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, differences in *agricultural* labor productivity are far larger than those of the aggregate. Caselli (2005), for example, reports that the ratio of labor productivity in the 90th and 10th percentile of countries is a factor of 22 in the aggregate, and a factor of 45 in agriculture. Because developing countries have most of their workers in agriculture, understanding why productivity differences in agriculture are so large is key to understanding world income inequality.

There are a number of reasons one may be skeptical, however, about the agricultural productivity data underlying these conclusions. One may worry that agricultural output data may be badly measured, particularly in the world's poorest countries, where statistical agencies often have limited resources (Jerven, 2013). The international prices used to aggregate agricultural goods to the country level may be better suited to richer countries than poorer countries, since international prices are quantity-weighted. More generally, skepticism may arise simply from the sheer magnitudes in agriculture, which dwarf the (already large) productivity differences in the aggregate.

* 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, Diego Restuccia, Nora Trapani and Chris Udry for helpful comments, and we thank Jonathan Greenland and Glenn Farley for excellent research assistance. All potential errors are our own.

In this paper we re-examine the cross-country agricultural productivity data using new evidence from disaggregate sources. We focus on physical measures of labor 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. Because productivity in these crops is easy to measure, we avoid the "black-box" nature of output data from aggregate sources. Moreover, because crop yields are observed at many levels, from individual production units to national aggregates, we are able to cross-check macro-level productivity statistics with micro-level estimates of productivity for these crops.

We find that cross-country differences in the quantity of grain produced per worker are enormous, and at least as large as those of the agricultural sector as a whole. Moreover, we show that independent micro-level estimates of grain yields correspond very closely to their counterparts in aggregate data. We conclude that the large disparities in agricultural labor productivity are real, at least for staple grains, and are not merely an artifact of mis-measurement or poor data quality.

I. Cross Country Differences in Agricultural Output per Worker

We begin by presenting the most recent available data on cross-country agricultural productivity differences, derived from aggregate data from the Food and Agriculture Organization (FAO) for 2007. 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. We draw on data for the 80 countries for which data on gross agricultural output per worker and GDP per capita are available, and in which at least one hundred thousand people are employed in agriculture.

The results are striking. Countries in the top ten percent of the world income distribution produce 50.1 times as much agricultural output per agricultural worker as countries in the bottom ten percent. Countries in the top quarter of the income distribution produce 29.9 times as much agricultural output per worker as countries in the bottom quarter. The United States in particular produces more than one hundred times as much agricultural output per worker as the sub-Saharan countries in the data (Ethiopia, Malawi, Madagascar, Zambia, Tanzania, Uganda, Ghana, Botswana, and Nigeria). These differences are substantially larger than the differences in GDP per worker for the same set of countries, and they are of similar magnitude to the agricultural productivity differences reported in Caselli (2005) and Restuccia, Yang and Zhu (2008).[1]

Are these differences plausible? Or are they, perhaps, an artifact of poor measurement? To address this question we look at physical measures of productivity for the world's three staple grains: maize, rice, and wheat. Focusing on physical measures of productivity in these grains is useful for several reasons. First, these three grains account for roughly half of the calories consumed by the average individual.[2] Thus, they constitute a substantial fraction of all agricultural output in their own right. Second, measuring productivity for grains does not require aggregation with international prices, thus avoiding one source of possible bias. Third, yields are a measure of production, not sales, and thus are an estimate of productivity whether the good is sold in the market or consumed at home. Finally, productivity in these grains is carefully measured by natural scientists. Thus, we can use independent estimates to cross-check the productivity numbers recorded in aggregate data.

For what follows, it is useful to decompose output per worker into two components: output per unit of land and land per unit of labor:

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

The value of this decomposition is that both of these measures are reported separately in the FAO data, and are independently estimated in micro-level studies by natural scientists or economists. In what follows, we cross-check the FAO data on agricultural output per worker by studying estimates of output per unit of land, and then land per worker.

Table 1 reports the average yield of the three staple grain crops across countries, measured as tons of output per hectare of land, by country income group.[3] Table 1 shows that there are important differences in yields across countries. For maize, countries at the top of the income distribution produced about nine tons per hectare, while the bottom ten percent produced just two. 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 produce more than twice as much grain per hectare as the poorest countries.

The right panel of Table 1 shows the

[1] We emphasize again 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 because intermediate input data are not available for most of the world's poorest countries.

[2] FAOSTAT reports that these three foods together account for approximately 43 percent of total calorie consumption. For the category of *least developed countries* these three food sources account for 47 percent of the total.

[3] Comparisons of crop yields across locations are necessarily complicated, as countries face different agroclimatic conditions and local market conditions. The United States, for example, is among the world's leaders in yield of maize and rice, but its wheat yields are relatively 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 are higher. But then, wheat is a very minor crop in most of the region, with nearly an order of magnitude fewer acres devoted to it than maize.

Table 1—Tons Produced per Hectare and Hectares per Agricultural Worker

	Tons Produced per Hectare			Hectares per Worker
	Maize	Rice	Wheat	
Top Ten Percent	9.2	8.1	4.9	44.6
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

Note: From FAO. Land is measured as hectares of arable land. Workers are agricultural workers are measured as the total number of economically active persons involved in agriculture.

distribution of arable land per agricultural worker across countries.[4] Table 1 reports very large differences in land per worker. Countries at the top of the income distribution have on average 44.6 hectares per worker, while the countries at the bottom of the income distribution 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 percent. 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.

A striking result arises when combining the left panel of Table 1 with the right panel. Combining an output-per-hectare difference of two (we will be conservative here) with a hectare-per-worker difference of 30 implies a yield per worker difference of 60 between the richest and poorest ten percent of countries. That is, cross-country differences in output per worker of these staple grains are of the same order of magnitude (or larger) than the agriculture-sector gross output-per-worker differences reported above.

[4]Measurement here is inevitably complicated, as land is not homogenous, and the data (by construction) 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.

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 Jerven (2013). One advantage of our decomposition, however, is that both output per unit of land and land per worker are readily observed in micro data. This section reports our efforts to check the macro observations against independent micro measures.

We begin by comparing FAO measures of national crop yield against independent micro measures of crop yield. These measures are from a variety of micro surveys and studies, reported in a variety of sources, including some studies carried out by economists and others conducted by agricultural scientists.[5] 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 plot them in Figure 1. This figure shows a compelling correlation between the two data sources, with almost all of the observations very close to the 45-degree line. The data also appear to support the FAO observation that

[5]Specifically, we draw on estimates from Burke et al. (2011), Duffy (2014), Erenstein et al. (2003), Fermont and Benson (2011), Gouse, Piesse and Thirtle (2006), Huang et al. (2000), Husain, Hossain and Janaiah (2001), Lobell, Cassman and Field (2009), Livezey and Foreman (2004), Quinones and Diao (2011), Ragasa et al. (2013), Smale (1993), Suri (2011), Tittonell and Giller (2013), Tittonell et al. (2008), Indian Ministry of Agriculture (2007), Department for Environment, Food and Rural Affairs of the U.K. (2012) and Williams (1998).

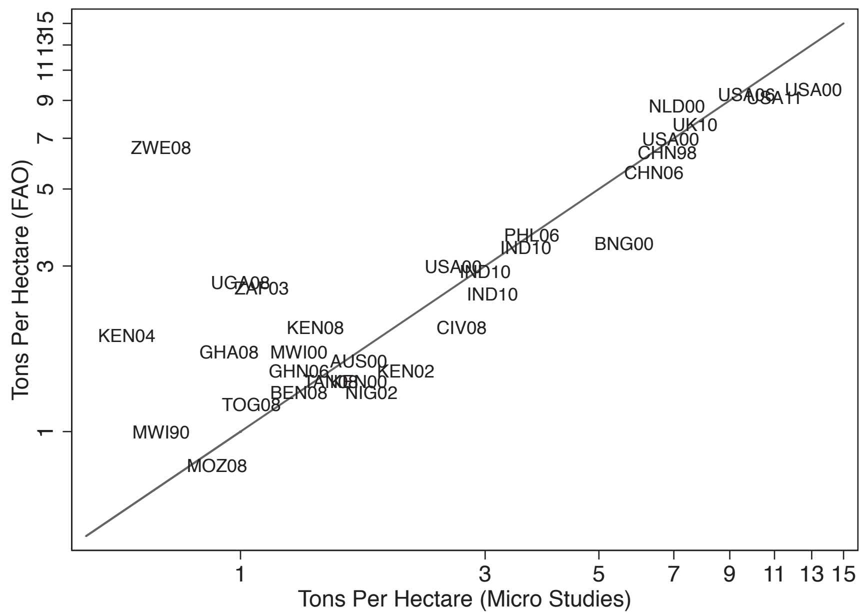

Figure 1. Grain Yields from Aggregate and Micro Sources.

Note: Micro sources correspond to a single micro-level estimated yield for maize, rice or wheat. Aggregate data come from the FAO for the same crop and year as the corresponding micro estimate.

rich countries attain higher yields than poor countries, with the magnitude of the differences broadly consistent between micro and macro data sources.[6] 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 these micro estimates of grain yields that we compiled.

What about land per worker? In richer countries, there is less concern about the quality of data on land and total agricultural employment, due to the large and regular censuses of agriculture and population. For the United States, for example, we found effectively the same land per worker when comparing census data from the US Department of Agriculture (USDA) in the year 2007 to the FAO data reported for the US in 2007. The USDA, which uses satellite imagery (at least in part) to measure crop land (Boryan et al., 2011) show 48 hectares per worker versus 53 hectares per worker in the FAO data.

As a check of the FAO land-per-worker data for poor countries, we draw on the Living Standards Measurement Surveys (LSMS) which are regarded as one of the highest quality sources of micro data available for developing countries. Furthermore, the LSMS are independent surveys of households that are not used in the construction of national aggregate statistics. 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), Guatemala (2003), Nigeria (2010) and Tanzania (2010) show 0.3, 1.9, 0.5, and 0.3 hectares per worker in rural areas. In short, these LSMS

[6]The FAO data are intended to be nationally representative. Few of our micro studies aim at nationally representative samples though, which may account for some of the discrepancies between the two sets of numbers.

surveys show that land per worker is indeed in the ballpark of one hectare per worker in developing countries.[7]

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 data. This serves as evidence that the large productivity differences in the agricultural sector are real, and not an artifact of poor measurement.

So why is agricultural output per worker so much lower in the developing world than in rich countries?[8] One hypothesis is that policies that distort farm size lead to a misallocation of farmland to farm operators (Adamopoulos and Restuccia, Forthcoming). Another theory is that farm operators in poor countries avoid using productivity-enhancing intermediates, such as fertilizers, because doing so increases their consumption risk (Donovan, 2013). A third theory is that the agriculture sector in developing countries tends to employ the lowest-ability workers (Lagakos and Waugh, 2013). Surely these are just some of the economic forces at work. There is much room for future research on this important issue.

[7] There are several caveats in order here. First, not all rural workers are agricultural workers in these countries. So the above calculations likely underestimate the amount of land per agricultural worker. Second, LSMS data are the product of household surveys, which means that they may miss large commercial farms. Still, the total acreage devoted to large commercial farms tends to be a small fraction of total farm land in most developing countries.

[8] A distinct but related question is why, within the typical developing country, the value of output produced per worker in agriculture is substantially lower than the value of output produced by non-agricultural workers, on average. This is the focus of our related work (Gollin, Lagakos and Waugh, 2014), and of Herrendorf and Schoellman (2013).

REFERENCES

Adamopoulos, Tasso, and Diego Restuccia. Forthcoming. "The Size Distribution of Farms and International Productivity Differences." *American Economic Review.*

Boryan, Claire, Zhengwei Yang, Rick Mueller, and Mike Craig. 2011. "Monitoring US agriculture: the US Department of Agriculture, National Agricultural Statistics Service, Cropland Data Layer Program." *Geocarto International*, 1: 1–18.

Burke, W., M. Hichaambwa, D. Banda, and T. Jayne. 2011. "The Cost of Maize Production by Smallholder Farmers in Zambia." Working Paper No. 50, Food Security Research Project.

Caselli, Francesco. 2005. "Accounting for Cross-Country Income Differences." In *Handbook of Economic Growth.* , ed. P. Aghion and S. Durlauf.

Department for Environment, Food and Rural Affairs of the U.K. 2012. "Agriculture in the United Kingdom, 2012." https://www.gov.uk/government/publications/agriculture-in-the-united-kingdom-2012.

Donovan, Kevin. 2013. "Agricultural Risk, Intermediate Inputs, and Cross-Country Productivity Differences." Unpublished Manuscript, University of Notre Dame.

Duffy, M. 2014. "Estimated Costs of Crop Production in Iowa." http://www.extension.iastate.edu/agdm/crops/pdf/a1-20.pdf.

Erenstein, Olaf, Frederic Lancon, S.O. Akande, S.O. Titilola, G. Akpokodje, and O.O. Ogundele. 2003. "Rice Production Systems in Nigeria: A Survey." West Africa Rice Development Association.

Fermont, A., and T. Benson. 2011. "Estimating Yield of Food Crops Grown by Smallholder Farmers: A Review in the

Uganda Context." IFPRI Discussion Paper 01097.

Gollin, Douglas, David Lagakos, and Michael E. Waugh. 2014. "The Agricultural Productivity Gap." *Quarterly Journal of Economics*, 129(2).

Gouse, Marnus, Jenifer Piesse, and Colin Thirtle. 2006. "Monsanto's Adventures in Zulu Land: Output and Labour Effects of GM Maize and Minimum Tillage." International Association of Agricultural Economists Conference.

Herrendorf, Berthold, and Todd Schoellman. 2013. "Why is Measured Productivity So Low in Agriculture?" Unpublished Manuscript, Arizona State University.

Huang, J., F. Qiao, L. Zhang, and S. Rozelle. 2000. "Farm Pesticide, Rice Production, and Human Health." International Development Research Center.

Husain, A.M. Muazzam, Mahabub Hossain, and Aidas Janaiah. 2001. "Hybrid Rice Adoption in Bangladesh: A Socioeconomic Assessment of Farmers' Experiences." BRAC Research Monograph Series No. 18.

Indian Ministry of Agriculture. 2007. "Cost of Cultivation of Principal Crops in India." Ministry of Agriculture of Government of India, New Delhi.

Jerven, Morten. 2013. *Poor Numbers: How We Are Misled by African Development Statistics and What to Do about It.* Cornell University Press.

Lagakos, David, and Michael E. Waugh. 2013. "Selection, Agriculture and Cross-Country Productivity Differences." *American Economic Review*, 103(2): 948–80.

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.

Quinones, E.J., and X Diao. 2011. "Assessing Crop Production and Input Use Patterns in Ghana–What Can We Learn from the Ghana Living Standards Survey?" Ghana Strategy Support Program Working Paper No. 24, IFPRI.

Ragasa, C., A. Dankyi, P. Acheampong, P. Wiredo, A. Chapoto, M. Asamoah, and R. Tripp. 2013. "Patterns of Adoption of Improved Rice Technologies in Ghana." IFPRI Working Paper 35.

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.

Smale, Melinda. 1993. "Maize is Life: Maize Reesarch and Smallholder Production in Malawi." USAID Research Report.

Suri, Tavneet. 2011. "Comparative Advantage in Technology Adoption." *Econometrica*, 79(1): 159–209.

Tittonell, Pablo, and Ken E. Giller. 2013. "When Yield Gaps are Poverty Traps: The Paradigm of Ecological Intensification in African Smallholder Agriculture." *Field Crops Research*, 143: 76–90.

Tittonell, Pablo, B. Vanlauwe, M. Corbeels, and Ken E. Giller. 2008. "Yield Gaps, Nutrient Use Efficiencies and Response to Fertilisers by Maize across Heterogeneous Smallholder Farms in Western Kenya." *Plant and Soil*, 313: 19–37.

Williams, J. 1998. "Sample Costs to Produce Rice: Sacramento Valley, Rice Only Rotation." University of California Cooperative Extension.