

Agricultural Productivity Differences Across Countries

By DOUGLAS GOLLIN, DAVID LAGAKOS AND MICHAEL E. WAUGH*

Economists have long recognized that cross-country differences in aggregate labor productivity are enormous. Recently, Caselli (2005) and Restuccia, Yang and Zhu (2008), among others, have shown that these differences have a strong sectoral dimension. In particular, measured differences in *agricultural* labor productivity are far larger than those of the aggregate. Using data from Caselli (2005) and taking the ratio of labor productivity in the 90th and 10th percentile of countries, for example, one finds a ratio of 22 in the aggregate, and a ratio of 45 in agriculture.

Is productivity variation across countries really so much larger in agriculture than in the rest of the economy?¹ Or is it an artifact of measurement error in aggregate

data? In this paper, we provide new evidence on the existence of large agricultural productivity differences across countries. We do so by focusing on physical measures of productivity for the world's three staple grains—maize, rice, and wheat—which together account for roughly half of the calories consumed by the average individual.² Because productivity in these crops is relatively easy to measure, we are able to open up the “black-box” of national income and product accounting and focus on an easily measurable object. Moreover, because crop yields are observed at many levels, from individual production units to national aggregates, we are able to validate macro-level productivity statistics with micro-level estimates of productivity for these crops.

We draw several conclusions from our study. First, cross-country differences in the quantity of grain produced per worker are enormous, and are largely consistent with the magnitude of the differences in agricultural output per worker discussed above. Second, independent *micro* estimates of grain output per unit of land, and land per agricultural worker, largely agree with the aggregate estimates. This suggests that the large disparities in agricultural labor productivity are real and not merely an artifact of mis-measurement or of poor data quality. Finally, we conclude that the main source of disparities in agricultural productivity—in an accounting sense—is the very large difference across countries in land per worker. These differences in land per worker are so large that even if output per unit of land were

* Gollin: University of Oxford, Department of International Development, Queen Elizabeth House, 3 Mansfield Road, Oxford OX1 3TB, United Kingdom, Douglas.Gollin@qeh.ox.ac.uk. Lagakos: Department of Economics, 0508, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093 and NBER, lagakos@ucsd.edu. Waugh: New York University, Kaufman Management Center, 44 West Fourth St. 7-160. New York, NY 10012-1126, mwaugh@stern.nyu.edu. We thank Kevin Donovan for helpful comments. All potential errors are our own.

¹There are a number of reasons to be skeptical. Many studies have uncritically calculated sectoral productivity levels in agriculture by using nominal shares from national accounts data, which raises important concerns about cross-country productivity. The more careful studies by Caselli (2005) and Restuccia, Yang and Zhu (2008) used productivity data that were constructed by the UN Food and Agriculture Organization in 1985, when the quality of input and output data may have been low. In addition, the data are constructed using Rao (1993)'s international prices for agriculture, which are arguably better suited to constructing measures of gross output value than for quantifying sectoral value added. At a basic level, a compelling rationale for skepticism arises simply from the sheer magnitudes in agriculture which dwarf the (already large) productivity differences in the aggregate. Moreover, if the differences in agriculture are so large, then an implication is that the differences in non-agriculture are correspondingly small, which seems to run counter to a long tradition in growth theory that associates industrialization with aggregate productivity growth.

²FAOSTAT reports that these three foods together accounted for 1209 calories per person per day at the global level in 2009, relative to a total food supply of 2831 calories, for approximately 43 percent of total calorie consumption. For the category of “least developed countries,” these three food sources accounted for 1088 calories against a total supply of 2298 per person per day, or 47 percent of the total.

identical across countries, there would still be large differences in output per worker.³ Combined with the evidence from Gollin, Lagakos and Waugh (2014), who find large gaps in the value of the marginal product of labor in the agricultural and non-agricultural sectors of most countries, the findings of the current paper suggest that understanding why so many people in poor countries continue to work in what appear to be low-productivity agricultural activities is central to understanding why there are large *aggregate* cross-country differences in output per worker.

I. Cross Country Differences in Agricultural Output per Worker

A starting point for our analysis is the widely cited disparity in agricultural output per worker across countries. We begin by presenting the most recent available data on cross-country productivity differences, derived from aggregate Food and Agriculture Organization (FAO) data. These data provide a measure of gross output per worker in agriculture, where output is valued at international prices. These numbers do not correspond to the national accounts concept of value added, as they do not adjust for intermediate inputs used in agricultural production. The international prices used here are only for outputs; neither input quantities nor input prices are represented here. Nevertheless, these numbers are useful because they can be calculated for a very large set of countries. The results are striking. Table 1 shows the ratios of gross agricultural output per worker by country income group. Countries are ranked by GDP per capita at PPP for 2007. It is evident that there are vast disparities between the top and bottom countries and country groupings in the data. For instance, countries in the top ten percent of the world income distribution produce on average 50.1 times as much agricultural output per agricultural worker as countries in the bottom ten percent. Coun-

tries in the top quarter of the income distribution produce 29.9 times as much agricultural output per worker as countries in the bottom quarter. These differences are substantially larger than the differences in GDP per worker for the same sets of countries.⁴

Are these differences plausibly real? To address this question it is useful to decompose output per worker into two components: the output per unit of land and the land per unit of labor:

$$(1) \quad \frac{\text{Output}}{\text{Worker}} = \frac{\text{Output}}{\text{Land}} \times \frac{\text{Land}}{\text{Worker}}.$$

Both of these measures are reported separately in the FAO data, and they can be more easily verified against external micro data sources. The left panel of Table 2 reports the average yield of the three staple grain crops – maize, rice, and wheat – across countries, where crop yield is a measure of physical output of harvested grain per unit of land. In principal, crop yields are measured by national agricultural statistics offices using geographically representative sampling frames and sample crop cuts at harvest time; FAO compiles these data, which are widely viewed as meaningful and (broadly) accurate by agricultural scientists. In the next section, we will review micro evidence on crop yield, but for the moment, we simply note that there are meaningful differences across country groups in crop yield. Comparisons of crop yields across locations are necessarily complicated; different countries face different agroclimatic conditions for these three crops, and local market conditions may also vary. This means that any particular country or country group may rank differently in yield comparisons for the three crops. The United States, for example, is among the world’s leaders in yield of maize and rice, but its wheat yields are rel-

³This conclusion echoes the recent work of Adamopoulos and Restuccia (Forthcoming), which highlights the enormous differences in average farm size between rich and poor countries.

⁴We emphasize again however that these ratios reflect gross output per worker, without any adjustment for intermediates, and hence are not strictly comparable to GDP per worker. In this paper we leave aside the issue of intermediates mostly because intermediate input data are not available for most of the world’s poorest countries.

TABLE 1—CROSS-COUNTRY DIFFERENCES IN GROSS AGRICULTURAL OUTPUT PER WORKER

	Ratio
Richest Ten Percent over Poorest Ten Percent	50.1
Richest Quarter over Poorest Quarter	29.9

Note: Gross Agricultural Output is measured at PPP using international agricultural prices of the FAO. Agricultural Workers are measured as the total number of economically active persons involved in agriculture, and is taken from the FAO. Both are for 2007 or the closest available year. Percentiles refer to the distribution of countries by GDP per capita at PPP for 2007.

atively low, reflecting the fact that wheat is grown as a dryland crop and a winter crop in many parts of the country. Similarly, in sub-Saharan countries, maize and rice yields are low, but wheat yields almost as well as in the United States— but then, wheat is a very minor crop in most of the region, grown primarily in the cool highlands of Ethiopia and other East African countries and also in the temperate climate of South Africa; wheat area is nearly an order of magnitude smaller than maize area in sub-Saharan Africa.⁵

The crop yield data show that there are important differences in land productivity across countries, but these are relatively modest compared to the differences in gross output per worker. For maize, the top ten percent of countries produce 9.2 tons per hectare, while the bottom ten percent produce just 2.0 tons per hectare. Thus, the top ten percent of countries produce just under five times more maize per hectare than the bottom ten percent. Large differences in maize yields are present between the top and bottom quarter of countries by income as well. For rice and wheat, the richest ten percent of countries produce 8.1 and 4.9 tons per hectare, while the bottom ten percent produce 2.9 and 2.0 tons per hectare. This implies that the richest countries output per unit of land is more than twice as many tons of rice and wheat per

hectare as the poorest countries.

In an accounting sense, it must be the case that the remaining disparities are due to differences in land per agricultural worker. The right panel of Table 2 shows the distribution of this variable across countries. Measurement here is inevitably complicated, as land is not homogenous; apart from natural differences in soil fertility and slope, there are also variations in the extent of land improvements such as irrigation and drainage. Nevertheless, we abstract from these differences and use a measure of arable land per worker, while recognizing that this is an imperfect measure that fails to account for differences in land quality and the intensity of farming.⁶

The final column of Table 2 shows that there are very large differences in arable land per agricultural worker. The top ten percent of countries in the income distribution have on average 44.6 hectares per worker. The bottom ten percent of countries in contrast have just 1.4 hectares per worker. Putting these together, the richest ten percent of countries use 31.2 times as much land per worker as the bottom ten

⁵It is not uncommon to find that crop yields are very high in countries where the area harvested is quite small; rather than representing a violation of the laws of comparative advantage, this simply reflects the fact that these crops are often grown under highly atypical conditions. For instance, wheat yields in Namibia are as high as those obtained in Europe—but only 2000 hectares were harvested in 2012, less than a sixth of the wheat area of Luxembourg.

⁶We would prefer to use a measure of area harvested per worker for each of the three crops that we are considering— but no aggregate data sources report agricultural labor by crop, making it essentially impossible to carry out our decomposition at the crop level (i.e., to decompose output per worker into output per unit of land and land per unit of labor). We also note that our measure of arable land is based on a classification of country area; as such, we do not address the phenomenon of multiple cropping, in which the same plot of land may be harvested two or even three times per year, under the same crop or different crops. For instance, in some parts of Southeast Asia, farmers may harvest three crops of rice each year from the same plot of land; similarly, in parts of the United States, farmers may plant wheat and soybeans or wheat and maize on the same land.

TABLE 2—TONS PRODUCED PER HECTARE AND HECTARES PER WORKER

	Tons Produced per Hectare			Hectares per Worker
	Maize	Rice	Wheat	
Top Ten Percent	9.2	8.1	4.9	44.6
Top Quarter	8.2	6.8	4.8	23.9
Bottom Quarter	2.4	3.2	2.1	1.3
Bottom Ten Percent	2.0	2.9	2.0	1.4
Ratio of Top to Bottom Ten Percent	4.7	2.8	2.5	31.2
Ratio of Top to Bottom Quarter	3.4	2.1	2.2	18.0

Note: From FAO. Land is measured as hectares of arable land.

percent. Comparing the richest and poorest quarter of countries by income, we find a ratio of 18.0 in arable land per agricultural worker. Note that this is not so much a function of natural endowments, as many rich countries (e.g., Belgium, Netherlands or Japan) have relatively modest land area relative to population. Instead, the variation here reflects the fact that agriculture accounts for a much smaller fraction of the labor force in rich countries than in poor countries.

Because our measure of land per worker in agriculture does not coincide exactly with our measures of output per unit of land, we cannot carry out a full decomposition in an arithmetic sense. However, we interpret these results as demonstrating, at a somewhat trivial level, that the large differences across countries in gross agricultural output per worker reflect modest differences in land productivity combined with very large differences in land per worker. This is important in that it suggests that, from the perspective of growth economics, we may learn less from focusing on the *intensity* of agricultural production (e.g., the use of capital and intermediate inputs per unit of land) than from gaining a richer understanding of the reasons why so many people work in agriculture in poor countries. The two are of course not unrelated, as capital and intermediates can substitute for agricultural labor. But we can argue with some confidence that theories of agricultural output-per-worker differences across countries should have dif-

ferences in land per worker playing a far greater role than differences in output per unit of land.

II. Micro Evidence on Agricultural Output per Worker

One concern that arises with the aggregate data above is that these data may simply reflect poor statistical procedures and inaccurate reporting, as suggested by Jevins (ref). If this is true, then we cannot assign much weight to the macro numbers derived above. One advantage of our decomposition, however, is that both output per unit of land and land per worker are readily observed in micro data. The following paragraphs report our efforts to check the macro observations against independent micro measures of these key variables.

We begin by comparing FAO measures of national crop yield against independent micro measures of yield, taken from a variety of micro surveys and experiments, reported in a variety of sources, including some studies carried out by economists and others conducted by agricultural scientists.⁷ Each observation is of a particular crop (either maize, rice or wheat) in a particular country in a particular year. We compare these crop-country-year observations with the corresponding FAO observations and derive Figure 1. This figure shows a compelling correlation between the two data sources, with almost all of the observations

⁷Specifically, we draw on evidence from Livezey and Foreman (2004), Lobell, Cassman and Field (2009), **Need to fill out all the rest of these**.

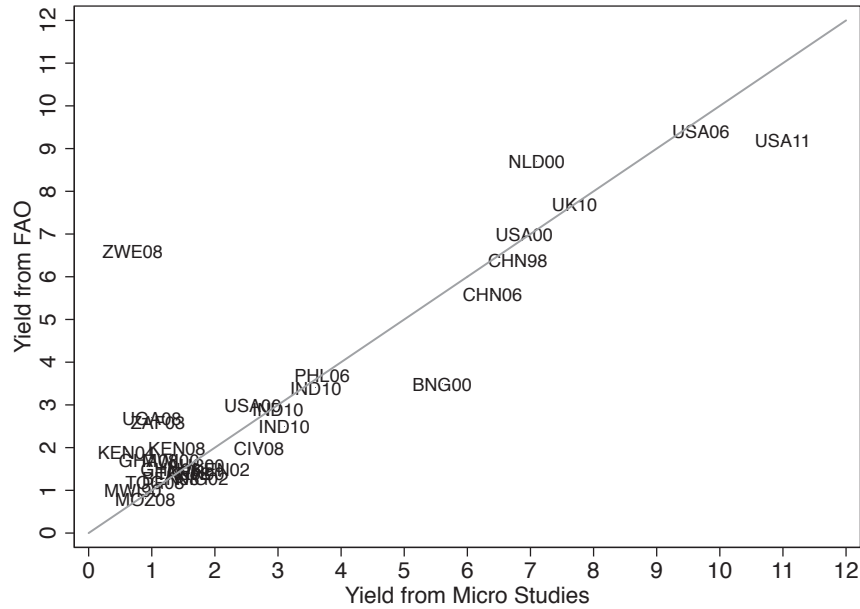


FIGURE 1. GRAIN YIELDS FROM AGGREGATE AND MICRO SOURCES.

Note: Micro sources come from studies in economics or the natural sciences and correspond to a single estimated yield for maize, rice or wheat in a single year. Aggregate data come from the FAO for the same crop and year as the micro source in question.

very close to the 45-degree line. The data also appear to support the FAO observation that rich countries attain higher yields than poor countries, with the magnitude of the differences broadly consistent between micro and macro data sources.⁸ We cannot claim that the selection of data points here is statistically representative or that the FAO yield data are therefore accurate for all purposes or questions. What we can say, however, is that we find essentially no disagreement between the FAO yield data and the many micro estimates of grain yields that we compiled.

Next, we turn to measures of land per worker. Here we compare the FAO data on arable land per agricultural worker, as described above, with data drawn from Living Standards Measurement Surveys (LSMS), which are widely regarded as one of the

highest quality sources of micro data available for developing countries. The LSMS data are the product of household surveys, which does mean that they are not necessarily representative of all farm establishments—in particular, they may miss some large commercial farms. But most evidence suggests that these surveys provide a reasonably accurate view of the rural economy in developing countries. The evidence is quite compelling: farm size is extremely small in most developing countries, with many or most families operating farms smaller than 5 hectares (Eastwood et al., 2004, p. 8). In many countries, land holdings are even smaller. For instance, Malawi's 2006-07 National Census of Agriculture and Livestock found over 2.5 million agricultural land holdings, with an average size of 0.96 ha per holding. Only 8 percent of holdings were larger than two hectares. In this setting, a small number of very large farms might alter the land labor ratio slightly but would still not move it as

⁸The FAO data are intended to be nationally representative. Few of our micro studies aim at nationally representative samples, so we would not expect the numbers to lie exactly on the 45-degree line even if the two data sources were perfectly consistent.

high as two hectares per holding.⁹

So far, we find that the LSMS studies show very small amounts of land per worker in developing countries, as in the FAO data. In Malawi in 2010, for example, the LSMS data show that the average rural household had 0.9 hectares, and 2.3 workers in the labor force. This implies that in Malawi, there are on average 0.4 hectares per worker in rural areas. Similar calculations in Ethiopia (2011), Tanzania (2010) and Nepal (2010) show 0.7, 0.3 and 0.3 hectares per worker in rural areas. Of course not all rural workers are in agricultural workers in these countries (though the vast majority certainly are), and it is unlikely that the farmland used in large commercial farms or publicly owned farmland are counted in these surveys. Still, with these important caveats aside, these surveys do point to very low amounts of land per worker in developing countries.

In developed countries, there is no serious debate about the amount of arable land per agricultural worker. Large censuses of agriculture lead to accurate estimates of the total land used in crop production, and employment surveys and population censuses provide good estimates of the size of the agricultural labor force. The FAO estimates of arable land and agricultural employment in the United States, for example, come from the USDA. Combining this lack of debate in rich countries with our findings from the LSMS data for developing countries, we conclude that aggregate and micro estimates of land per worker across countries tell essentially the same story.

III. Conclusion

A recent literature has claimed that cross-country differences in *agricultural* labor productivity are even larger than cross-country differences in *aggregate* productivity. In this paper we re-visit the data underlying this claim. We focus on the world's

three staple grains – maize, rice and wheat – for which direct measures of physical productivity are readily available. We conclude that productivity differences in grain are enormous according to both macro and micro evidence. This serves as evidence against the hypothesis that the vast cross-country productivity differences in agriculture are largely an artifact of measurement error.

REFERENCES

- Adamopoulos, Tasso, and Diego Restuccia.** Forthcoming. “The Size Distribution of Farms and International Productivity Differences.” *American Economic Review*.
- Caselli, Francesco.** 2005. “Accounting for Cross-Country Income Differences.” In *Handbook of Economic Growth*, ed. P. Aghion and S. Durlauf.
- Gollin, Douglas, David Lagakos, and Michael E. Waugh.** 2014. “The Agricultural Productivity Gap.” *Quarterly Journal of Economics*, 129(2).
- Livezey, Janet, and Linda Foreman.** 2004. “Characteristics and Production Costs of U.S. Rice Farms.” *USDA Statistical Bulletin Number 974-7*.
- Lobell, David B., Kenneth G. Cassman, and Christopher B. Field.** 2009. “Crop Yield Gaps: Their Importance, Magnitudes, and Causes.” *Annual Review of Environment and Resources*, 34.
- Rao, P.** 1993. “Intercountry comparisons of agricultural output and productivity.” *Etude FAO: Developpement Economique et Social (FAO)*.
- Restuccia, Diego, Dennis Tao Yang, and Xiaodong Zhu.** 2008. “Agriculture and Aggregate Productivity: A Quantitative Cross-Country Analysis.” *Journal of Monetary Economics*, 55: 234–250.

⁹Note also that a holding is typically operated by a family using predominantly family labor, often with more than one effective worker per holding. This means that the estimates of average holding size most likely overstate the land-labor ratio.