

Fisher-Post, Matthew

Working Paper

Inequality and agricultural structural change: Evidence from macro and microdata, 1950-present

WIDER Working Paper, No. 2025/5

Provided in Cooperation with:

United Nations University (UNU), World Institute for Development Economics Research (WIDER)

Suggested Citation: Fisher-Post, Matthew (2025) : Inequality and agricultural structural change: Evidence from macro and microdata, 1950-present, WIDER Working Paper, No. 2025/5, ISBN 978-92-9267-562-2, The United Nations University World Institute for Development Economics Research (UNU-WIDER), Helsinki, <https://doi.org/10.35188/UNU-WIDER/2025/562-2>

This Version is available at:

<https://hdl.handle.net/10419/315099>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

WIDER Working Paper 2025/5

Inequality and agricultural structural change

Evidence from macro and microdata, 1950–present

Matthew Fisher-Post *

February 2025

Abstract: Since 1950, agricultural productivity has been increasing even as labourers leave agriculture. However, while average productivity of the sector has been converging, within-sector inequality has been increasing. Agricultural income inequality is still less than overall income inequality, but it measures significantly higher when we use higher-quality and more comprehensive survey data. This means not only to observe the entirety of household farm income, but also to measure the magnitude of capital income and corporate profits in the sector. Given the likely increase in agricultural inequality during the process of structural change, I show also the extent to which social protection programmes are both insufficient and poorly targeted for rural populations.

Key words: agricultural structural change, food systems transformation, inequality

JEL classification: D3, N5, O1, Q1

Acknowledgements: For helpful discussions and comments, the author would like to thank Ammad Bahalim, François Bourguignon, Derek Byerlee, Ruben Echeverría, Doug Gollin, El Iza Mohamedou, Vasco Molini, Jamie Morrison, Fabrizio Moscatelli, Thomas Piketty, Dani Rodrik, Akiko Suwa-Eisenmann, and Alberto Zezza, as well as organizers and participants at the UNU-WIDER ‘Kuznets at 70’ symposium and the ‘Reducing Inequalities for Food Security and Nutrition’ event of the United Nations Committee on World Food Security. I acknowledge financial support from ERC Grant #856455, EUR Grant #ANR-17-EURE-0001, the American Foundation for the Paris School of Economics, and the New Venture Fund via Bill & Melinda Gates Foundation (Grant #016354).

Note: tables and figures at the end of the paper, after the reference list

* Paris School of Economics, Paris, France, and United Nations Food and Agriculture Organization, Rome, Italy; mfp@psemail.eu

This study is published within the UNU-WIDER project [Inequality and structural transformation – Kuznets at 70](#), which is part of the [Reducing inequalities across and within countries](#) research area.

Copyright © UNU-WIDER 2025

UNU-WIDER employs a fair use policy for reasonable reproduction of UNU-WIDER copyrighted content—such as the reproduction of a table or a figure, and/or text not exceeding 400 words—with due acknowledgement of the original source, without requiring explicit permission from the copyright holder.

Information and requests: publications@wider.unu.edu

ISSN 1798-7237 ISBN 978-92-9267-562-2

<https://doi.org/10.35188/UNU-WIDER/2025/562-2>

Typescript prepared by Siméon Rapin.

United Nations University World Institute for Development Economics Research provides economic analysis and policy advice with the aim of promoting sustainable and equitable development. The Institute began operations in 1985 in Helsinki, Finland, as the first research and training centre of the United Nations University. Today it is a unique blend of think tank, research institute, and UN agency—providing a range of services from policy advice to governments as well as freely available original research.

The Institute is funded through income from an endowment fund with additional contributions to its work programme from Finland and Sweden, as well as earmarked contributions for specific projects from a variety of donors.

Katajanokanlaituri 6 B, 00160 Helsinki, Finland

The views expressed in this paper are those of the author(s), and do not necessarily reflect the views of the Institute or the United Nations University, nor the programme/project donors.

1 Introduction

Economic growth usually implies structural change, the movement of labour across sectors of employment.¹ Perhaps the most striking structural change in modern history has been the movement of labour from primarily agricultural employment, to primarily non-agricultural employment. (And similarly, from rural areas to urban.) However, little is known about the relationship between agricultural structural change and economic inequality—another hallmark of modern economic growth.

Simon Kuznets (1955) argued that inequality initially increases with development (as structural change)—as labourers move from the lower-productivity and lower-inequality (traditional, agricultural, rural) sector to the higher-productivity and higher-inequality (modern, manufacturing/services, urban) sector—before decreasing as the majority of the labour force reaches the modern sector. It is worth noting, though, that Kuznets did not necessarily foresee any changes in the mean income of the two sectors; if anything, he predicted a widening of the between-sector average productivity gap (Baymul and Sen 2020; Anand and Kanbur 1993). The dynamic trajectories of within-sector inequality, during structural change, were difficult to predict.

W. Arthur Lewis (1954) argued that the agricultural sector in developing countries is characterized by an ‘unlimited’ supply of labour working at close to subsistence wages. Growth in the non-agricultural sector proceeds by drawing agricultural labourers at close to (only slightly higher than) the same fixed rate of wages, and reinvesting the profits as capital accumulates. Again, the agricultural sector remains somewhat stagnant (and low in inequality) until its prevailing wage rate would equal that of the modern sector (Gollin 2014). The logic of structural change in this framework is compelling, but its view of within-agriculture changes is perhaps similarly limited.

For all of their elegance, these early models of structural change leave important questions on the dynamics of productivity growth and income distribution within the agricultural sector.

This paper makes several contributions. First, I show that agricultural productivity (relative to productivity outside the sector) has been increasing, even as labourers leave agriculture.²

¹ Canonically, labour (and capital) moves from lower-productivity to higher-productivity sectors.

² Whether this is because low-productivity workers leave, or staying workers increase their productivity, is a more difficult question.

Second, across agriculture and the rest of the economy, I document that between-sector inequality has been declining since 1980 and perhaps earlier, but within-sector inequality has been increasing. It is likely that both the non-agricultural sector and the agricultural sector have been increasing in inequality.

The first part of this narrative—that average productivity levels are converging across sectors and between-sector inequality is narrowing—could fit into the Lewis (1954) framework. The second part—that inequality is growing within the agricultural sector—matches the Kuznets (1955) intuition.

In any case, I show that agricultural structural change—the movement of labour from agriculture to other sectors of the economy—explains very little of the overall growth in inequality that has prevailed in recent decades.³ Instead, within-sector inequality in agriculture likely mirrors that in the non-agricultural sectors, with one caveat.

Much of agriculture remains a household business. I document that the corporate sector in agriculture has not been increasing. Then again, to the extent that agriculture is a corporate endeavor (or, one may infer, when it becomes so), its income distribution widens. Evidence in cross-section shows a compelling relationship between overall inequality and the size of the corporate sector in agriculture. I also show that the size of the corporate sector in agriculture is related to the extent of inequality in agriculture.

Given that agricultural structural transformation likely implies an increase in inequality in the sector, I discuss the adequacy and targeting of social protection among rural populations.

Unfortunately, to assess global changes in agricultural income distributions over time, demands more than existing data sources have captured. In fact, household surveys are notoriously poor at capturing capital incomes, and the income of those at the top of the distribution. The agriculture sector is perhaps an extreme case of this. Very few agricultural surveys capture any information about the corporate sector. Distinct income concepts, missing data on income, and missing individuals are at the heart of the empirical measurement challenge.

2 Data

Data for this study draws from several sources.

³ Note that between-country inequality has actually converged somewhat, so the overall world income distribution has largely compressed (except at the top). At the same time, within-country inequality has increased almost everywhere. See Bourguignon (2015) and Chancel et al. (2022).

For sector-level macroeconomic aggregates and their decomposition in section 3.1, I combine data from the United Nations System of National Accounts (UN SNA 2023), both online and in a custom data retrieval from archival records,⁴ with data from the GGDC/UNU-WIDER Economic Transformation Database (ETD) (Kruse et al. 2022),⁵ as well as data from ILOSTAT (2024) and the World Bank (2024). For aggregate inequality statistics I rely on the World Inequality Database (see wid.world and Chancel et al. 2022).

For contemporary microdata on agricultural income distributions, I turn to a custom retrieval from the universe of ILO labour force surveys (ILOSTAT 2024)—covering more than 100 countries and more than 1000 country-years—and to the Rural Livelihoods Information System (RuLIS), a publicly available dataset from FAO, IFAD and the World Bank (FAO 2024). The latter covers fewer than 40 countries and 100 country-years, but does so with the highest-quality survey data on agricultural household incomes, harmonized according to the criteria of the World Bank Living Standards Measurement Study’s Integrated Surveys on Agriculture (LSMS-ISA; see, e.g., (Carletto and Godoiay 2019)). Similar datasets to that of the ILO include the World Bank’s International Income Distribution Database (I2D2, accessed offline but now accessible at [link](#)) and the Luxembourg Income Study (LIS 2024). However, these suffer from a similar incomplete coverage of agricultural labour incomes as the (ILOSTAT 2024) data, so I prefer the RuLIS data for benchmark estimates of agricultural household labour income inequality.⁶ The choice of microdata and its consequences are discussed further in section 4.

To understand the role of the corporate sector in agriculture—and the place of capital income along the agricultural income distribution—I draw on UN SNA (2023) as above, as well as the Georgia Survey of Agricultural Holdings (Geostat 2020), whose unique advantage will be discussed further in section 5.

Finally, to highlight the current and potential role of social protection in agricultural structural transformation, I draw on novel data from Fisher-Post and Gethin (2023).

From this discussion of data sources and their uses, I turn to several sets of results in the sections that follow.

⁴ For more on the latter, see Bachas et al. (2022).

⁵ Data available for download at <https://www.wider.unu.edu/database/etd-economic-transformation-database>.

⁶ For family farms, household farm income comprises the majority of agricultural incomes, so it would be conceptually fraught (and empirically challenging) to estimate an individualized labour income distribution rather than an equivalized household labour income distribution. More consequentially, we must unfortunately remain on the distribution of labour incomes, because capital income in agriculture is largely unobserved in reference surveys. This is discussed further in section 5.

3 Long-run trends, 1950–2021

3.1 Agricultural employment and productivity

In this section I analyze long-run trends in agricultural employment and productivity, to understand the implications of structural change for inequality in the sector.

I start with accounting identities for total production Y and the workforce L , where Y comprises both agricultural and non-agricultural production and L comprises both agricultural and non-agricultural workers:

$$\begin{aligned} Y &= Y_{agr} + Y_{nonagr} \\ L &= L_{agr} + L_{nonagr} \end{aligned} \tag{1}$$

so the share of agriculture in the workforce is given by $\theta_{agr} = \frac{L_{agr}}{L}$. Total productivity is $\frac{Y}{L}$; the productivity of agricultural labour is $\frac{Y_{agr}}{L_{agr}}$; and $\frac{Y_{nonagr}}{L_{nonagr}}$ is the productivity of non-agricultural labour.

It is perhaps worth noting that some industrial activities—e.g., the production of agricultural inputs; or value-added on food processing—are adjacent but not included in the agriculture sector strictly defined in the UN SNA (2023) and ILOSTAT (2024). Nonetheless, the ratio between agricultural productivity and non-agricultural productivity can be expressed as:

$$productivity\ ratio = \frac{Y_{agr}/L_{agr}}{Y_{nonagr}/L_{nonagr}} \tag{2}$$

Figures 1 and 2 (and Appendix Figures A1–A5) show the statistics these equations represent. In both high- and low-income countries, agricultural employment (as a share of total employment) has declined since 1950. In developing countries, the average proportion of agricultural employment today is approximately what it was in 1950 for developed countries: Roughly 30% of the workforce remains in agriculture in developing countries. And this proportion remains above 50% in sub-Saharan Africa. Meanwhile in developed countries today there is little room for agricultural employment to decline any further, as it represents only approximately 5% of total employment.

While employment in agriculture has declined, its relative productivity (relative to other industries) has increased.⁷ In developed countries, this convergence happened during the mid- to

⁷ Note that the strict macroeconomic accounting of equations (1) and (2), does not account for differences in hours worked across sectors of employment. The strong form of the implicit assumption presented here is that that

late 20th century, and since then agriculture has remained on average roughly 70% as productive as non-agriculture. Meanwhile in developing countries, on average, such convergence was delayed until the turn of the millennium. Only in recent years has the average productivity of agricultural employment risen above 20% of the average productivity of non-agricultural employment, now nearing 40%. This ratio is closer to 50% in Latin America and the Caribbean, and in the Middle East and North Africa (with earlier increases); while it is closer to 30% (and only recently increasing) in Asia and the Pacific and in sub-Saharan Africa.

Just as developing countries' agricultural employment (as a share of total employment, i.e., 30%) mirrors today what it was in developed countries in 1950, so too is developing countries' agricultural productivity (relative to non-agricultural productivity, i.e., 40%) similar to what it was in developed countries in 1950.

3.2 The inverse relationship in focus

The global and regional trends on agricultural employment and productivity beg the question of whether and to what extent agriculture's relative decline in employment is related to its relative increase in productivity: Can we be sure that the decline in employment and the increase in relative productivity are happening in the same countries at the same time? To settle this question, I set up the following regression model:

$$y_{it} = \beta_0 + \alpha_i + \gamma_t + \beta_1 x_{it} + \beta_2 nnipc_{it} + \epsilon_{it} \quad (3)$$

where y represents dependent variables on agricultural productivity or value added; x is agricultural employment as a share of the workforce; with country α_i and time γ_t fixed effects; and controlling for per capita net national income $nnipc$.

Results are presented in Table 1. We can observe that a 1% decrease in the workforce share of agricultural employment is associated with 0.69% *increase* in the productivity of agriculture, relative to non-agricultural employment. Similarly, a 1% decrease in the workforce share of agricultural employment is associated with only a 0.34% decrease in the agriculture's share of total value added. (If agriculture were equally productive as non-agricultural employment, a 1% decrease in agricultural employment would decrease agricultural value added by precisely 1%.⁸)

ratio is 1.0—that workers in both sectors work the same number of hours. The weak form of this assumption is that the ratio of average hours worked in agricultural employment to average hours worked in non-agricultural employment, has not changed.

⁸ If the two sectors were equally productive [with geometrically similar inequality, and a representative sample of workers switching sectors], a change in relative employment share would not be associated with any change in relative productivity.

Meanwhile, if we examine the same question in levels rather than in shares, I find that a 1% decrease in total agricultural employment (regardless of whether non-agricultural employment increases or decreases, in levels) is also associated with a 0.34% *increase* in average agricultural productivity per worker (in levels; i.e., also regardless of non-agricultural productivity per worker). And a 1% decrease in total agricultural employment is associated with only a 0.61% decrease in total agricultural production (where, if the marginal worker were of average productivity, the decrease would be 1%).

Using country and time fixed-effects in this regression approach ensures that we are not observing a spurious correlation globally. Rather, *within* countries over time, a decrease in agricultural employment is associated with an increase in agricultural productivity—both in absolute terms, and relative to non-agricultural employment.

The result holds for all regions, and for both developing and developed countries, and for all time periods in our sample. In all, these headline results are consistent both with a Lewis (1954) process, in which the marginal agricultural worker is essentially underemployed; and with a Gollin et al. (2014) process, in which a significant share of labour is misallocated (from an economic efficiency standpoint) when it is allocated to agriculture. It would require micro-data to quantify changes in productivity at the level of the workers who change sectors,⁹ but the results are highly suggestive that those who leave agriculture are at the lower end of the agricultural income distribution.

Agricultural employment declines as per-capita income rises, and relative agricultural productivity (relative to non-agricultural) increases with development, as well.

3.3 Agricultural structural change

Another way to look at this is via decomposition à la McMillan et al. (2014) and de Vries et al. (2015).

Total productivity is:

$$P_t = \sum_{i=1}^2 \theta_{it} \cdot P_{it} \quad (4)$$

⁹ It would be necessary to show that the same workers leaving agriculture become more productive (and to adjust for self-selection between sectors), and not only to show as these results do, that on average the per-worker agricultural productivity increases as labourers leave agriculture. In fact, with these results it would remain possible, if unlikely, that workers who remain in agriculture become more productive as a direct result of others leaving agriculture. *A priori*, this seems less likely than the composition effect, that the workers who leave agriculture are, on average, below mean productivity, both in the agriculture sector they leave and of course in the non-agriculture sector they join.

where total productivity P at time t is equal to the weighted mean of productivity levels across sectors i , weighted by their share in the workforce θ_{it} , as in equation (1) above.

Whereas those authors (above) were concerned with several sectors (including to distinguish between manufacturing and services), here I limit the analysis to a two-sector economy, agriculture (or primary production; i.e., including forestry and fishing) vs. non-agriculture.

It can then be shown that changes in overall productivity are decomposable into changes in the productivity within sectors, vs. structural change as the compositional change of labour allocation across sectors:

$$\Delta P = \underbrace{\sum_{i=1}^2 \Delta P_i \theta_{i,t+k}}_{\text{within-sector}} + \underbrace{\sum_{i=1}^2 P_{i,t} \Delta \theta_i}_{\text{structural change}} \quad (5)$$

The within-sector term explains how much of overall productivity growth owes to sectoral productivity growth, holding sectoral labour allocation constant; while the structural change term explains how much of overall productivity growth owes to labour reallocation across sectors, holding sectoral productivity constant.

Following this framework,¹⁰ Table 2 shows productivity growth over time, by region, and the extent to which it is explained by sector-level changes in productivity vs. labour reallocation across sectors.

These results show that overall productivity growth is largely explained by within-sector productivity growth—above all in the developed countries of Europe and anglophone countries—but with the notable exception of sub-Saharan Africa, where much of per-capita growth over the past 60 years can be explained rather by changes in sectoral composition rather than by changes in within-sector productivity.

Of course, these labour productivity statistics can be misleading if they are taken to represent the contributions of both labour and capital to value-added, whereas they only discuss labour productivity by industry, from labour force statistics. Capital income (within an industry) accrues to individuals or households whose employment may be listed outside the industry, or to individuals outside the labour force. In agriculture, any returns to capital may accrue to investors who are not included among agricultural workers in labour force statistics. We overstate

¹⁰ Note that de Vries et al. (2015) go one step farther and further decompose the ‘structural change’ term into a ‘static’ and a ‘dynamic’ term. However, for the purposes of this analysis we can agree with McMillan et al. (2017) on the difficulty of interpreting that further disaggregation, and the sufficiency of considering structural change itself as a dynamic process.

the average (income and) productivity of farmworkers if we include capital income that accrues to individuals whose primary employment is elsewhere. The labour productivity statistics in agriculture may in fact be overstated, then, if owners of capital are disproportionately found outside of agriculture. I will return to this point on capital income in Section 5 below.

In short, the results from this decomposition complement those from sections 3.1 and 3.2 above. In the previous section we observed that relative agricultural productivity increases as labourers leave agriculture, and here we observe that this is explained more by within-sector changes in productivity than by between-sector changes in the composition of the workforce—despite the fact that, globally, workers are leaving agriculture and that agriculture is the less-productive sector. The productivity gap between sectors is shrinking, but the productivity gap within sectors is expanding.

However, as above, these results do not tell us whether we are observing a *real* productivity increase in the same workers who leave the agriculture sector (or in those who stay behind), or whether it is simply the case that low-productivity (i.e., below-average productivity within the sector) farm workers are the ones leaving agriculture. Nor do these results tell us what is happening to the income distribution within each sector, as labour leaves agriculture. It is likely that inequality is increasing within the non-agricultural sector (and especially within certain industries such as finance and natural resource production), but the effect on the income distribution within agriculture is more theoretically ambiguous. The long-run macroeconomic (and total income distribution) data does not allow us to observe sector-level income distributions directly.

3.4 What does this mean for agricultural inequality?

A naive cross-section shows that the same countries with the highest agricultural productivity gap are also the countries where inequality is highest. (See Appendix Figure A6.) Of course, there are many explanations for why developing countries (where agriculture is the most significant part of the economy) also exhibit the highest levels of inequality.

To better understand the role of sector-level changes on inequality within and between sectors, we can also decompose total inequality (as the variance of incomes), into ‘within-sector’ vs. ‘between-sector’ components, as follows:

$$\underbrace{V(y)}_{\text{total inequality}} = \underbrace{[\theta_a V(\mathbf{y}_a) + \theta_b V(\mathbf{y}_b)]}_{\text{within-sector}} + \underbrace{[V(\bar{y}_a, \bar{y}_b)]}_{\text{between-sector}} \quad (6)$$

where $V(y)$ is the overall economywide variance of all incomes, decomposed into within- and between-sector variances of agricultural and non-agricultural incomes; θ is the share of

the workforce (as in equation [1] above); and a is the agricultural sector and b is the non-agricultural sector (rather than agr and $nonagr$ above).

Over the time period for which we can compare sector-level productivity statistics with total inequality statistics (i.e., since 1980),¹¹ this decomposition can be seen visually in Figure 3.

Inequality between sectors—the productivity gap—explains very little of overall inequality, despite their correlation. Where perhaps it did so in Asia and sub-Saharan Africa in earlier decades, more recently all regions have converged to an equilibrium state where most of inequality is explained by disparities within sectors. Overall inequality has increased (Chancel et al. 2022)—and, globally, agriculture is becoming more like non-agriculture. Within-sector inequality explains overall inequality. A more disaggregated decomposition (with detailed sector-level income data, worldwide) would highlight which industries have the greatest dispersion of incomes. In this setup, where we observe total income distributions and the average labour productivity by sector¹² but do not observe the dispersion of incomes (or productivity) within a sector, we can only conclude that within-sector inequality matters more than between-sector inequality, but we cannot observe in which sector it matters most.

Another way to see this is to run a two-way fixed effect regression in the form of equation (3) above, examining whether changes in overall inequality are explained by changes in the productivity ratio between sectors. Reassuringly, the result of this regression agrees with the results in Figure 3 from equation (6) above. The coefficient is only marginally significant ($0.05 < p < 0.10$), with a negligible effect size: A 1% increase in the ratio of agricultural to non-agricultural productivity¹³ is associated with only a 0.06% reduction in inequality (expressed as the ratio of top 10% to bottom 50% average incomes). The implication is clear, as above, that agricultural structural change is not explaining much of overall inequality.

Since the (average) productivity gap between sectors does not explain much of overall inequality, let alone its well documented increase in recent decades, it is unfortunate that the vast majority of available, long-run, global income data—on gross productivity levels, by sector; and

¹¹ Reliable, economywide inequality statistics are largely unavailable prior to 1980, especially in developing countries (see Chancel et al. 2022).

¹² Again, there is the notable caveat that we do not here observe capital income allocation by sector—neither the capital share by sector, nor the sector from which investors owe their wealth. To the extent that capital income (corporate profits) in agriculture accrues to earners who do not declare agriculture as their sector of employment, then we are overestimating the average labour productivity of agriculture, because we are underestimating the number of people who earn income from the sector (or overestimating the amount of income that accrues to *labour* in the sector).

¹³ Usually this is a convergence from relatively lower-productivity agriculture to relatively higher-productivity non-farm employment, as in Figures 1 and 2.

on overall economywide income distributions; and particularly in developing countries—does not let us examine patterns of income distribution within sectors.

However, there are some datasets that do allow us to look more closely at agriculture-sector inequality.

4 Contemporary agricultural inequality

4.1 Agricultural household income distributions

To understand how inequality has changed within sectors, then—and within the agricultural sector specifically—it is not sufficient to look at headline trends in aggregate production, employment, and average productivity levels; nor to look at economywide inequality patterns. Survey microdata (with industry markers) allows us to estimate within-sector income distributions.

One such source—the universe of labour force surveys from the ILO—would seem to suggest that agricultural income inequality is quite similar to overall income inequality, as in Figure 4. However, several caveats are in order. First, to say that within-agriculture inequality is similar to inequality overall does not suggest that the income levels are the same. (See Appendix Figure A7.) Second, this may be an artifact of imperfect data.

If instead we take another source of survey microdata, for the same country-years, and attempt to apply the same income definition (income among agricultural employees), we end up with a different result. Figure 5 shows the difference between the ILO estimates of Figure 4 and those which are given by household income and expenditure surveys in RuLIS. If the two income distributions across surveys within the same country-year were perfectly aligned (ILO vs. RuLIS), their Gini coefficients would sit perfectly on the 45-degree line of the graph in Figure 5. That they do not suggests measurement error in one or both sources. In principle, the former (ILO) is based on a 30-day recall (seasonally adjusted) while the latter (RuLIS) is based on a 12-month recall. This could be one of the principal drivers of the discrepancies (in which the 30-day recall measurement of labour force surveys [ILO] exhibits wider variance and therefore more inequality than does the 12-month recall measurement of household income surveys [RuLIS]). Beyond data collection differences, data cleaning differences (e.g., prevalence and treatment of outliers) could also explain the systematic (if not universal) difference across institutional sources.

In any case, the ‘employee’ definition of agricultural income is insufficient. Agriculture, more than any other industry, is largely practiced as a household enterprise, with less of its total production in employer-employee relationships (whether formal or informal) than in other in-

dustries. It is difficult to predict *a priori* what effect this improvement in measurement should have on inequality. Missing agricultural ‘self-employment’ household income could be found disproportionately among the poor, or among the comparatively richer, so the issue is theoretically ambiguous and becomes an empirical question.

Empirically, Figure 6 shows the effect on inequality of including a ‘household income’ definition on agricultural income,¹⁴ rather than the ‘employee income’ (wage income) definition of Figures 4 and 5. Agricultural inequality is generally (if not necessarily or always) higher under the ‘household income’ definition. This means that comparatively richer agricultural households disproportionately earn agricultural income outside of employee relationships.¹⁵ When agricultural household income is well measured, inequality in the sector is higher.

But by how much does agricultural ‘employee’ income inequality underestimate agricultural ‘household’ income inequality? Figure 7 shows the magnitude for sub-Saharan African countries in RuLIS. It turns out that household farm income is more than 60% of total income for agricultural households—regardless of where on the sector-level income distribution the household is found; whereas agricultural wage income is never more than 25% of total household income, on average. As a share of household income, agricultural wage income is (relatively) increasing with total income—but not by as much as total income itself is (absolutely) increasing. That difference in the rates of increase in the share of agricultural wage income vs. in the levels of total household income, would explain the result in Figure 6: To only consider wage income underestimates inequality in the sector, as it underestimates total household income from agriculture, and particularly among higher-income households in the agricultural income distribution.

4.2 Agricultural vs. overall income inequality

We can also place agricultural household income into perspective, on the total household income distribution, for the 40 RuLIS countries and the household income concepts it covers. While there would not be any re-ranking among agricultural households (whose sources of income even outside of agriculture, are already accounted for and included in Figure 7), Figure 8 shows the relative place of agriculture in the total household income distribution (here including both agricultural and non-agricultural households, the latter of which earn less than

¹⁴ The agricultural income distribution comprises those (households) for whom agricultural income represents more than 30% of total income, following the definition offered by RuLIS itself (FAO 2024).

¹⁵ Of course there can also be a re-ranking effect: It is not necessarily the same agricultural households and individuals who appear richest under the agricultural household-income definition, as earned the most under the agricultural employee-income definition.

30% of their income from agriculture, by construction).¹⁶ It is clear that agriculture is the primary employment of the poorest—whether as ‘employees’ or in ‘self-employed’ household farm income—whereas agriculture does not matter as much for (as many of) those who earn more. Off-farm self-employment remains an important source of income even for the richest households (indeed its importance increases with income), and off-farm employee income also increases in importance, but agricultural income’s share declines drastically as total income increases.

Of course, it is also important to note that these survey datasets do not capture capital income—corporate profits, income from investment, and passive income from ownership—neither in agriculture nor in the non-agricultural sector. This would likely not change the results very much for the left-hand side of the distribution (capital income generally only matters at the top of the distribution), but it could have an exaggerated effect both on income component shares and even on the ranking of earners at the top of the income distribution. In survey data we largely do not observe capital income—a point which I will return to in section 5 below.

Meanwhile, the data from RuLIS allows us to revise the results of Figure 4 above, from ILO labour force surveys. The revised results are presented in Figure 9. Agricultural inequality now no longer appears to mirror overall inequality. The agricultural sector is less unequal than the economy as a whole. Important caveats remain, most notably on the magnitude of capital income in the sector, but the revised picture is striking. On the face of it, inequality is less significant in the agriculture sector than in other sectors. This would support the Kuznets (1955) logic, where agriculture is the low-inequality sector and other sectors are higher-inequality.¹⁷ For Kuznets this was abstract theorizing (only partly testable in the data, for a small sample of developed countries over the course of their history), but the most rigorous agricultural income survey data on developing countries confirms it.

¹⁶ Appendix Figure A8 shows the same graph for Asia, with similar trends. (Note that there are fewer Latin American countries in RuLIS, so the region-wide average is less salient.)

¹⁷ Note that the further implications of Kuznets (1955) is that overall inequality widens before it shrinks or stabilizes, as labourers leave agriculture (the low-inequality sector) for other sectors (higher in inequality)—implying an effect in both between-sector inequality (to the extent that agriculture is the lower-income sector on average) and within-sector inequality: Wven if the latter held constant (although it might itself increase if only low-wage labourers leave agriculture), the change in labour-force weights toward the inequality levels of non-agriculture would substantially increase the overall measure of inequality. For more on this, see Kanbur (2019).

Even if agriculture is the lower-inequality sector, inequality in agriculture remains important, because it is the lower-income sector.^{18,19} I will also discuss in section 5 the relative concentration of capital income in agriculture (relative to other sectors), which would (disproportionately) increase inequality in the sector to the extent that it is found among the top earners.

Unfortunately, with the high-quality data in RuLIS (and, to a lesser extent with labour force surveys), long-run time series do not exist to show us how inequality within agriculture (or within non-agriculture) has been changing over time, so we cannot put empirical results on the entire chain of Kuznets' reasoning. Kuznets (1955) did not emphasize the dynamics of within-sector inequality, beyond the mechanical relationship between labour's exit from the low-inequality sector and its entry to the high-inequality sector. Whether inequality has been changing in the sector—on the intensive margin, from dynamic advantages or disadvantages among those who remain in agriculture; or on the sector's extensive margin, from simple exit of agricultural labour at one end of the distribution or the other—remains an open question.

5 Capital income in agriculture

5.1 Corporate-sector agricultural value added

As household survey data does not observe firm-level data or corporate profits (and observes little other capital income), it is very difficult to say anything about the role of the corporate sector (and capital income) in agricultural inequality.

As an entry point to this discussion, Figure 10 shows the relationship between the size of the corporate sector within agricultural production, and overall inequality. Of course, this does not tell us much about inequality within agriculture.²⁰ Overall, there is no global trend correlating corporate agriculture and overall inequality. However, some regional patterns are striking. For example, Latin American countries generally show a high level of inequality, and a sys-

¹⁸ Appendix Figure A7 above illustrated this point, comparing the median income in agriculture to that of non-agricultural sectors. That figure was using ILO labour force survey data to compare median income of employees; the effect would likely be even more pronounced with total household income among agricultural households in RuLIS.

¹⁹ The notion of the relative importance of inequality assumes that variations in income levels matter more among poorer people. At the limit, as average income levels go to zero (or to infinity), the same percentage (or standard deviation) of variation would matter much more (less).

²⁰ The data does not allow us to make that comparison, for lack of RuLIS data in the same country-years as SNA data on corporate agriculture. In any case, the income data in RuLIS reflects agricultural household income (including wage income) but largely ignores capital income, and excludes corporate-sector agricultural holdings entirely. To the extent that corporate-sector agriculture could be expected to correlate with higher inequality in the sector, this would likely be by way of net corporate profits—observed in SNA aggregates, but unobserved in the survey microdata. This is discussed further in section 5.

tematically large share of corporations in agricultural value added. This is suggestive of large land holdings of corporations, and one would expect higher agricultural inequality in Latin America than in other regions (after taking corporate production into account). Similarly, several countries of Eastern Europe and the former Soviet Union show a high level of corporate value added in agriculture, although these countries are relatively lower in overall inequality. Some of the countries in Africa where agriculture is most important (both in North Africa and in sub-Saharan Africa) show very little presence of corporations in agricultural value added. Botswana is a notable exception. While agriculture no longer represents more than 2% of total value added in Botswana (down from 50% at independence and more than 20% through the 1970s), the presence of corporations in the sector likely reflects a legacy of large-scale land-holding.

Meanwhile, Figure 11 shows that corporate profits in agriculture do indeed follow the size of the corporate sector in agriculture. It is not surprising that corporate profits would nearly match the importance of corporate production (both expressed as a share of total value added), although perhaps the employee share would be expected to be higher.²¹ That corporate profits so nearly match corporate value added shows the extent to which corporate profits outweigh employee compensation in agricultural corporate production: The labour share of agriculture within its corporate sector is very low. Of course, this concentration of income among the owners of corporations, has significant implications for inequality in the sector. If we add corporate agricultural income to (the people to whom it accrues, along) the agricultural income distribution, inequality of the sector would increase.

Some of these countries—with national accounts aggregate statistics on the corporate share in agriculture, and on the share of profits within the agricultural corporate sector—make these statistics available as time series in the national accounts database. A few of these series are notable (if not suspicious) for the volatility of their data, but there is little discernible time trend in the global panel. On aggregate, for the select subsample of countries which provide such statistics in their national accounts, we cannot rule out the null hypothesis that the global corporate share of agricultural value-added has not increased since the early 1990s. (The same is true for the profit share of agricultural corporations' value-added.) Of course, the global average statistics may mask significant heterogeneity within countries and among corporations.

5.2 Capital in the agricultural income distribution

Even if we cannot tell whether the corporate sector in agriculture is expanding relative to the size of the agricultural sector overall, the role of juridical holdings in agricultural income in-

²¹ Note that corporate value added is equal to corporate operating surplus (essentially corporate profits, net of depreciation) *plus* compensation of employees *plus* net indirect taxes on production and imports (i.e., net of subsidies on the same).

equality is still worth examining. However, standard agricultural survey microdata is again largely insufficient to respond to this question.

Few agricultural surveys collect income data on juridical holdings, and fewer still share this data publicly.^{22,23} One such survey was recently released by the Geostat (2020). In Georgia, agriculture represents only 8% of economywide value-added in Georgia, but more than 40% of employment (World Bank 2024). And while the corporate sector represents a very small fraction of agricultural output in Georgia—from more than 50% of agricultural output in the Soviet era, to less than 5% today (see, e.g., Lerman 2004)—the results of this exercise can be instructive nonetheless.

Figure 12 compares the upper tail of the agricultural income distribution with and without the inclusion of corporate-sector agriculture. A few caveats are in order, but the headline result is striking. The inclusion of a mere 667 holdings (and 9,204 workers)—approximately 0.1% of total agricultural holdings (and 0.7% of agricultural workers) covered by this survey—raises the Gini coefficient by two percentage points.

This is likely a lower bound on the true effect, for several reasons. First of all, the data is measured as revenue per worker.²⁴ Of course, most workers do not receive as their wage anything near to the average revenue per worker of the farm.²⁵ Moreover, the profits that accrue to the owners of corporate-sector farms are entirely unobserved.²⁶ If we could observe the labourers' real wages and the profits per individual or household, the income distribution presented in Figure 12 would likely be skewed even further at the top.

Second, the farm survey is underestimating its national sample. In comparison to data from the World Bank (2024), holdings in the survey with revenue data represent only 54% of agricultural employment (including both household and employee labour) and at most 33% of agricultural value-added.²⁷ Since the share of missing income (in total income) is larger than the share

²² Juridical holdings are those managed by (government or) private enterprises.

²³ See, for example, the Food and Agriculture Microdata Catalog (FAM, [link](#)), maintained by the United Nations Food and Agriculture Organization.

²⁴ Cost data in this survey presents enough challenges that I do not calculate profits—the mean and median agricultural profits per holding are negative.

²⁵ Labour is but one cost among many, and the income (net of other costs—and perhaps net of profits) per worker would be a more appropriate measure.

²⁶ The survey data reports incomes by source, and some cost data, but only reports the number of workers and not the expenditure that was paid to these labourers. More importantly, even if labour costs were directly observed, we do not observe to whom accrue the profits of the corporate-sector agricultural holdings, after payments to labour and after all other costs (nor how many people would count among the owners and shareholders to split this return on investment).

²⁷ In principle, to use strictly revenue as an income concept in the sector would be closer to the national accounts concept of 'output'—which is always larger than value added, due to intermediate transactions and the same goods

of missing people (in total agricultural employment), the missing persons would be of above-average income: To observe these missing incomes would likely increase overall inequality in the sector.

Finally, the survey—even this survey, which is best-in-class for its coverage of incomes and juridical holdings—does not capture asset values very well, nor the owners of assets. Investment income from agricultural enterprise is simply not a feature of this data, nor of similar survey datasets. In principle, if we observed the total of agricultural value-added in the survey (i.e., all incomes from all transactions in the sector, including all intermediate expenditures), it would be almost superfluous to concern oneself with the value of assets in the sector—their annual returns would be captured within the income flows that comprise value-added. However, since we are not observing more than one-third of aggregate sector-level value added in the survey data, it is also worth pointing out how little we observe of the value of productive assets in these surveys. Land value is not reported (although land area is), nor the value of transactions on land sales.²⁸ Ideally, investments (and profit margins) up and down each agricultural value chain would also be captured—but surveys (of holdings, or households) are notoriously poor at capturing value chain transactions (Barrett et al. 2022). To observe wages and profits along the value chain would significantly improve the measure of agricultural income inequality.

These thorny income concepts and data discrepancies speak to the challenges of adequately capturing all of agriculture-sector incomes and workers. Even so, when we observe corporate-sector agricultural holdings—and even in a country where these are rare—the measurement of agricultural inequality increases significantly.

6 Social protection

If agricultural income inequality increases during structural transformation, what is to be done?

Social protection spending represents a singularly effective fiscal policy tool to combat poverty and inequality, at least in the short run (Grosh et al. 2022). Within social protection generally, non-contributory social assistance expenditures (direct transfers in cash or in kind, whether conditional or unconditional) are particularly targeted toward vulnerable populations—at least in principle—whereas contributory social insurance programmes (largely pensions from formal-sector employment) are by construction less progressive and redistributive.

being sold twice—but in this case we do not observe intermediate transactions by which to calculate value-added as such; and yet the farmgate output observed remains far below total value-added. We are likely observing less than 33% of value-added.

²⁸ For more on land inequality, see Bauluz et al. (2020).

Long established in OECD countries, social safety net programmes play a significant role in high-income countries, at more than 12% of GDP on average (of which more than 5% is for social assistance alone)—a share of the economy that has remained roughly constant for the past 40 years. (See Figure 14, with data from (Fisher-Post and Gethin 2023).) Meanwhile, social assistance has increased in many Latin American countries since the 1990s (notably in conditional cash transfer programmes), and social protection is also increasing in Asia. However, despite a heralded movement toward cash transfer programmes in sub-Saharan African countries (Devereux 2016), the magnitude of this spending remains weak: African countries average only 1% of GDP in social assistance spending.

More than the overall magnitude of this spending, however, it is worth asking whether social assistance spending reaches the poorest and most vulnerable, and among rural and agricultural households in particular.

One graph may be instructive. Figure 13 shows the concentration of social assistance spending, along the household income distribution, globally (weighted by population) for all developing countries represented in the RuLIS data. We see that, overall, social assistance spending is generally progressive, targeted to the poorest half of the population more than to the richer half of the population (if, however, not very progressive within the poorest 50 percent of the population).

Within the agricultural population, however, the story is different. Among agricultural households, social assistance increases with income. This is true both in their propensity to access social assistance, and in the total amount transferred.

Meanwhile, emerging evidence suggests that social assistance spending can catalyze productive investments among smallholder farmers (Daidone et al. 2018). Even if this were not so, there is a case to be made for universal basic income (Banerjee et al. 2019). Taken together, for both efficiency and equity, the evidence on social protection implores governments and international organizations to increase social assistance spending—and to improve targeting to the rural poor.

Note that these long-run macroeconomic social assistance aggregates (and comprehensive survey microdata incidence estimates) do not reflect the most recent spike of social protection spending during the COVID-19 pandemic. However, the early studies on pandemic spending reveal a similar dynamic: High-income countries temporarily increased their social protection spending and (more or less) effectively targeted their poorest citizens, while developing countries did not have the capacity to do so (World Bank 2022).

In sum, social protection should be seen as an effective and underused (or misused) tool for governments to confront inequality in agricultural communities.

7 Conclusion

With these considerations on agricultural income inequality, it is possible to draw conclusions in several directions.

We can first of all conclude that there is a need for more and better data. There is an urgent need to expand comprehensive agricultural household income surveys globally, to understand the nature of changes to the agricultural income distribution.

Second of all, and despite the data's insufficiency to draw long-run global trends on within-sector inequality (strictly within the agriculture sector, sans comparison to overall inequality), we can still point to several important results in cross-section and at country level.

Among these is the importance of precision in measuring agricultural household income. Labour force surveys do not capture the bulk of income in the sector. When surveys measure agricultural household income well, they generally show higher inequality than in the strictly 'employee' concept of agricultural income. They also show that agricultural household income inequality is less than overall inequality—although the households in agriculture remain significantly poorer.

While agricultural household income surveys represent a significant improvement over labour force surveys and other income surveys that do not pay special attention to the sector, even agricultural household income surveys do not always capture well the role of capital income in the sector. To the extent that corporations play an important role in an economy's agricultural sector, even the best-measured agricultural household survey will likely underestimate inequality in the sector. The lack of long-run, individual-linked data on corporate-sector profits in agriculture, means that it remains difficult to understand the nature of value chains, and to whom accrues the most direct income effect from transformations in the agricultural sector.

Further data collection and policy research must focus on vulnerable populations. For example, in the RuLIS data we can observe that, globally, female-headed households are found disproportionately among the poorest rural households. (See Appendix Figure A9.) It was similarly unsurprising—but worth quantifying—to show in section 6 to what extent existing social protection policies miss their targets, and above all among vulnerable agricultural populations.

Agricultural inequality may be less than overall inequality, but it is still substantial, and all the more so for agriculture's continuing role as the employment-of-last-resort for hundreds of millions of the world's poorest people. While there remains little data and fewer analytical efforts to estimate the nature of inequality in agriculture—including an inability to systematically doc-

ument changes over time—we understate the case if we ignore (because we cannot observe) the magnitudes of capital income accruing to top earners in the sector.

Long-run time series data on agricultural income distributions will be able to show in more detail the impact of different types of agricultural structural transformation on employment, productivity, and inequality in the sector.

References

- Anand, S., and Kanbur, R. (1993). ‘The Kuznets process and the inequality—development relationship’. *Journal of Development Economics*, 40(1): 25-52. [https://doi.org/10.1016/0304-3878\(93\)90103-T](https://doi.org/10.1016/0304-3878(93)90103-T)
- Bachas, P., Fisher-Post, M. H., Jensen, A., and Zucman, G. (2022). ‘Capital taxation, development, and globalization: Evidence from a macro-historical database’. NBER Working Paper 29819. Cambridge, MA: National Bureau of Economic Research. <https://doi.org/10.3386/w29819>
- Banerjee, A., Niehaus, P., and Suri, T. (2019). ‘Universal basic income in the developing world’. *Annual Review of Economics*, 11(August): 959–983. <https://doi.org/10.1146/annurev-economics-080218-030229>
- Barrett, C. B., Reardon, T., Swinnen, J., and Zilberman, D. (2022). ‘Agri-food value chain revolutions in low-and middle-income countries’. *Journal of Economic Literature*, 60(4): 1316–1377. <https://doi.org/10.1257/jel.20201539>
- Bauluz, L., Govind, Y., and Novokmet, F. (2020). ‘Global land inequality’. WID.world Working Paper 2020/10. Paris: World Inequality Lab.
- Baymul, C., and Sen, K. (2020). ‘Was Kuznets right? New evidence on the relationship between structural transformation and inequality’. *The Journal of Development Studies*, 56(9): 1643-1662. <https://doi.org/10.1080/00220388.2019.1702161>
- Bourguignon, F. (2015). *The globalization of inequality*. Princeton: Princeton University Press. <https://doi.org/10.1515/9781400865659>
- Carletto, C., and Godoiay, S. (2019). ‘A thing of the past? Household surveys in a rapidly evolving (agricultural) data landscape: Insights from the LSMS-ISA’. *Agricultural Economics*, 50(S1): 51–62. <https://doi.org/10.1111/agec.12532>
- Chancel, L., Piketty, T., Saez, E., and Zucman, G. (2022). *World inequality report 2022*. Harvard University Press. <https://doi.org/10.4159/9780674276598>
- Daidone, S., Davis, B., Handa, S., and Winters, P. (2018). *The household and individual-level economic impacts of cash transfer programmes in Sub-Saharan Africa: Synthesis report*. Rome: United Nations Food and Agriculture Organization (FAO). <https://doi.org/10.18356/005ab3f5-en>
- Devereux, S. (2016). ‘Social protection for enhanced food security in sub-Saharan Africa’. *Food Policy*, 60(April): 52-62. <https://doi.org/10.1016/j.foodpol.2015.03.009>
- de Vries, G., Timmer, M., and de Vries, K. (2015). ‘Structural transformation in Africa: Static gains, dynamic losses’. *The Journal of Development Studies*, 51(6): 674-688.

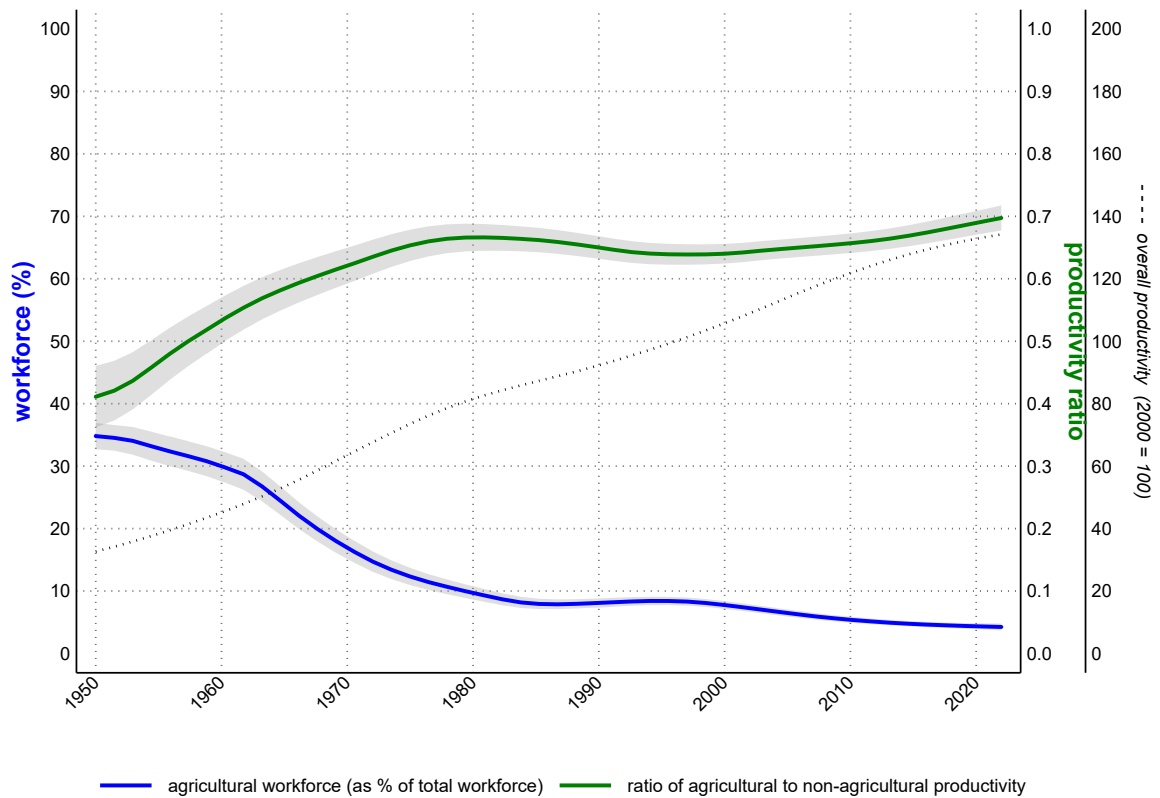
<https://doi.org/10.1080/00220388.2014.997222>

- FAO (2024). *Rural livelihoods information system*. Food and Agriculture Organization of the United Nations (FAO).
- Fisher-Post, M., and Gethin, A. (2023). ‘Government redistribution and development: Global estimates of tax-and-transfer progressivity, 1980-2019’. Working Paper 2023/17. Paris: World Inequality Lab.
- Geostat (2020). *Survey of agricultural holdings*. National Statistics Office of Georgia.
- Gollin, D. (2014). ‘The Lewis model: A 60-year retrospective’. *Journal of Economic Perspectives*, 28(3): 71–88. <https://doi.org/10.1257/jep.28.3.71>
- Gollin, D., Lagakos, D., and Waugh, M. E. (2014). ‘The agricultural productivity gap’. *The Quarterly Journal of Economics*, 129(2): 939–994. <https://doi.org/10.1093/qje/qjt056>
- Grosh, M., Leite, P., Wai-Poi, M., and Tesliuc, E. (2022). *Revisiting targeting in social assistance: A new look at old dilemmas*. World Bank Publications. <https://doi.org/10.1596/978-1-4648-1814-1>
- ILOSTAT (2024). *International Labour Organization, ILOSTAT Database*. International Labour Organization.
- Kanbur, R. (2019). ‘Kuznets and beyond’. *The Oxford Handbook of Structural Transformation*,(): .
- Kruse, H., Mensah, E., Sen, K., and de Vries G. J. (2022). ‘A manufacturing renaissance? Industrialization trends in the developing world’. *IMF Economic Review*, 71(June): 439–473. <https://doi.org/10.1057/s41308-022-00183-7>
- Kuznets, S. (1955). ‘Economic growth and income inequality’. *The American Economic Review*, 45(1): 1–28.
- Lerman, Z. (2004). ‘Successful land individualization in Trans-Caucasia: Armenia, Azerbaijan, Georgia’. In D. Macey, W. Pyle, and S. Wegren (eds), *Building Market institutions in post-communist agriculture: land, credit, and assistance* (pp. 53–75). Lanham: Lexington Books.
- Lewis, W. A. (1954). ‘Economic development with unlimited supplies of labour’. *The Manchester School*, 22(2): 139–191. <https://doi.org/10.1111/j.1467-9957.1954.tb00021.x>
- LIS (2024). *Luxembourg Income Study*. The Luxembourg Income Study.
- McMillan, M., Rodrik, D., and Sepúlveda, C. (2017). ‘Structural change, fundamentals and growth: A framework and case studies’. NBER Working Paper 23378. Cambridge, MA: National Bureau of Economic Research. <https://doi.org/10.3386/w23378>
- McMillan, M., Rodrik, D., and Verduzco-Gallo, Í. (2014). ‘Globalization, structural change, and productivity growth, with an update on Africa’. *World Development*, 63(November): 11–32. <https://doi.org/10.1016/j.worlddev.2013.10.012>
- UN SNA (2023). *National accounts statistics: Main aggregates and detailed tables, 2022*. New York: Statistics Division, United Nations.
- World Bank (2022). *Poverty and shared prosperity 2022: correcting course*. The World Bank.
- World Bank (2024). *World Development Indicators 2024*. The World Bank.

Figures

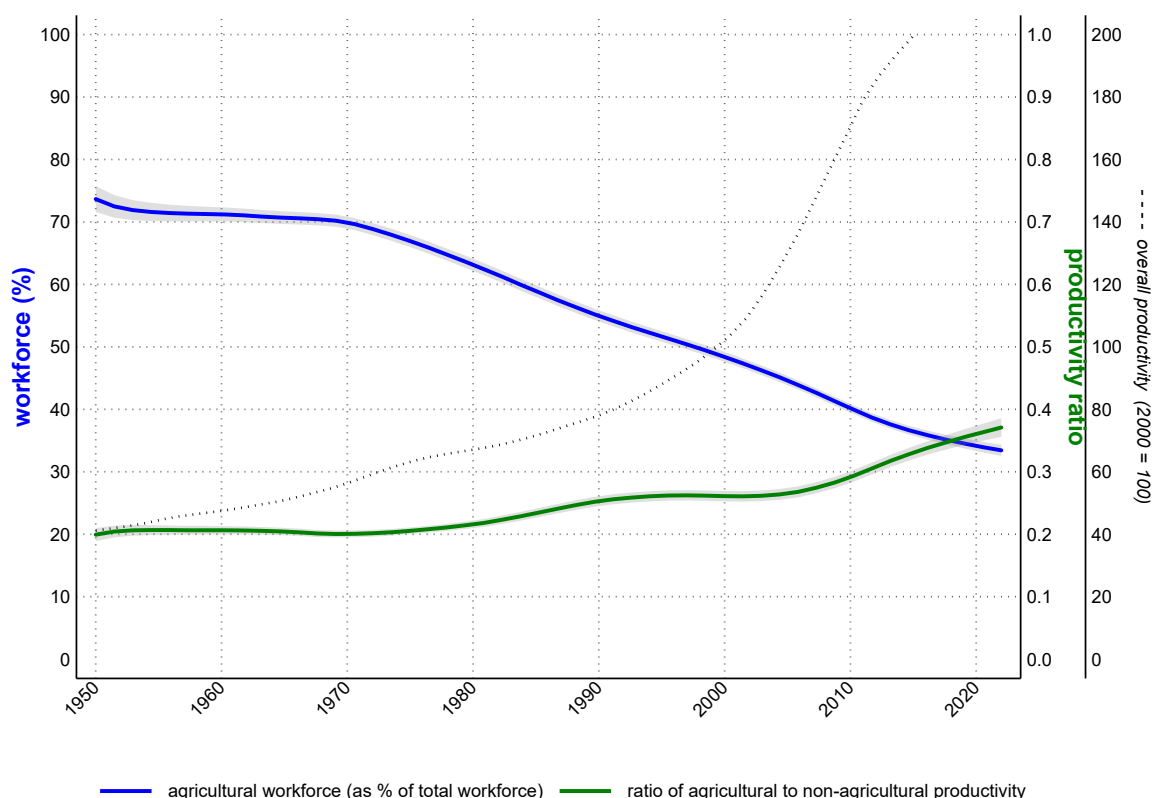
All the figures are the author's own illustrations based on data described in Section 2.

Figure 1: Relative size and productivity of agricultural workforce high-income countries



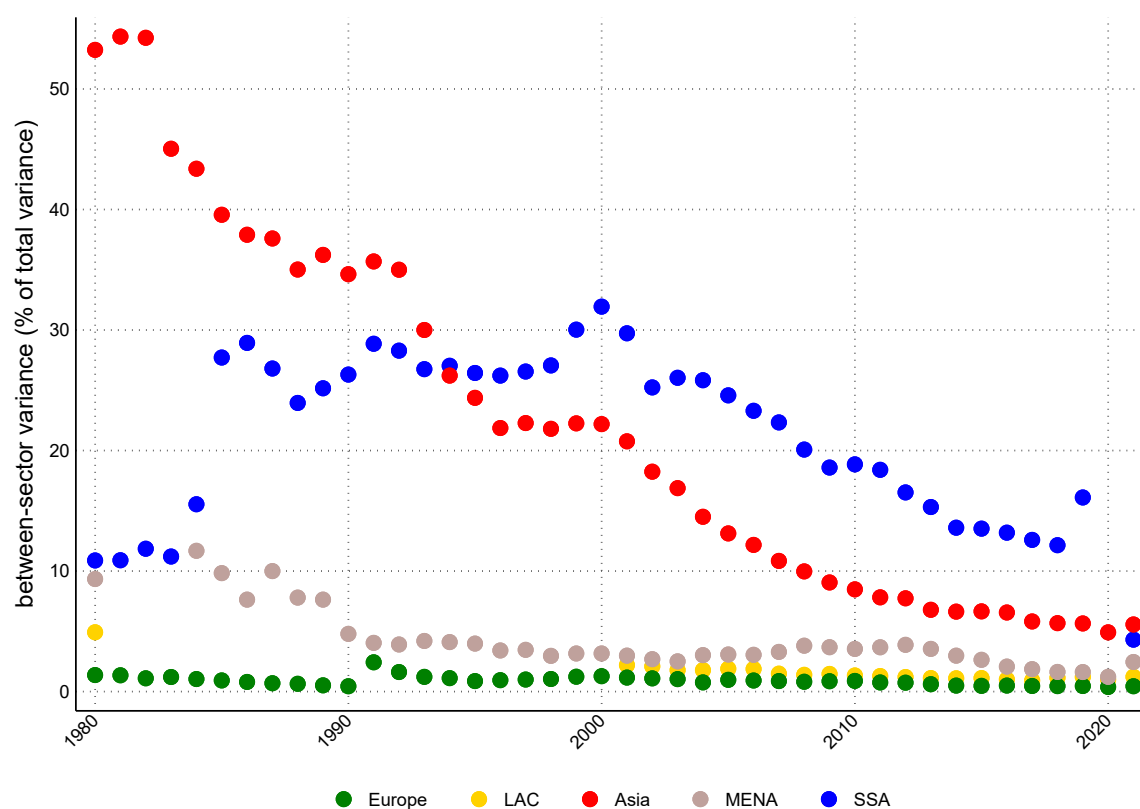
Note: lines represent local polynomial smoothing (with 95% confidence interval) on the size of the agricultural workforce (relative to total workforce) and the ratio of relative productivity in the agricultural sector to relative productivity in rest of the economy. Relative productivity is the ratio of the sector's share of total value added to its share of the workforce (such that economy-wide average productivity is always 1.0), so the ratio of agriculture's productivity to that the rest of the economy shows the agricultural sector's productivity gap. Total productivity (dashed line) is national income per capita (in constant 2022 USD at PPP), with each country's productivity level indexed to 100 at the turn of the century (average in the years 1999–2001). All estimates are weighted by size of national populations. The years 1970 and 1991 show significant structural breaks (sharp expansions) in the sample of countries for which we observe labour force and productivity data.

Figure 2: Relative size and productivity of agricultural workforce —Developing countries



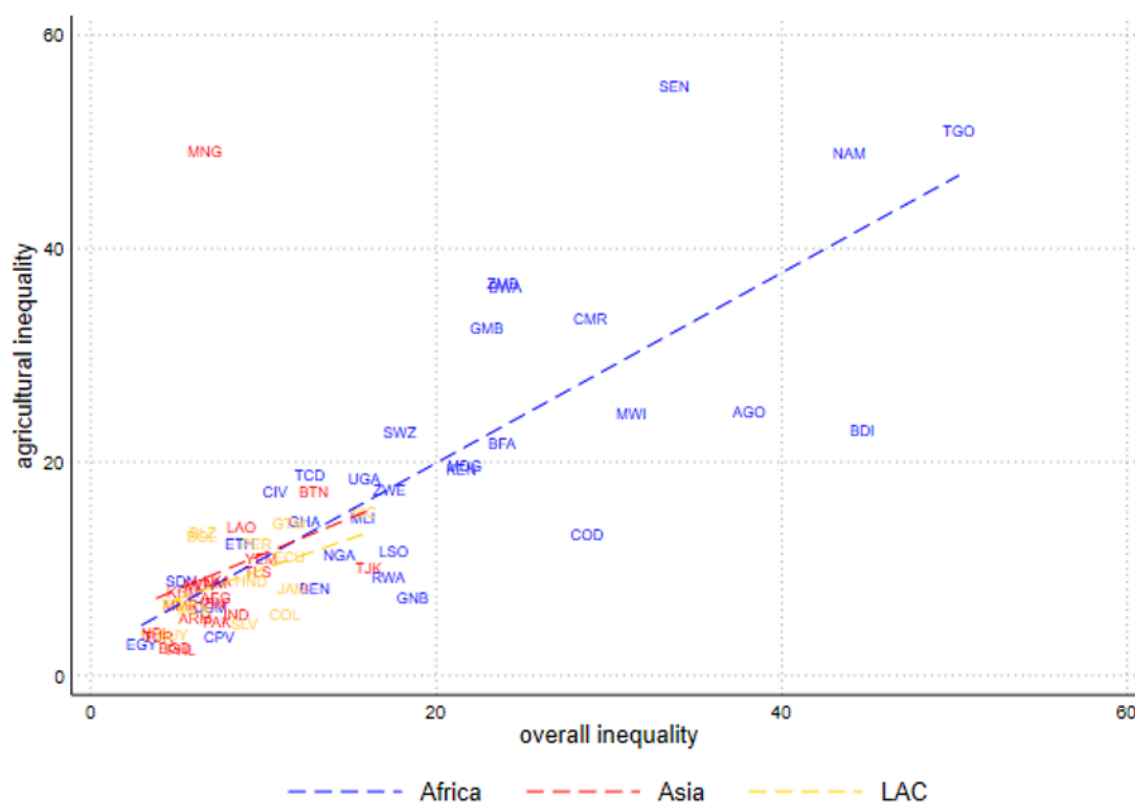
Note: lines represent local polynomial smoothing (with 95% confidence interval) on the size of the agricultural workforce (relative to total workforce) and the ratio of relative productivity in the agricultural sector to relative productivity in rest of the economy. Relative productivity is the ratio of the sector's share of total value added to its share of the workforce (such that economywide average productivity is always 1.0), so the ratio of agriculture's productivity to that the rest of the economy shows the agricultural sector's productivity gap. Total productivity (dashed line) is national income per capita (in constant 2022 USD at PPP), with each country's productivity level indexed to 100 at the turn of the century (average in the years 1999–2001). All estimates are weighted by size of national populations. The years 1970 and 1991 show significant structural breaks (sharp expansions) in the sample of countries for which we observe labour force and productivity data.

Figure 3: Inequality within vs. between sectors



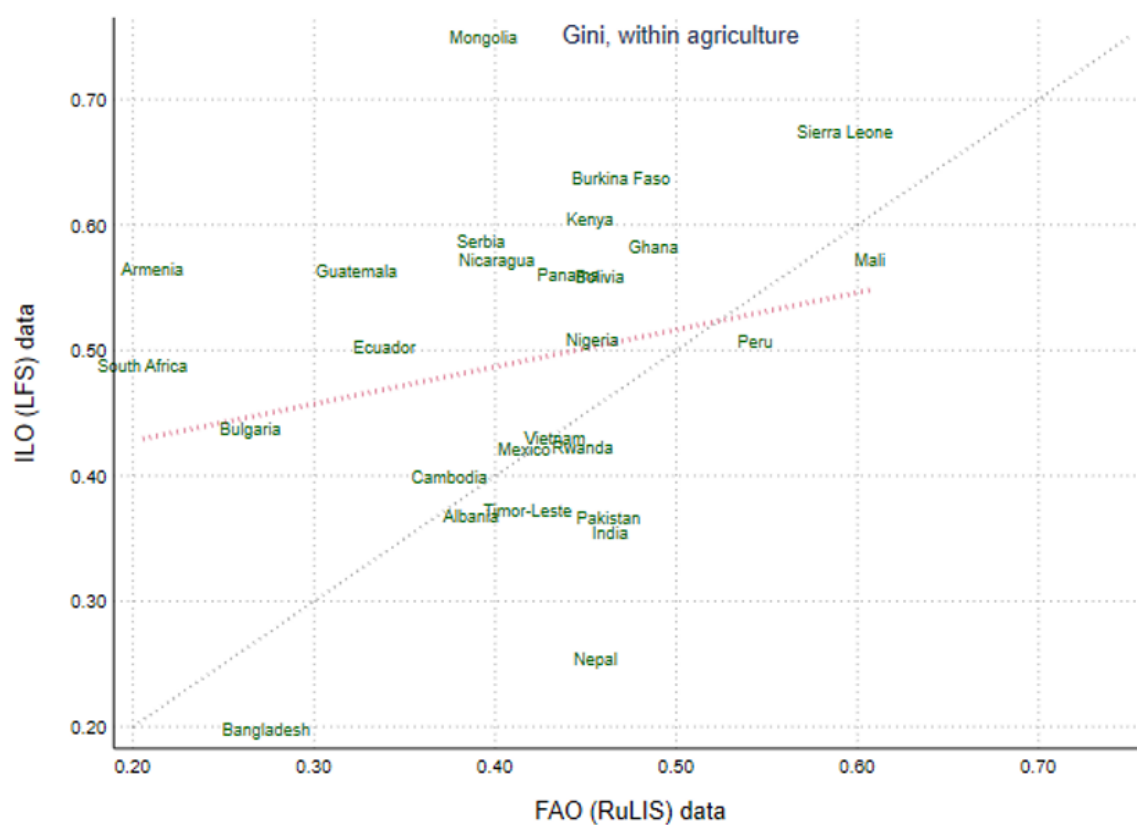
Note: this binned scatterplot presents between-sector inequality from equation (6), by region and over time. Overall inequality is the variance of all incomes, observed in the World Inequality Database since 1980 for most developing countries (see Chancel et al. 2022). Decomposed into within-sector vs. between-sector inequality (as variance), the y-axis values can be interpreted as the percent of overall variance explained by the difference in average incomes between sectors, with the residual inequality explained by the variance of incomes within sectors (weighted by sector shares in the workforce).

Figure 4: Inequality in agriculture vs. inequality overall—ILO labour force surveys



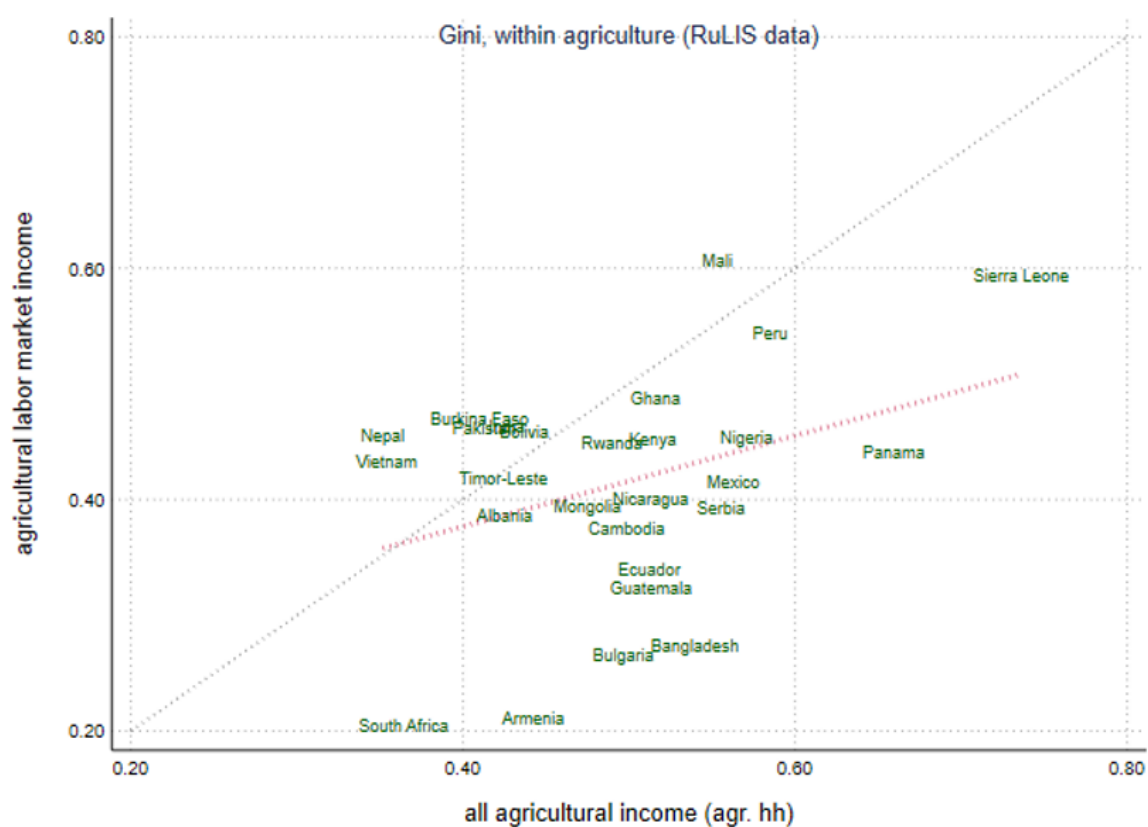
Note: inequality is measured here as the Gini coefficient. The x-axis shows overall economywide inequality (among all employees, in ILO labour force surveys), while the y-axis shows inequality among agricultural employees. Dashed lines represent population-weighted linear trends.

Figure 5: Inequality in agriculture, ILO vs. RuLIS



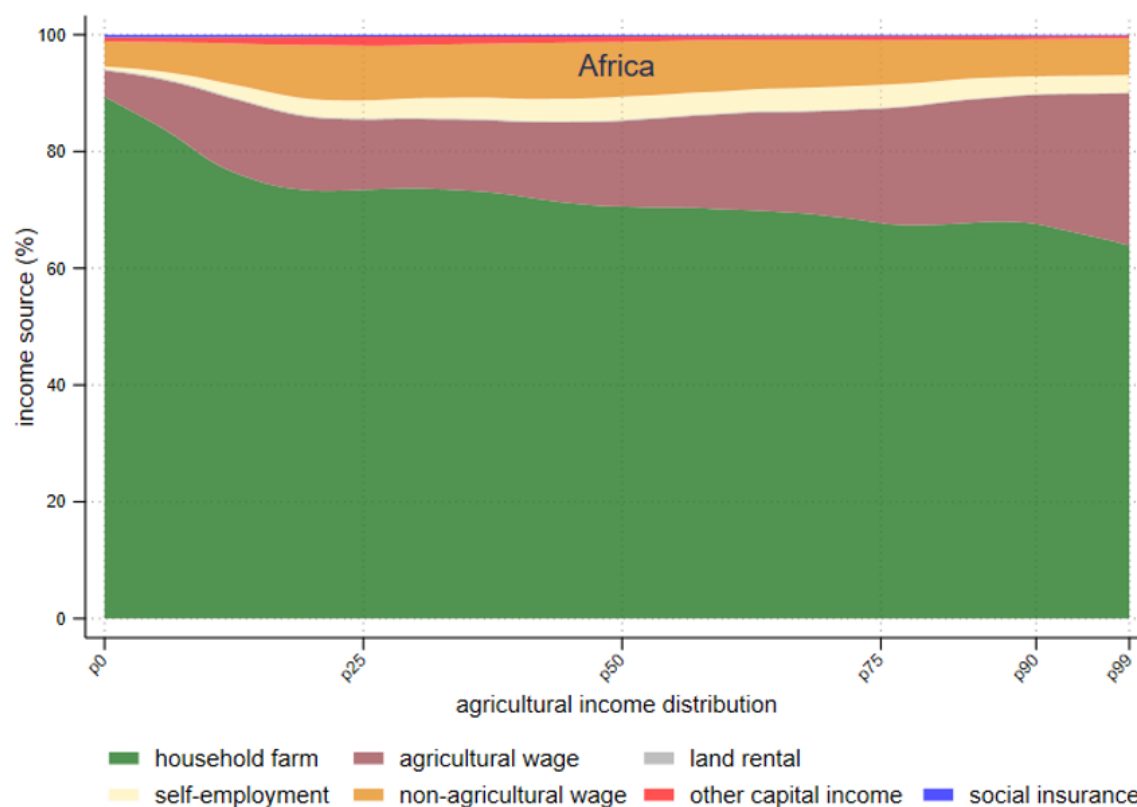
Note: inequality in the agriculture sector is measured as the Gini coefficient for agricultural employees' income, in ILO data compared to that of RuLIS—separate surveys for the same country-years.

Figure 6: Inequality in agriculture, employee income vs. household income



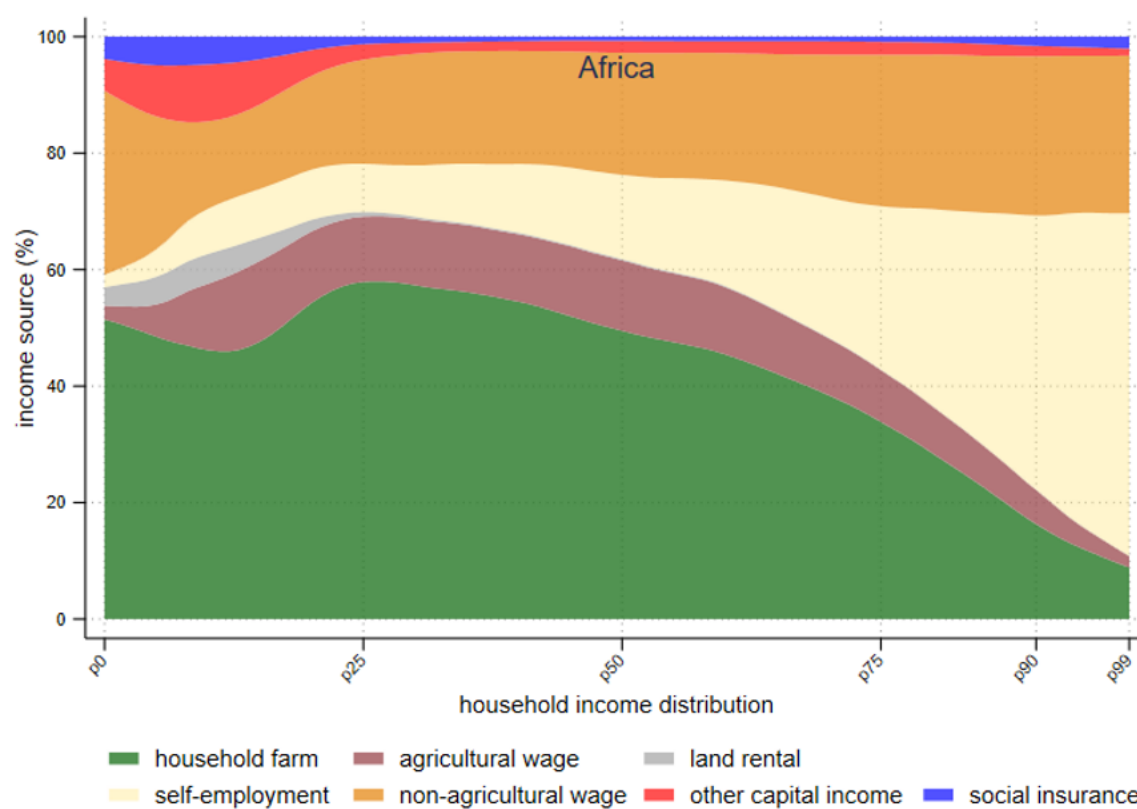
Note: on the y-axis, inequality in the agriculture sector is measured as the Gini coefficient for agricultural employees' income. On the x-axis, inequality in the agriculture sector is measured as the Gini coefficient for agricultural household income. The data source for both graphs is RuLIS (the same country-years' surveys).

Figure 7: Agricultural household income, by income components



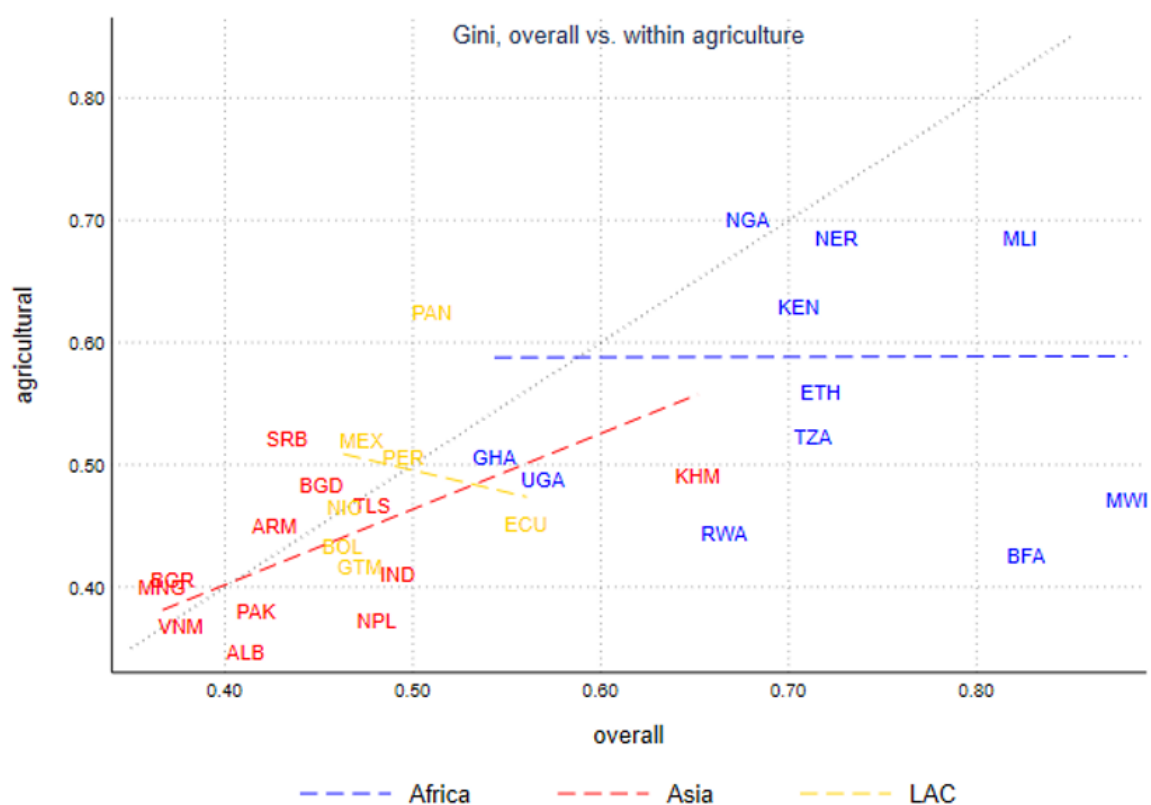
Note: the figure depicts agricultural income inequality, as the distribution of income available in RuLIS, and ranked by total income, but restricted to agricultural households (households whose source of income is at least 30% from agriculture, following the RuLIS definition of agricultural households).

Figure 8: Overall household income, by income components



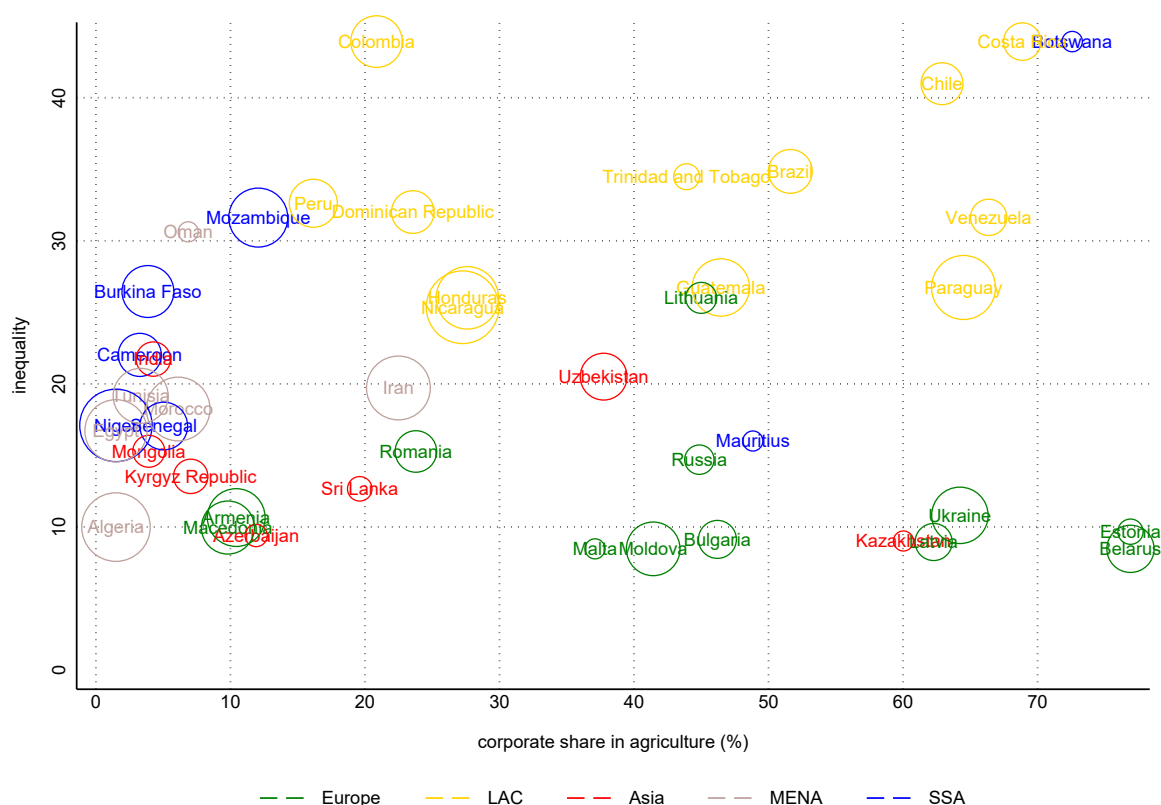
Note: the figure depicts overall household income inequality, as the distribution of all income available in RuLIS, and ranked by total income, for all households (not only those whose income comes from agriculture).

Figure 9: Inequality in agriculture vs. inequality overall—RuLIS household income



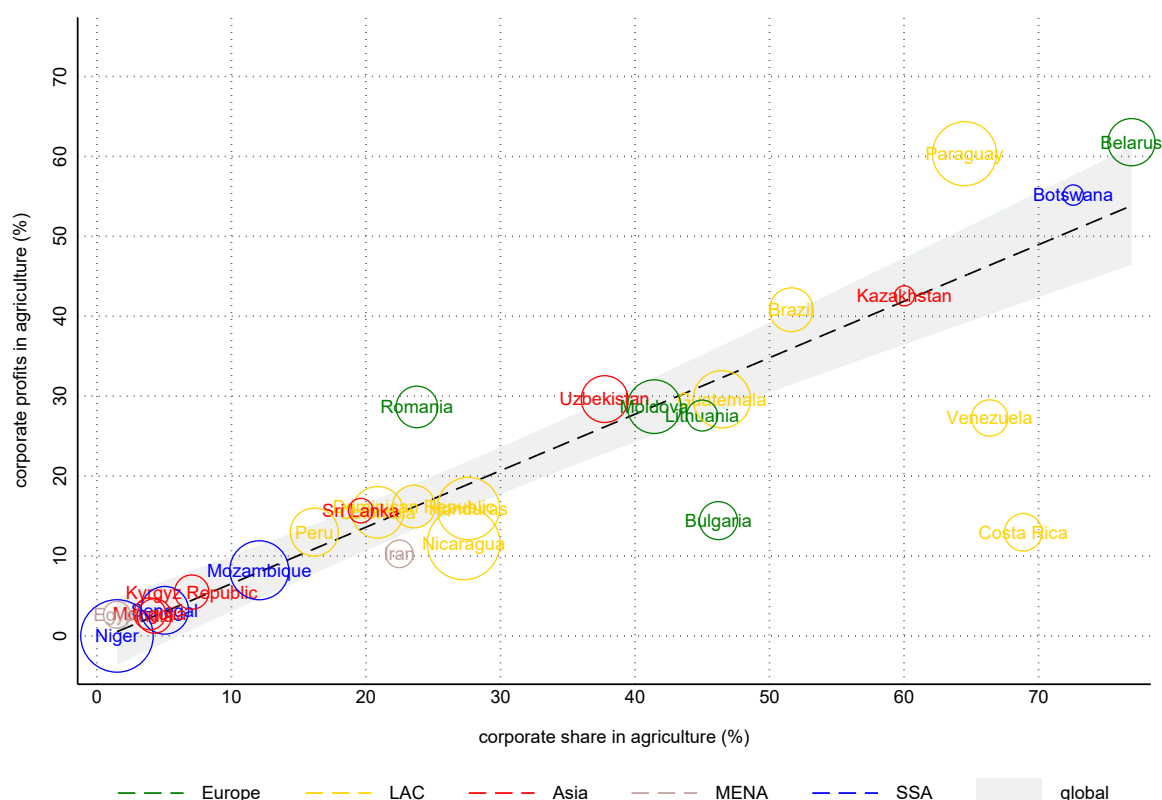
Note: inequality is measured here as the Gini coefficient. The x-axis shows overall economywide inequality (from the World Inequality Database), while the y-axis shows inequality among agricultural households (from RuLIS)—a revised version of Figure 4 above—here with all agricultural income accounted for. Dashed lines represent population-weighted linear trends.

Figure 10: Inequality and the corporate share of agricultural value added



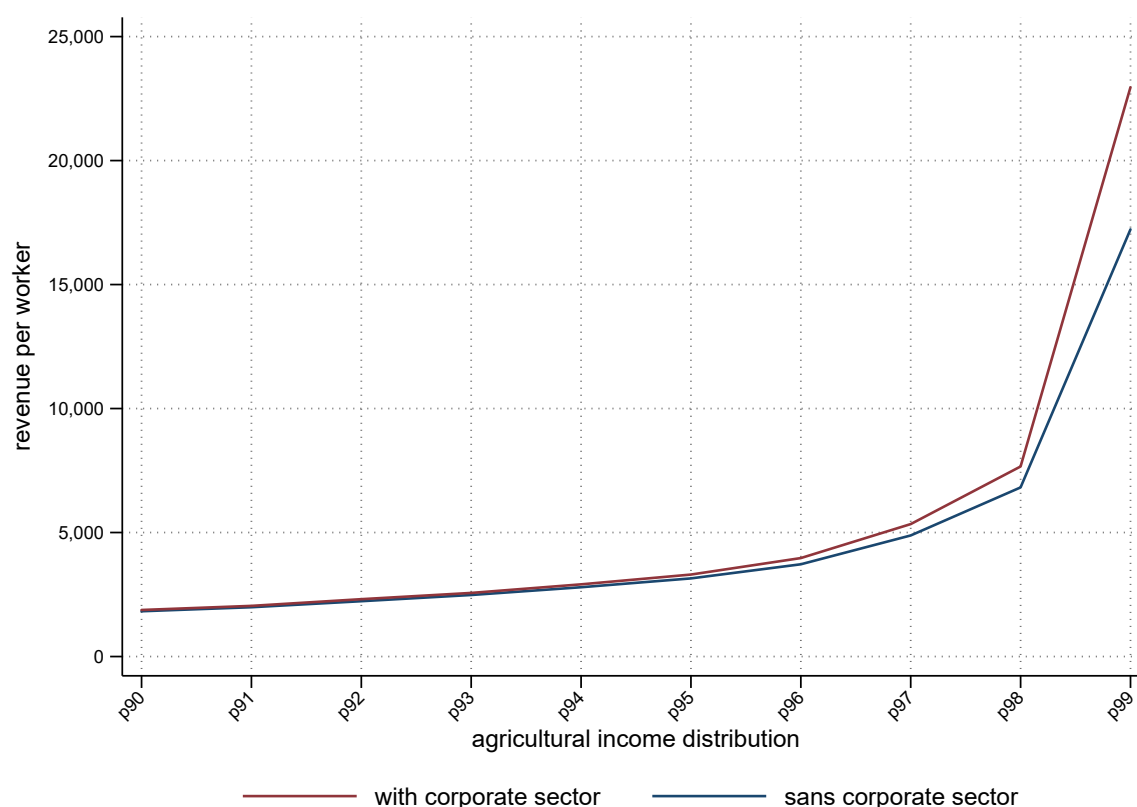
Note: the scatterplot shows total inequality (y-axis) against the corporate share of agricultural value added (x-axis). The hollow-circle overlay shows the weight of the agricultural sector in a country's total value added (ranging in this sample from 1% [Oman] to 40% [Niger] of total value added). Results represent the latest-year values for each country observed in our data, where 35 of 47 countries' values are since 2010 (41 of 47 since 2006; and the earliest year is 1996). Inequality data (from the World Inequality Database) represents total economywide inequality: the ratio of the average income of top 10% income earners to average income of bottom 50% earners. The corporate share of agricultural value added here also includes government agricultural value added (never more than 15% in our sample, and declining over time—less than 2.5% in all but five countries); the residual is the household sector, almost always the majority of agricultural value added.

Figure 11: Agricultural corporate profits and the corporate share



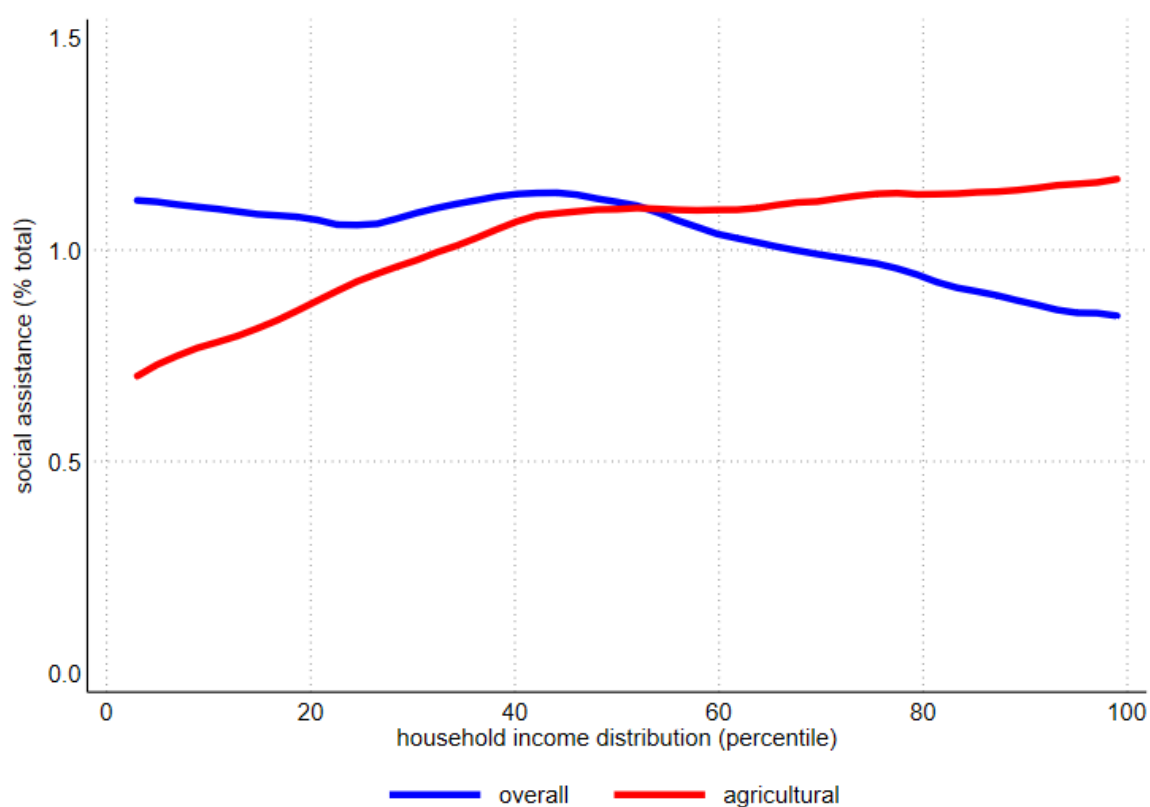
Note: the scatterplot shows agricultural corporate profits (y-axis) against the agricultural corporate share (x-axis), both as a percentage of agricultural value added. The hollow-circle overlay shows the weight of the agricultural sector in a country's total value added (ranging in this sample from 3% [Botswana] to 40% [Niger] of total value added). The dotted line represents global linear fit (and 95% confidence interval) between the two. Results represent the latest-year values for each country observed in our data, although the years do not perfectly match between SNA data sources (MADT Tables 2.3 and 2.6 for agricultural corporate profits; and MADT Tables 5.1 and 5.2 for agricultural corporate share). The corporate share of agricultural value added here also includes government agricultural value added (never more than 6.5% in our sample, and declining over time—less than 2.5% in all but two countries); the residual in both cases is the household sector (including compensation of employees, and mixed income), almost always the majority of agricultural value added.

Figure 12: Revenue per worker, agricultural holdings, Georgia 2020



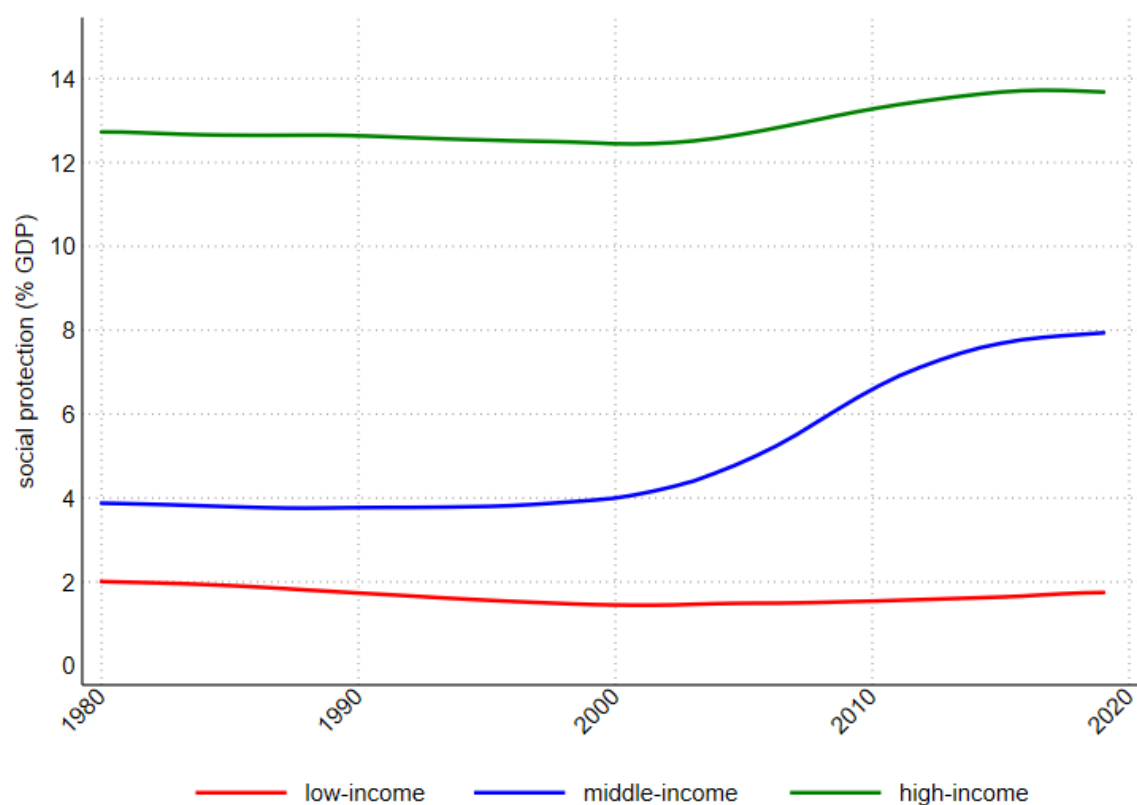
Note: the figure shows the average revenue per worker for the top ten percent of holdings (weighted by number of workers) in Georgia in 2020, drawn from the annual survey of agricultural holdings (Geostat 2020). The red line includes corporate agricultural holdings in the income distribution. The blue line excludes corporate agricultural holdings from the income distribution. [N = 1,379,344 workers (of which 9,204 on corporate holdings) on 571,869 holdings (of which 667 corporate).] Currency is expressed in current 2020 Georgian Lari (GEL). [NB Average per capita income in Georgia in 2020 was 13,254 GEL (World Bank 2024).]

Figure 13: Concentration of social assistance expenditure—Agricultural households vs. overall



Note: this figure shows the concentration curve aggregate social assistance spending, against the agricultural (red) and overall (blue) income distributions. If aggregate spending were equally allocated to each percentile of the distribution, the line would be flat at 1.0% of total spending for each percentile of the distribution. A decreasing (increasing) slope of the line indicates progressivity (regressivity) in social spending. The data source is the RuLIS universe of household income surveys.

Figure 14: Social protection spending by income level, 1980–2019



Note: the figure plots total social protection expenditure (including both social insurance and social assistance), as a percentage of GDP, for high- vs. middle- vs. low-income countries, with data from Fisher-Post and Gethin (2023) and country classification from World Bank (2024). Averages are weighted by national income.

Tables

Table 1: Correlation between agricultural employment and productivity

	productivity ratio (1)	value added share (2)	productivity level (3)	value added level (4)
employment share	-0.639*** (0.0841)	0.321*** (0.0641)		
employment level			-0.344*** (0.0429)	0.615*** (0.0486)
<i>N</i>	5923	5923	5923	5923

Note: this table presents the raw correlation between changes in agricultural employment and changes in agricultural productivity. All dependent and independent variables are in taken in logarithms. Column (1) represents the effect of a 1% change in agricultural employment share in the workforce, on the (per cent change in the) ratio of agricultural to non-agricultural productivity; and column (2) the same [i.e., the effect of a 1% change in agricultural employment share in the workforce], on the (per cent change in) agricultural value added, as a share of total value added. Column (3) represents the effect of a 1% change in total agricultural employment (regardless of its share in the workforce) on the (per cent change in) agricultural productivity (regardless of non-agricultural productivity), while column (4) presents the same [i.e., the effect of a 1% change in agriculture employment levels] on the (per cent change in) total agricultural production (value added). Results are weighted by population and control for log national income per capita. Standard errors in parentheses are clustered at the country level. * $p < 0.05$ ** $p < 0.01$ *** $p < 0.001$.

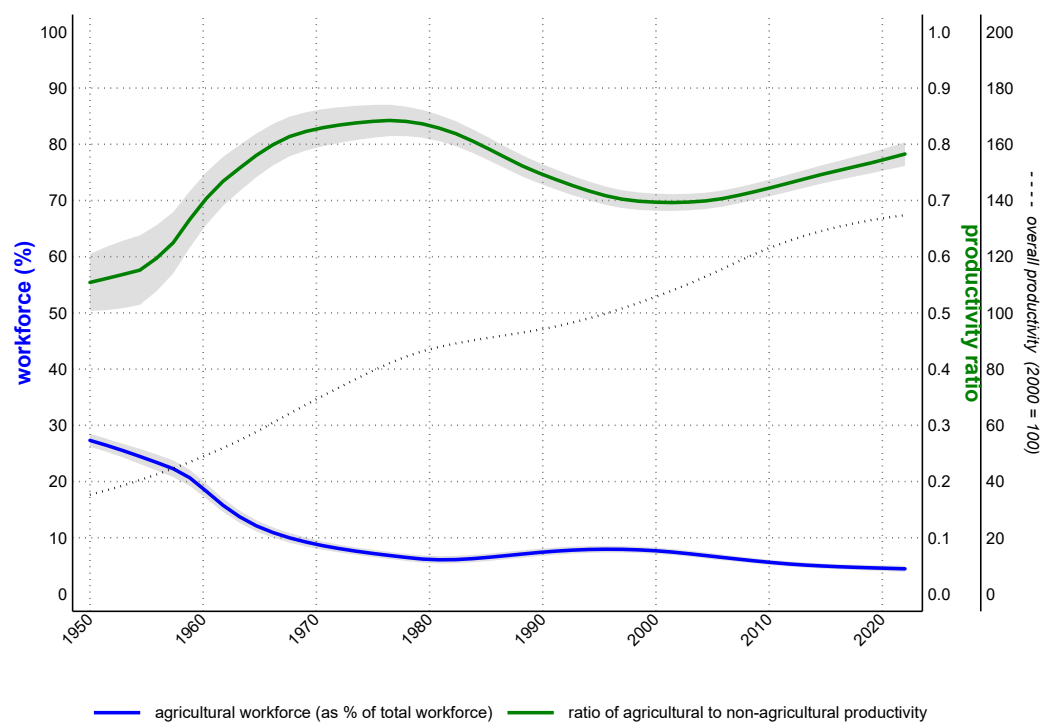
Table 2: Productivity growth and structural change, 1960–2021

region	period	N	productivity growth	of which	
				<i>within-sector</i>	<i>structural change</i>
Europe & Anglosphere	1960–69	1	52%	48%	4%
	1970–79	7	20%	19%	1%
	1980–89	10	18%	17%	1%
	1991–99	42	10%	9%	0%
	2000–10	46	12%	10%	2%
	2011–21	47	13%	12%	1%
Latin America & Caribbean	1960–69	8	29%	19%	10%
	1970–79	8	28%	17%	12%
	1980–89	8	-8%	-15%	7%
	1991–99	29	6%	3%	2%
	2000–10	31	10%	7%	3%
	2011–21	27	-5%	-6%	1%
Asia & Pacific	1960–69	4	34%	46%	-13%
	1970–79	7	22%	7%	15%
	1980–89	11	30%	18%	12%
	1991–99	32	16%	5%	11%
	2000–10	37	53%	38%	15%
	2011–21	30	45%	36%	10%
Middle East & North Africa	1960–69	2	24%	19%	5%
	1970–79	2	41%	29%	13%
	1980–89	2	27%	20%	7%
	1991–99	19	2%	-1%	3%
	2000–10	21	8%	3%	5%
	2011–21	20	8%	2%	5%
Sub-Saharan Africa	1960–69	5	17%	-9%	25%
	1970–79	18	9%	-2%	11%
	1980–89	18	-16%	-17%	1%
	1991–99	36	-4%	-7%	3%
	2000–10	40	29%	11%	17%
	2011–21	41	0%	-10%	10%

Note: this table presents the decomposition of productivity growth from equation (5). Overall productivity growth is defined as the cumulative change in value added per adult (in constant 2022 USD at PPP). The part explained by ‘within-sector’ productivity growth refers to average productivity growth in each of two sectors (holding the end year sectoral workforce shares constant), while ‘structural change’ refers to the part of productivity growth explained by labourers switching sectors (holding the beginning year sectoral productivity levels constant). Structural change, in this case, would usually reflect labourers moving from agricultural to non-agricultural employment (from the low- to the high-productivity sector). Results per region-decade are weighted by total value added (in constant 2022 USD at PPP), by construction. Only ‘panel’ countries—for which we observe data in both the beginning and end years of a given period—are included within each region.

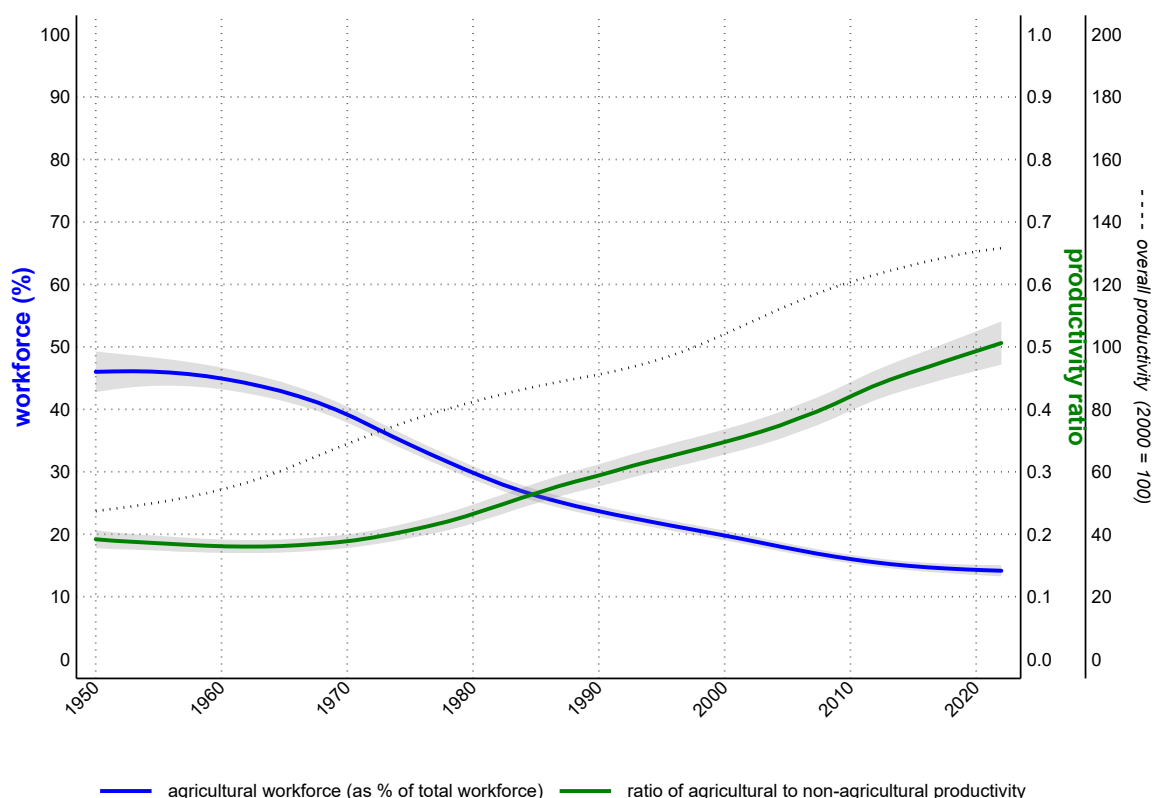
Appendix: additional figures and tables

Figure A1: Relative size and productivity of agricultural workforce —Europe & Anglosphere



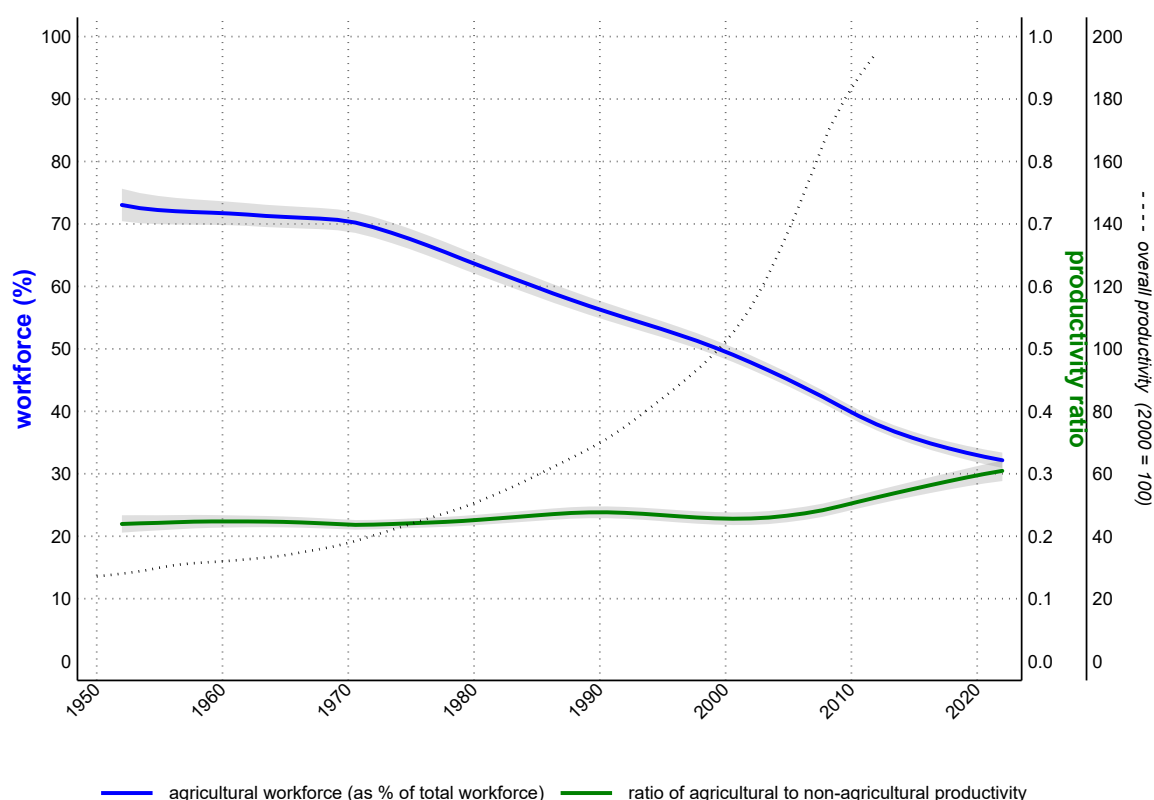
Note: lines represent local polynomial smoothing (with 95% confidence interval) on the size of the agricultural workforce (relative to total workforce) and the ratio of relative productivity in the agricultural sector to relative productivity in rest of the economy. Relative productivity is the ratio of the sector's share of total value added to its share of the workforce (such that economywide average productivity is always 1.0), so the ratio of agriculture's productivity to that the rest of the economy shows the agricultural sector's productivity gap. Total productivity (dashed line) is national income per capita (in constant 2022 USD at PPP), with each country's productivity level indexed to 100 at the turn of the century (average in the years 1999–2001). All estimates are weighted by size of national populations. The years 1970 and 1991 show significant structural breaks (sharp expansions) in the sample of countries for which we observe labour force and productivity data.

Figure A2: Relative size and productivity of agricultural workforce —Latin America & Caribbean



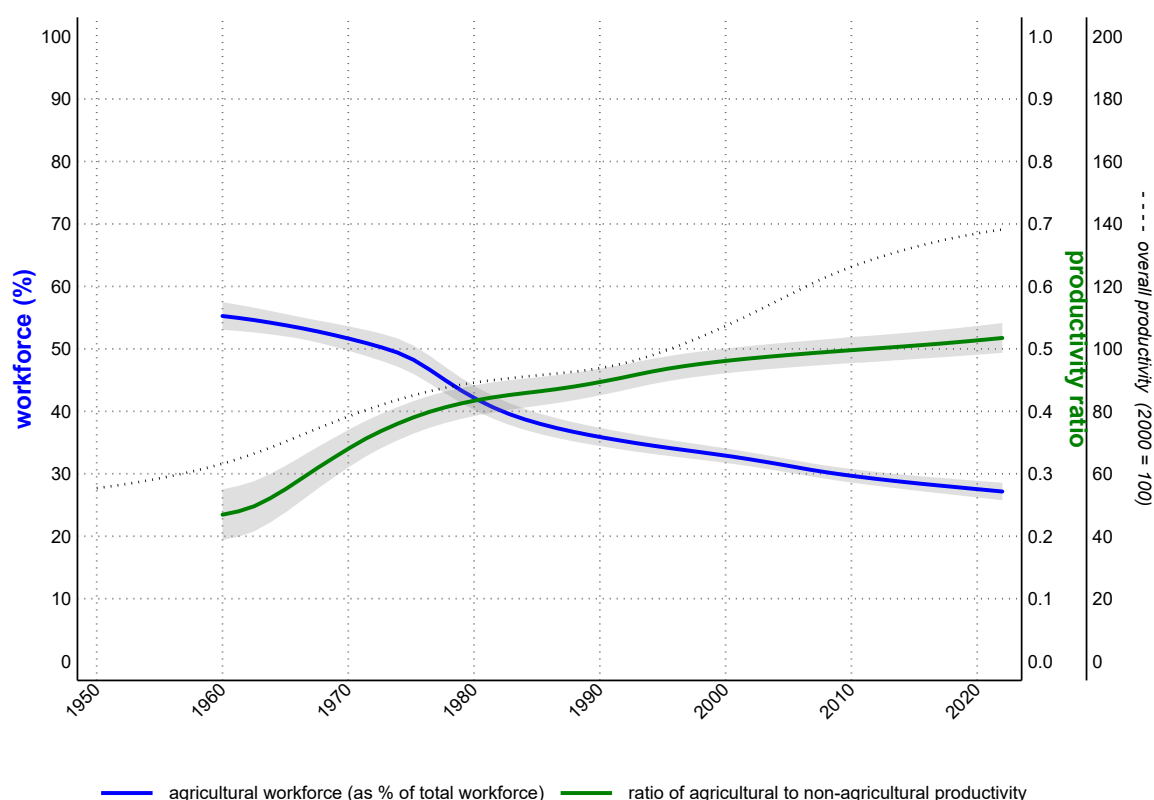
Note: lines represent local polynomial smoothing (with 95% confidence interval) on the size of the agricultural workforce (relative to total workforce) and the ratio of relative productivity in the agricultural sector to relative productivity in rest of the economy. Relative productivity is the ratio of the sector's share of total value added to its share of the workforce (such that economywide average productivity is always 1.0), so the ratio of agriculture's productivity to that the rest of the economy shows the agricultural sector's productivity gap. Total productivity (dashed line) is national income per capita (in constant 2022 USD at PPP), with each country's productivity level indexed to 100 at the turn of the century (average in the years 1999–2001). All estimates are weighted by size of national populations. The years 1970 and 1991 show significant structural breaks (sharp expansions) in the sample of countries for which we observe labour force and productivity data.

Figure A3: Relative size and productivity of agricultural workforce— Asia & Pacific



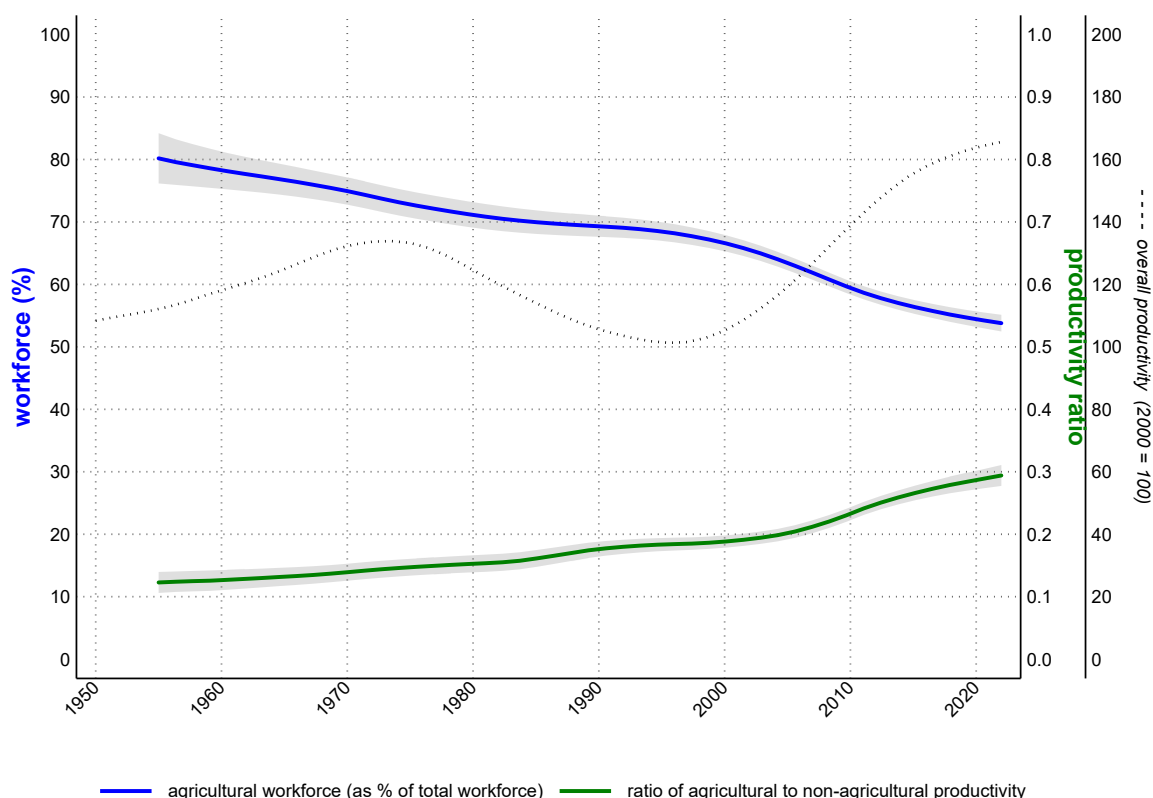
Note: lines represent local polynomial smoothing (with 95% confidence interval) on the size of the agricultural workforce (relative to total workforce) and the ratio of relative productivity in the agricultural sector to relative productivity in rest of the economy. Relative productivity is the ratio of the sector's share of total value added to its share of the workforce (such that economywide average productivity is always 1.0), so the ratio of agriculture's productivity to that the rest of the economy shows the agricultural sector's productivity gap. Total productivity (dashed line) is national income per capita (in constant 2022 USD at PPP), with each country's productivity level indexed to 100 at the turn of the century (average in the years 1999–2001). All estimates are weighted by size of national populations. The years 1970 and 1991 show significant structural breaks (sharp expansions) in the sample of countries for which we observe labour force and productivity data.

Figure A4: Relative size and productivity of agricultural workforce —Middle East & North Africa



Note: lines represent local polynomial smoothing (with 95% confidence interval) on the size of the agricultural workforce (relative to total workforce) and the ratio of relative productivity in the agricultural sector to relative productivity in rest of the economy. Relative productivity is the ratio of the sector's share of total value added to its share of the workforce (such that economywide average productivity is always 1.0), so the ratio of agriculture's productivity to that the rest of the economy shows the agricultural sector's productivity gap. Total productivity (dashed line) is national income per capita (in constant 2022 USD at PPP), with each country's productivity level indexed to 100 at the turn of the century (average in the years 1999–2001). All estimates are weighted by size of national populations. The years 1970 and 1991 show significant structural breaks (sharp expansions) in the sample of countries for which we observe labour force and productivity data.

Figure A5: Relative size and productivity of agricultural workforce —Sub-Saharan Africa



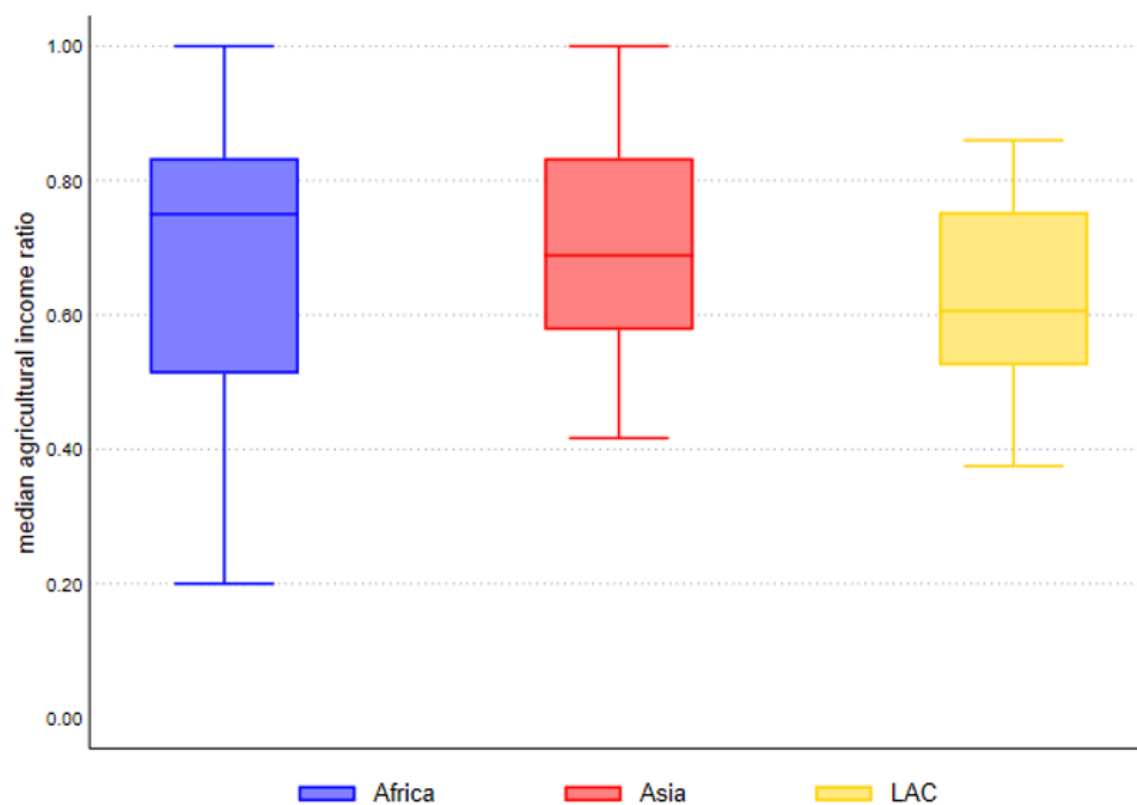
Note: lines represent local polynomial smoothing (with 95% confidence interval) on the size of the agricultural workforce (relative to total workforce) and the ratio of relative productivity in the agricultural sector to relative productivity in rest of the economy. Relative productivity is the ratio of the sector's share of total value added to its share of the workforce (such that economywide average productivity is always 1.0), so the ratio of agriculture's productivity to that the rest of the economy shows the agricultural sector's productivity gap. Total productivity (dashed line) is national income per capita (in constant 2022 USD at PPP), with each country's productivity level indexed to 100 at the turn of the century (average in the years 1999–2001). All estimates are weighted by size of national populations. The years 1970 and 1991 show significant structural breaks (sharp expansions) in the sample of countries for which we observe labour force and productivity data.

Figure A6: Inequality and the agricultural productivity gap, 2015



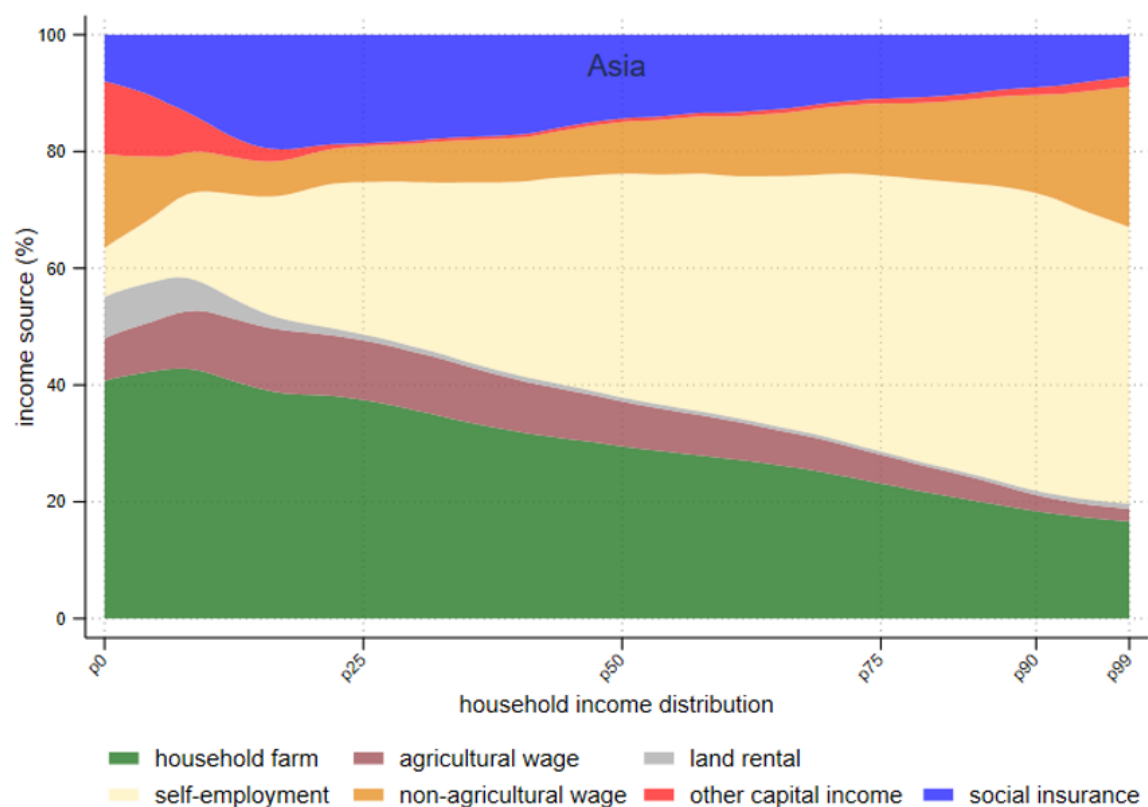
Note: inequality (y-axis) is expressed as the ratio of the average earnings of the top 10% of earners to the average earnings of the bottom 50% of the income distribution (World Inequality Database). The productivity gap (x-axis) is expressed as in equation (2) above. Data on labour productivity is taken for the year 2015. Dashed lines represent population-weighted linear trends.

Figure A7: Median income in agriculture vs. outside of agriculture



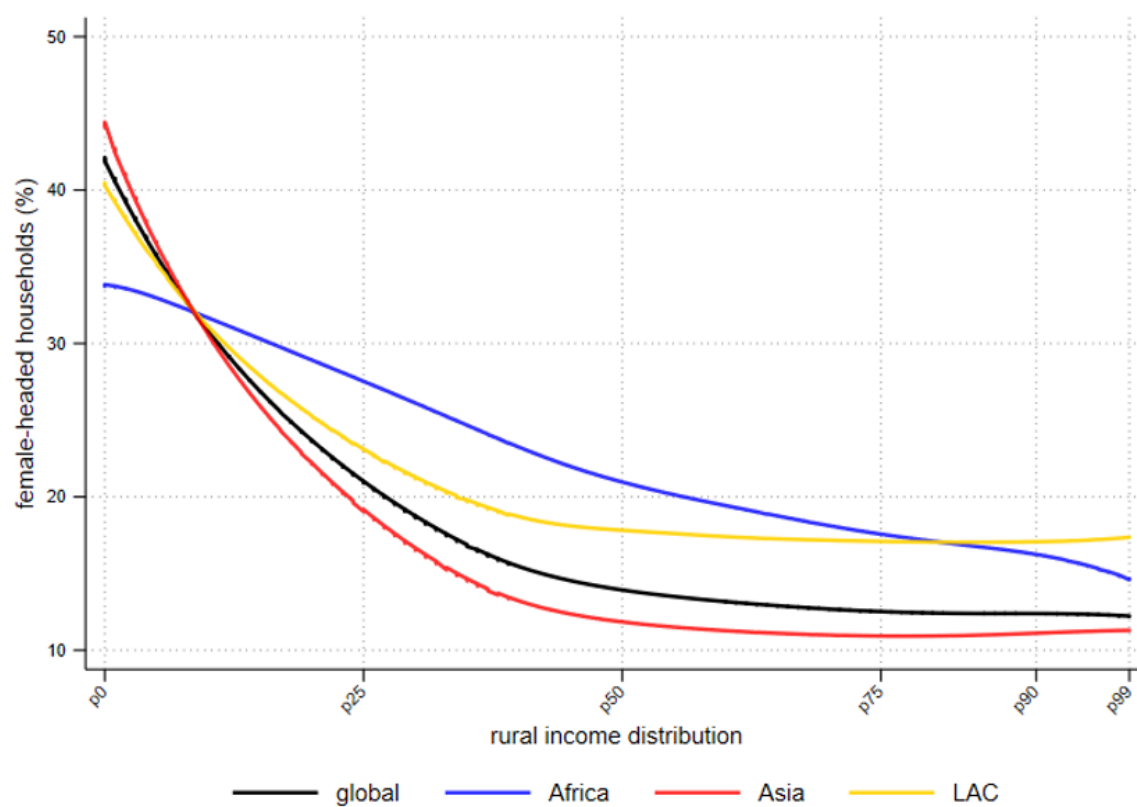
Note: this figure compares the median employee income in agriculture to the median employee income outside of agriculture, by region, using ILO labour force survey data.

Figure A8: Overall household income, by income components



Note: the figure depicts overall household income inequality, as the distribution of all income available in RuLIS, and ranked by total income, for all households (not only those whose income comes from agriculture).

Figure A9: Female-headed households in the rural income distribution



Note: the figure shows the propensity of agricultural households to be headed by women, in the rural household income distribution. Distributions are calculated at the country level, then averaged by region (weighted by population).