (2007): “Survey of Iowa Leasing Practices,” http://www.extension.iastate.edu/Publications/FM1811.pdf.

JACOBY, H., AND G. MANSURI (2009): “Incentives, Supervision, and Sharecropper Productivity,” *Journal of Development Economics*, 88(2), 232–241.

KUZNETS, S. (1971): *Economic Growth of Nations: Total Output and Production Structure*. Harvard University Press.

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.

MINNESOTA POPULATION CENTER (2011): “Integrated Public Use Microdata Series, International: Version 6.1,” Minneapolis: University of Minnesota.

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 (2009): “Dollar a Day Revisited,” *World Bank Economic Review*, 23(2), 163–184.

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.

VALENTINYI, A., AND B. HERRENDORF (2008): "Measuring Factor Income Shares at the Sectoral Level," *Review of Economic Dynamics*, 11, 820–835.

VOLLRATH, D. (2009): "How Important are Dual Economy Effects for Aggregate Productivity?," *Journal of Development Economics*, 88(2), 325–334.

——— (2010): "Measuring Aggregate Agricultural Labor Effort in Dual Economies," Unpublished Manuscript, University of Houston.

WILLIAMS, J. H. (2005): "Cross-National Variations in Rural Mathematics Achievement: A Descriptive Overview," *Journal of Research in Rural Education*, 20(5).

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

Table 1: Data Sources and Descriptions

Country	Variable	Year	Source	Ag/Non-Ag or Urban/Rural
Albania	Schooling[1]	2005	Multiple Indicator Cluster Survey (EPDC)	U
Argentina	Schooling	2001	Census of Population and Housing (IPUMS)	A
Armenia	Schooling[1]	2001	Population and Housing Census (IPUMS)	A
	Hours	2008	Report on Labour Force and Informality	A
Azerbaijan	Schooling[1]	2006	Demographic and Health Survey (EPDC)	U
Bangladesh	Schooling[1]	2007	Demographic and Health Survey (EPDC)	U
	Hours	1989	Labour Force Survey (ILO)	A
Belarus	Schooling	1999	Population Census (IPUMS)	A[3]
Belize	Schooling[1]	2006	Multiple Indicator Cluster Survey (EPDC)	U
Bhutan	Schooling[1]	2005	Population and Housing Census	U
	Hours[4]	2007	Living Standard Survey	U
Bolivia	Schooling[1]	2001	Census of Housing and Population (IPUMS)	A
	Hours	2000	Mecovi Survey	U
Botswana	Schooling[1]	1996	Labour Force Survey	A
	Hours	1996	Labour Force Survey	A
Brazil	Schooling	2000	Demographic Census (IPUMS)	A
	Hours	2000	Demographic Census (IPUMS)	A
Bulgaria	Schooling	2003	Living Standards Measurement Study	A
	Hours	2003	Living Standards Measurement Study	A
Burkina Faso	Schooling[1]	2006	Multiple Indicator Cluster Survey (EPDC)	U
Burundi	Schooling[16]	1998	Enquete Prioritaire	A
Cambodia	Schooling	1998	General Population Census (IPUMS)	A
	Hours[4]	2001	Labour Force Survey	U
Cameroon	Schooling[1]	2006	Multiple Indicator Cluster Survey (EPDC)	U
Central African Rep.	Schooling[1]	2006	Multiple Indicator Cluster Survey (EPDC)	U
Chad	Schooling[1]	2004	Demographic and Health Survey (EPDC)	U
Chile	Schooling	2002	Population and Housing Census (IPUMS)	A
	Hours	2002	National Employment Survey (ILO)	A
China	Schooling[1]	1990	National Population Census (IPUMS)	A
Colombia	Schooling	2005	General Census (IPUMS)	A
Costa Rica	Schooling	2000	Population and Housing Census (IPUMS)	A
	Hours[3]	2000	Multi-Purpose Household Survey (ILO)	A
Cote D’Ivoire	Schooling	1988	Living Standards Survey	A
	Hours	1988	Living Standards Survey	A

Cuba	Schooling[1]	2002	Population and Dwelling Census (IPUMS)	A
Dominican Republic	Schooling[1]	2007	Demographic and Health Survey (EPDC)	U
	Hours	2007	Encuesta de Fuerza de Trabajo (ILO)	A
Ecuador	Schooling	2001	Census of Population and Dwelling (IPUMS)	A
	Hours	2001	Census of Population and Dwelling (IPUMS)	A
Egypt	Schooling[1]	2000	Demographic and Health Survey (EPDC)	U
El Salvador	Schooling[1]	2006	Encuesta de Hogares de Propositos Multiples	U
Ethiopia	Schooling[1]	2005	Demographic and Health Survey (EPDC)	U
	Hours	2005	Labour Force Survey	U
Fiji	Schooling[17]	1996	Census of Population and Housing	U
	Hours[4]	2005	Employment and Unemployment Survey	U
Gabon	Schooling[1]	2000	Demographic and Health Survey (EPDC)	U
The Gambia	Schooling[1]	2006	Multiple Indicator Cluster Survey (EPDC)	U
Georgia	Schooling[1]	2005	Multiple Indicator Cluster Survey (EPDC)	U
Ghana	Schooling	2000	Population and Housing Census (IPUMS)	A
	Hours	2000	Population and Housing Census (IPUMS)	A
Guatemala	Schooling	2010	National Survey of Employment and Income	U
Guinea	Schooling[13]	1996	Census of Population and Housing (IPUMS)	A
Guyana	Schooling[1]	2005	Demographic and Health Survey (EPDC)	U
Honduras	Schooling[1]	2005	Demographic and Health Survey (EPDC)	U
India	Schooling[1]	2004	Socio-Economic Survey (IPUMS)	A
Indonesia	Schooling[1]	2007	Demographic and Health Survey (EPDC)	U
	Hours	2006	National Labour Force Survey (ILO)	A
Iran	Schooling[1]	2006	Census of Population and Housing (IPUMS)	A
Iraq	Schooling[1]	1997	Population Census (IPUMS)	A
	Hours	2007	Household Socio-Economic Survey (LSMS)	A
Jamaica	Schooling	2001	Population and Housing Census (IPUMS)	A
	Hours	2001	Population and Housing Census (IPUMS)	A
Jordan	Schooling[1]	2004	Population and Housing Census (IPUMS)	A
	Hours	2004	Population and Housing Census (IPUMS)	A
Kazakhstan	Schooling[1]	2006	Multiple Indicator Cluster Survey (EPDC)	U
Kenya	Schooling[13]	1999	Population and Housing Census (IPUMS)	A
	Hours	2006	Integrated Budget Household Survey	U
Kyrgyz Republic	Schooling[1]	1999	Population Census (IPUMS)	A
Lao PDR	Schooling[1]	2006	Multiple Indicator Cluster Survey (EPDC)	U
Lesotho	Schooling[1]	2004	Demographic and Health Survey (EPDC)	U
	Hours[4]	2008	Integrated Labour Force Survey	U

Liberia	Schooling[1]	2007	Demographic and Health Survey (EPDC)	U
	Hours[4]	2010	Labour Force Survey	U
Lithuania	Schooling[179]	2000	Population and Housing Census	U
Macedonia	Schooling[1]	2005	Multiple Indicator Cluster Survey (EPDC)	U
Madagascar	Schooling[1]	2008	Demographic and Health Survey (EPDC)	U
Malawi	Schooling[1]	2008	Population and Housing Census (IPUMS)	A
	Hours	2005	Second Integrated Household Survey	A
Malaysia	Schooling[1]	2000	Population and Housing Census (IPUMS)	A
	Hours	2007	Labour Force Survey (ILO)	A
Maldives	Schooling	2009	Demographic and Health Survey (EPDC)	U
Mali	Schooling	1998	Census of Population and Housing (IPUMS)	A
Marshall Islands	Schooling	1994	Multi-Subject Household Survey	A
Mauritius	Hours[5]	2009	Continuous Multi-Purpose Household Survey	A
Mexico	Schooling	2000	Population and Dwelling Count II (IPUMS)	A
	Hours	2000	Population and Dwelling Count II (IPUMS)	A
Moldova	Schooling[1]	2004	Demographic and Health Survey (EPDC)	U
Mongolia	Schooling[1]	2000	Population and Housing Census (IPUMS)	A
Morocco	Schooling[1]	2004	Demographic and Health Survey (EPDC)	U
Namibia	Schooling[1]	2006	Demographic and Health Survey (EPDC)	U
Nepal	Schooling[1]	2001	National Population Census (IPUMS)	A
	Hours[4]	2008	Labour Force Survey	A
Nicaragua	Schooling[1]	2001	Demographic and Health Survey (EPDC)	U
Nigeria	Schooling[1]	2008	Demographic and Health Survey (EPDC)	U
	Hours[4]	2009	Labour Force Survey	U
Pakistan	Schooling[1]	1998	Housing and Population Census (IPUMS)	U
	Hours[4]	2009	Labour Force Survey	U
Panama	Schooling[1]	2000	Census of Population and Housing (IPUMS)	A
	Hours	2001	Continuous Household Survey (ILO)	A
Papua New Guinea	Schooling[178]	2000	Census National Report	U
Paraguay	Schooling[17]	2002	Censo Nacional de Poblacion y Vivienda	U
Peru	Schooling	2007	Census of Housing and Population (IPUMS)	A
	Hours	2007	Estadisticas del Mercado de Trabajo	U
Philippines	Schooling	1990	Census of Population and Housing (IPUMS)	A
	Hours	1990	Labour Force Survey (ILO)	A
Romania	Schooling[1]	2002	Population and Housing Census (IPUMS)	A
	Hours[3]	2002	Population and Housing Census (IPUMS)	A

Rwanda	Schooling[1]	2002	Census of Population and Housing (IPUMS)	A
	Hours	2006	Integrated Living Conditions Survey	A
Saint Lucia	Schooling[1]	1991	Population and Housing Census (IPUMS)	A
	Hours	1991	Population and Housing Census (IPUMS)	A
Sao Tome and Principe	Schooling[1]	2009	Demographic and Health Survey (EPDC)	U
Senegal	Schooling[3]	2002	Census of Population and Housing (IPUMS)	A
Serbia	Schooling[1]	2006	Multiple Indicator Cluster Survey (EPDC)	U
Sierra Leone	Schooling[1]	2004	Population and Housing Census (IPUMS)	A
	Hours	1989	Labour Force Survey	A
South Africa	Schooling	2007	Community Survey (IPUMS)	A
	Hours	2009	Labour Market Dynamics in South Africa	A
Sri Lanka	Schooling[19]	2001	Census of Population and Housing	U
	Hours[4]	2009	Labour Force Survey	A
Sudan	Schooling[1]	2008	Population and Housing Census (IPUMS)	A
Suriname	Schooling[1]	2006	Multiple Indicator Cluster Survey (EPDC)	U
Swaziland	Schooling[1]	2006	Demographic and Health Survey (EPDC)	U
	Hours[5]	2008	Labour Force Survey	A
Syria	Schooling[1]	2006	Multiple Indicator Cluster Survey (EPDC)	U
	Hours	2010	Labour Force Survey	A
Tajikistan	Schooling[1]	2005	Multiple Indicator Cluster Survey (EPDC)	U
Tanzania	Schooling[1]	2002	Population and Housing Census (IPUMS)	A
	Hours	2009	Integrated Labour Force Survey	A
Thailand	Schooling	2000	Population and Housing Census (IPUMS)	A
Tonga	Schooling[17]	2006	Census of Population and Housing	U
	Hours[4]	2003	Labour Force Survey	A
Turkey	Schooling[1]	2003	Demographic and Health Survey (EPDC)	U
	Hours	2003	Household Labour Force Survey (ILO)	A
Uganda	Schooling	2002	Population and Housing Census (IPUMS)	A
	Hours	2006	National Household Survey	A
Ukraine	Schooling[1]	2007	Demographic and Health Survey (EPDC)	U
Uzbekistan	Schooling[1]	2006	Multiple Indicator Cluster Survey (EPDC)	U
Venezuela	Schooling	2001	Population and Housing Census (IPUMS)	A
	Hours	2001	Population and Housing Census (IPUMS)	A
Vietnam	Schooling	1999	Population Census (IPUMS)	A
	Hours	1999	Labour Force Survey (ILO)	A

Yemen	Schooling[1]	2006	Multiple Indicator Cluster Survey (EPDC)	U
Zambia	Schooling[1]	2007	Demographic and Health Survey (EPDC)	U
	Hours	2005	Labour Force Survey	A
Zimbabwe	Schooling[1]	2006	Demographic and Health Survey (EPDC)	U
	Hours[5]	2009	Labour Force Survey	A

Note: Hours worked and years of schooling data are for all economically active persons aged 15+ unless otherwise noted. IPUMS is the International Public-Use Microdata Series; EPDC is the Education Policy and Data Center; ILO is the International Labor Organization; LSMS are the Living Standards Measurement Surveys.

[1]Years of schooling imputed from educational attainment.
[2]Hours worked in main occupation.
[3]Agriculture status determined from occupation
[4]Computed from intervalled hours data.
[5]Sample consists of only employed persons.
[6]Sample consists of heads of households only.
[7]Sample includes economically inactive persons.
[8]Sample includes persons aged 5+.
[9]Sample includes persons aged 10+.

Table 2: Raw Agricultural Productivity Gaps

Measure	Weighted	Unweighted
5th Percentile	1.7	1.1
Median	3.7	3.0
Mean	4.0	3.6
95th Percentile	5.4	8.8
Number of Countries	112	112

Sample is developing countries, defined to be below the mean of the world income distribution. The weighted statistics weight each country by its population.

Table 3: Rural-Urban Education Quality Differences

Country	$\hat{\gamma}$
Argentina	0.87
Bolivia	0.95
Brazil	0.89
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
Venezuela	0.78
Vietnam	0.74
Average	0.87

The value $\hat{\gamma}$ is our estimate of the number of years of urban schooling equivalent to one year of rural schooling. The "average" is the simple unweighted average across countries.

Table 4: Non-Agriculture to Agriculture Averages by Region

Region	Years of Schooling	Human Capital	Hours Worked	Cost of Living
Africa	2.8	1.4	1.2	1.3
Asia	1.8	1.4	1.2	1.3
Americas	2.0	1.5	1.0	1.4
Europe	1.4	1.3	1.2	1.0
All Countries	2.0	1.4	1.2	1.3

Sample is developing countries, defined to be below the mean of the world income distribution. Averages are weighted by population.

Table 5: Adjusted Agricultural Productivity Gaps

Measure	Complete Data	All Countries
5th Percentile	1.0	0.8
Median	1.9	1.8
Mean	1.9	1.9
95th Percentile	2.9	2.6
Number of Countries	35	112

Sample is developing countries, defined to be below the mean of the world income distribution. “Complete data” means the set of countries with data on hours, human capital and the cost of living by sector. “All countries” means that when data is missing is it imputed as the mean ratio across all countries with data available. All averages are weighted by population.

Table 6: Agricultural Productivity Gaps from Micro and Macro Data

Country	Agriculture Share (%) of			APG	
	Employment	Value Added			
	Micro	Macro	Micro	Macro	Micro
Cote d'Ivoire (1988)	71.0	32.0	37.7	4.3	4.0
Guatemala (2000)	40.2	15.1	16.9	3.4	3.3
Pakistan (2001)	57.8	25.8	20.5	4.2	4.3
South Africa (1993)	11.8	4.3	8.2	1.6	1.5

"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.

Table 7: Agricultural Productivity Gaps: Alternative Measures from Micro Data

Country	VA/Worker	Income/Worker	Expenditure/Worker
Cote d'Ivoire (1988)	4.0	3.7	3.5
Guatemala (2000)	3.3	3.3	2.6
Pakistan (2000)	4.3	5.3	1.9
South Africa (1993)	1.5	1.6	1.7

All ratios are calculated using LSMS household survey data.

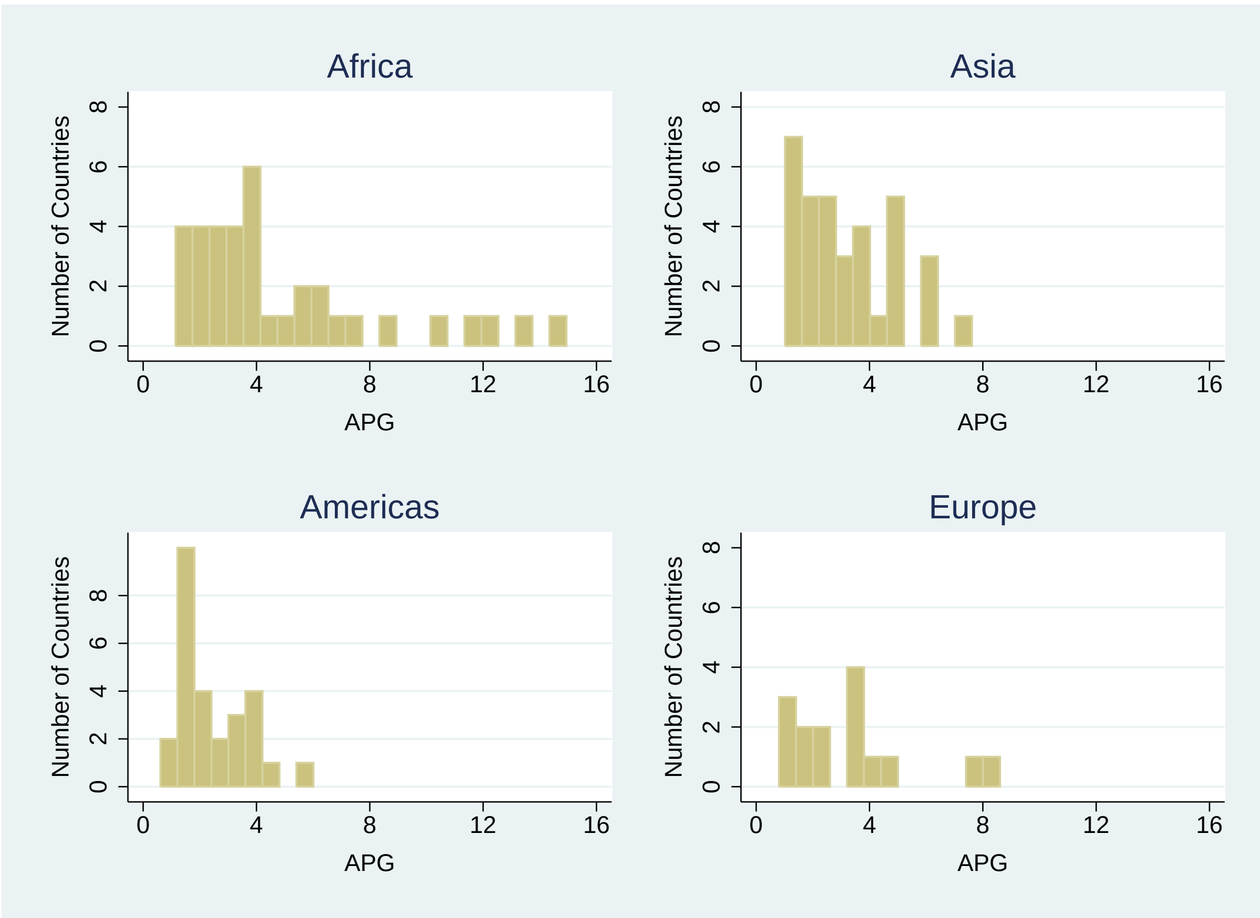

Figure 1: Distribution of APGs by Region

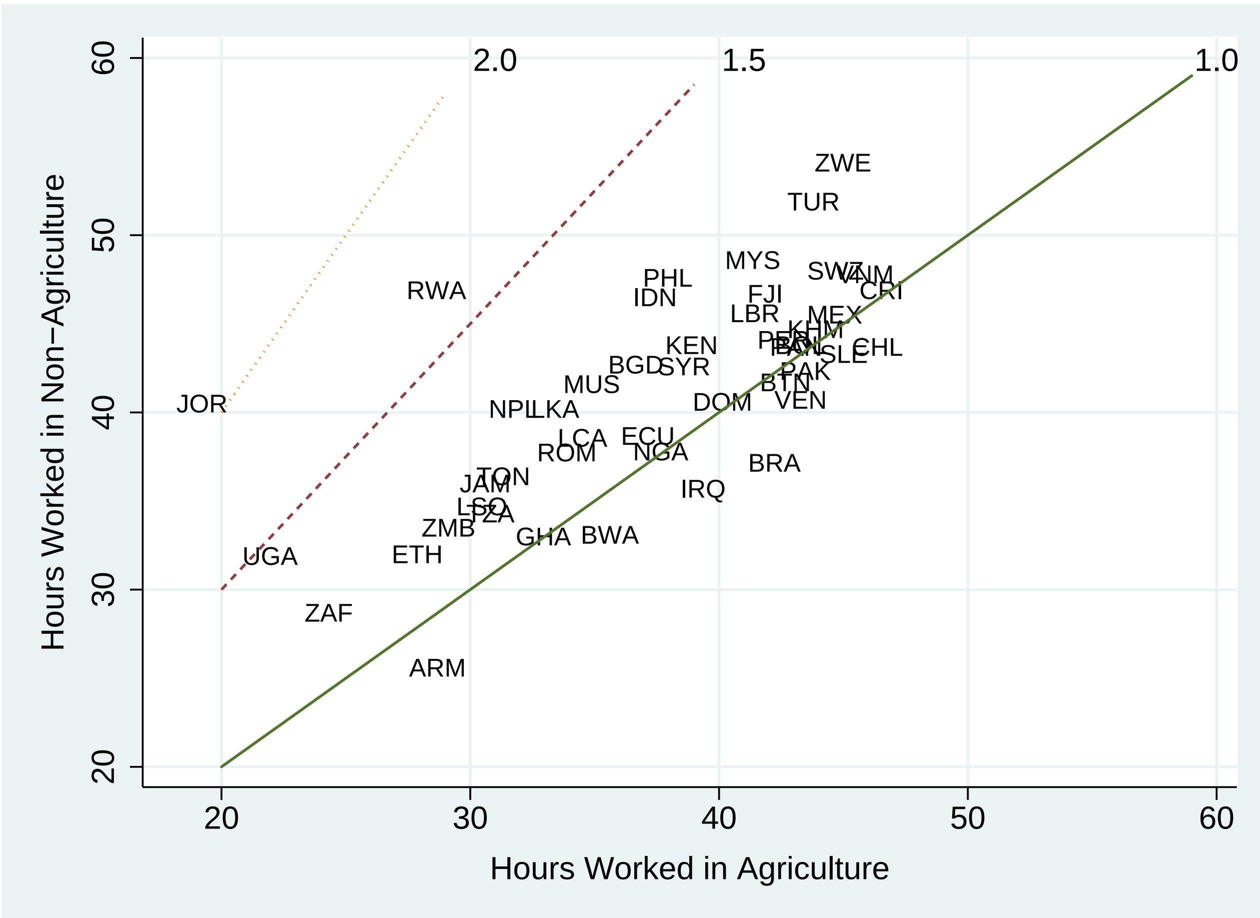

Figure 2: Hours Worked by Sector

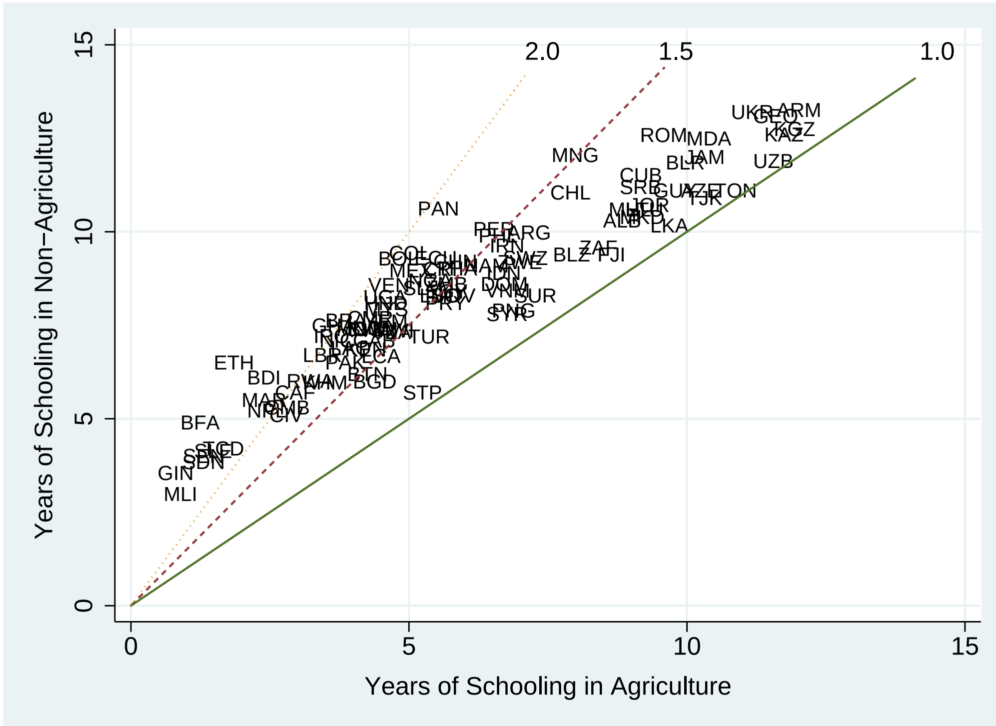

(a) Years of Schooling by Sector

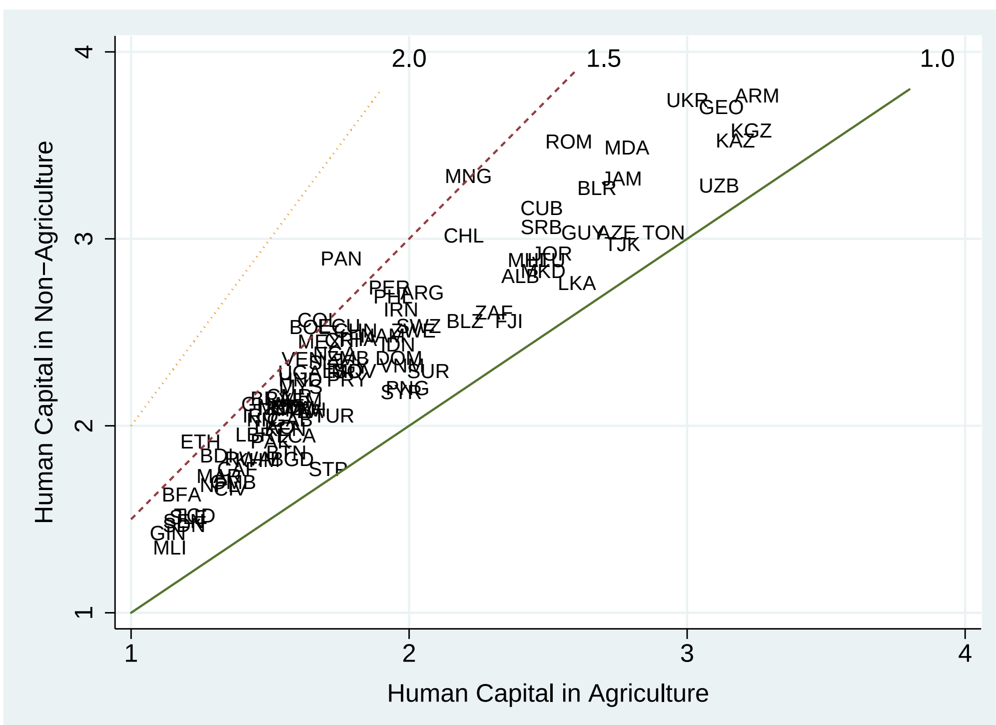

(b) Human Capital by Sector

Figure 3: Schooling and Human Capital by Sector

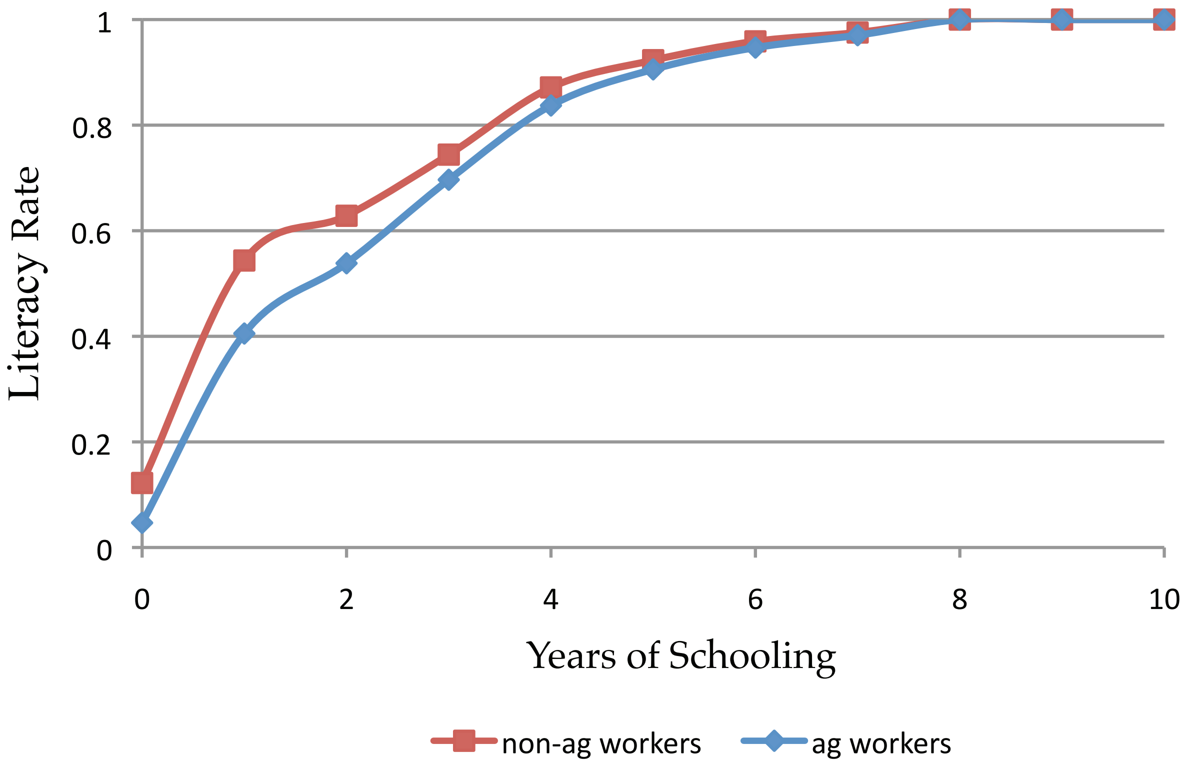

Figure 4: Literacy Rates by Years of Schooling, Uganda

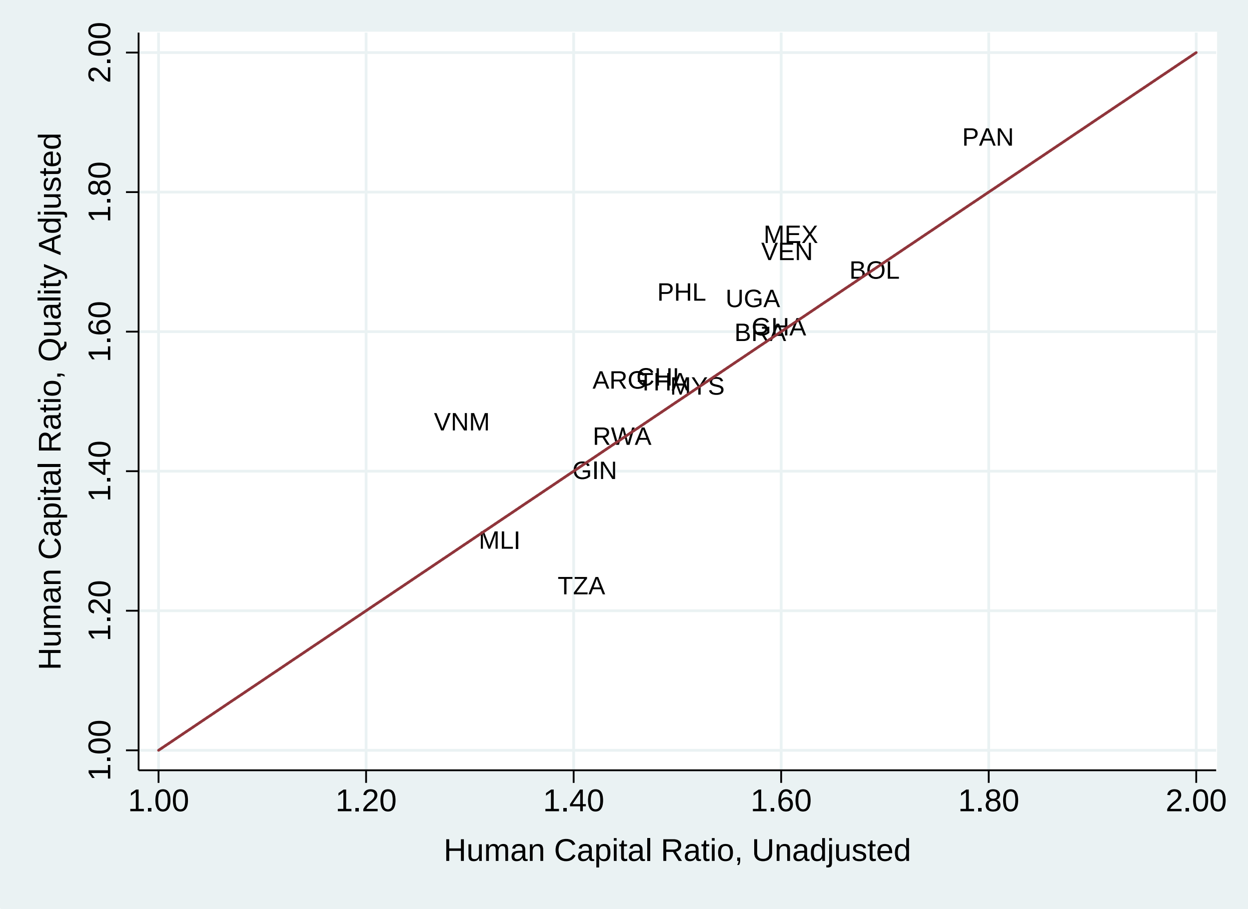

Figure 5: Human Capital Ratios, Adjusted for Quality and Unadjusted

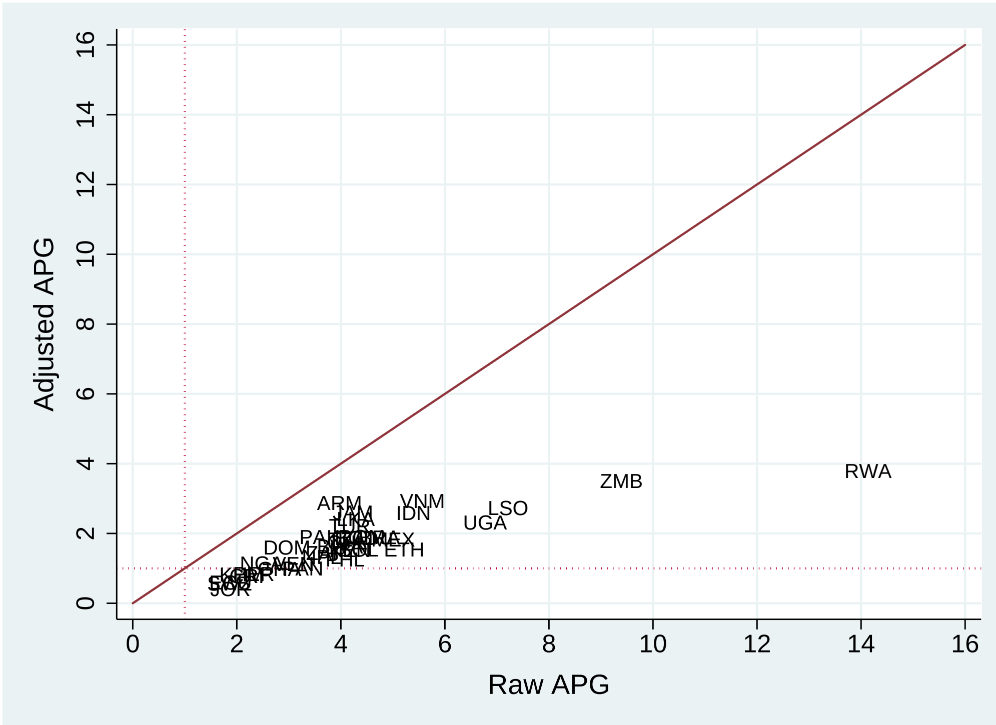

(a) All Countries with Complete Data

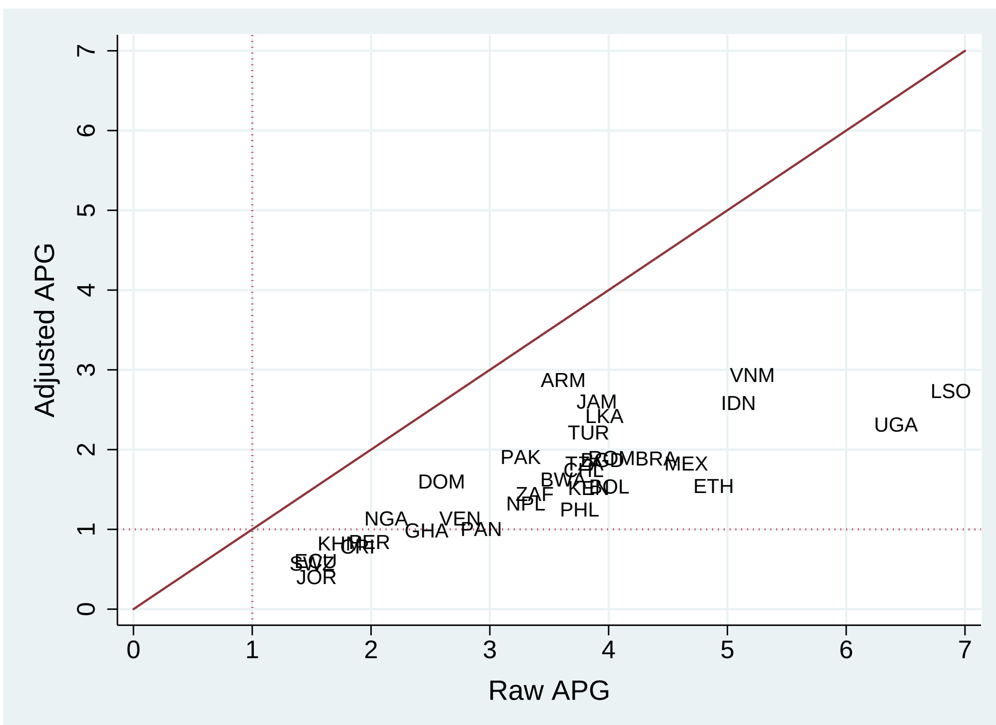

(b) Close up – Countries with Complete data and Raw APG ≤ 7

Figure 6: Raw and Adjusted Agricultural Productivity Gaps

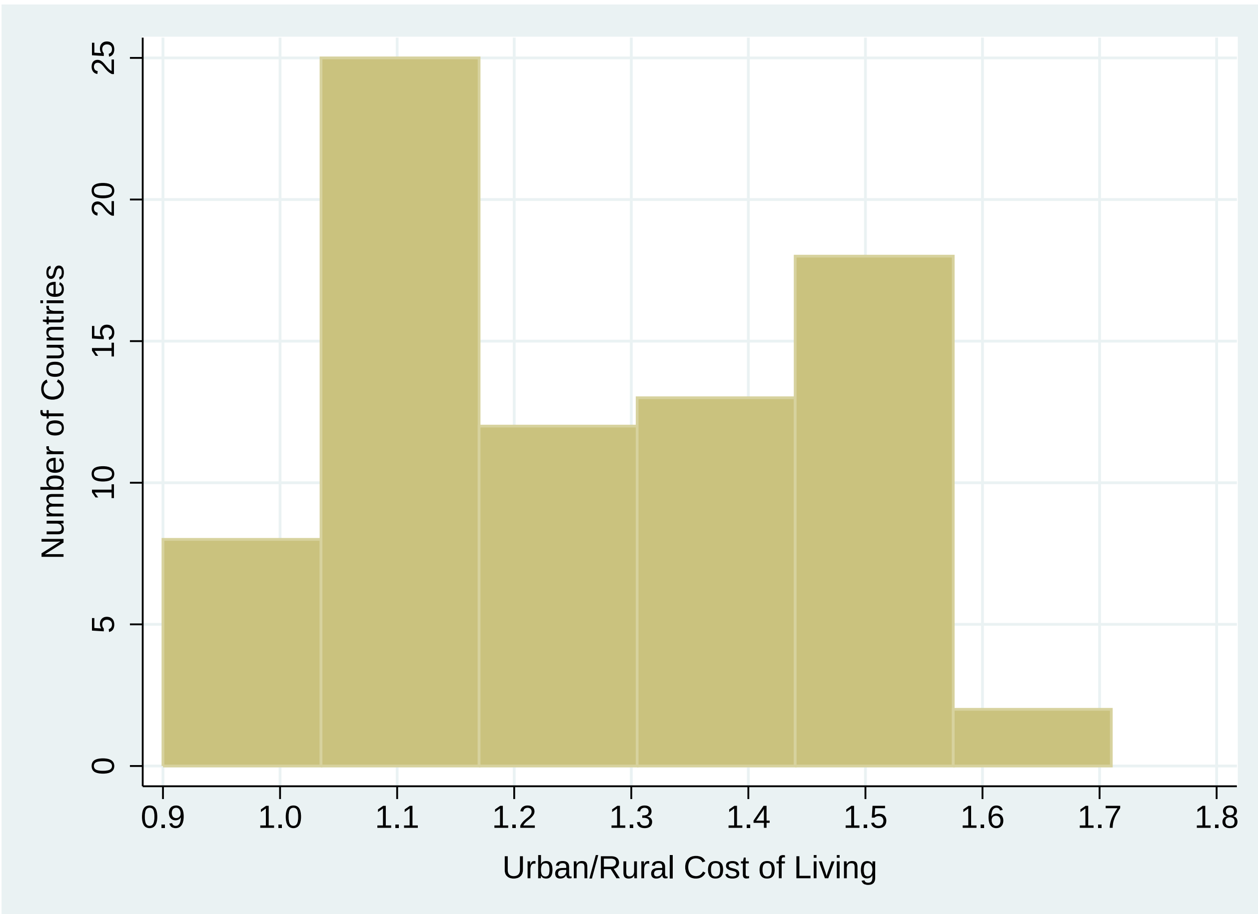

Figure 7: Cost-of-Living Ratio, Urban Areas/Rural Areas

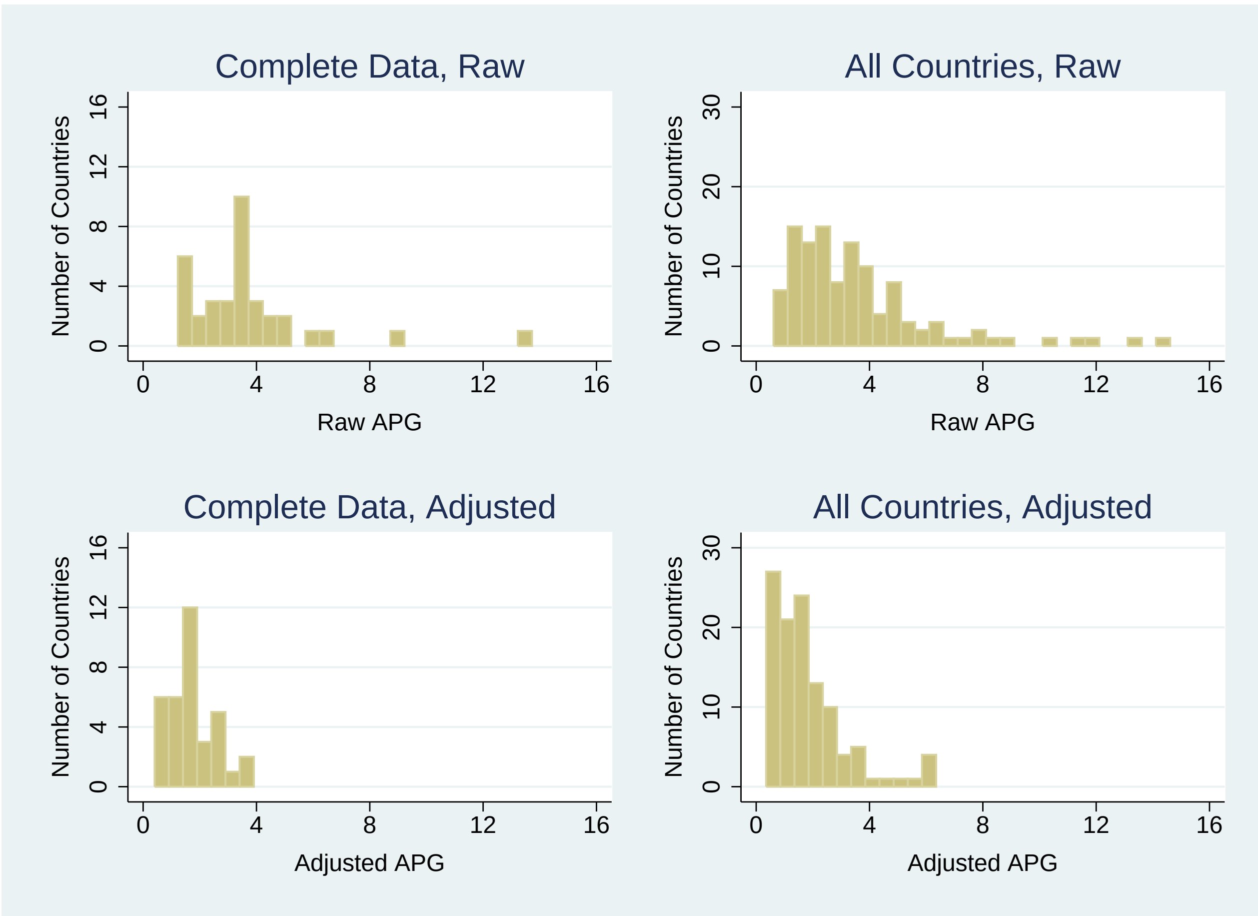

Figure 8: Distribution of Raw and Adjusted Agricultural Productivity Gaps

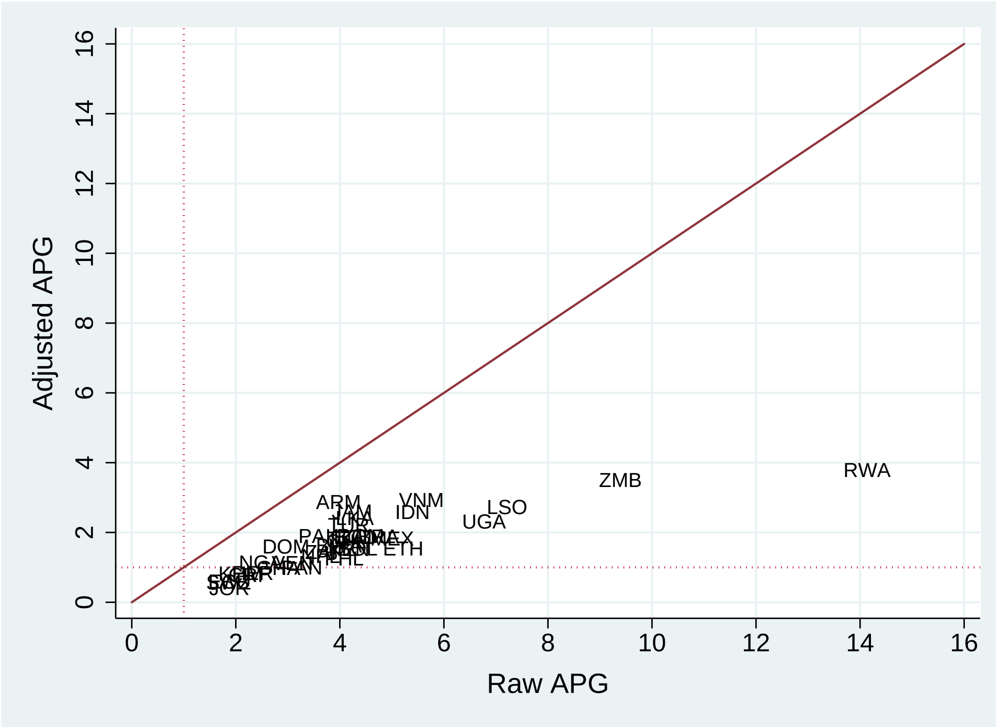

(a) All Countries with Complete Data

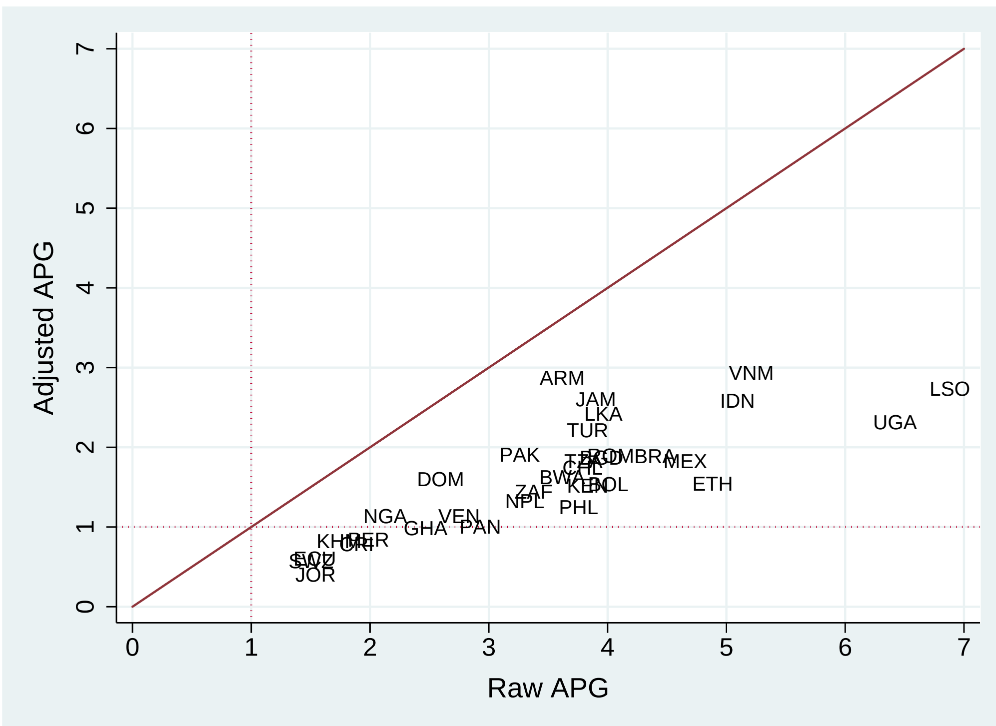

(b) Close up: Countries with Complete Data and Raw APG < 7

Figure 9: Raw and Adjusted Agricultural Productivity Gaps