

LAND INEQUALITY IN THE DEVELOPING WORLD

LUIS BAULUZ
YAJNA GOVIND
FILIP NOVOKMET
DANIEL SANCHEZ-ORDONEZ

WORKING PAPER N°2020/10

UPDATED VERSION
JULY 2026

WORLD
INEQUALITY
..... LAB

Land Inequality in the Developing World*

Luis Bauluz[†]

Yajna Govind[‡]

Filip Novokmet[§]

Daniel Sánchez-Ordóñez[¶]

First version: June 2020

This version: July 2026

Abstract

Agricultural land is central to rural livelihoods in the developing world. Existing cross-country evidence on land inequality relies primarily on area-based census measures. We harmonize household surveys from seventeen developing countries and show that value, not area, is the relevant economic margin of land wealth. We find that using a value-based measure has important implications for the measurement of land inequality and poverty: it leads to higher estimated inequality on average, smaller cross-country differences, and reveals higher land-based poverty rates. Including landless agricultural households reinforces these patterns.

JEL classification: Q15, O1

Keywords: Land Ownership, Inequality, Distribution

*We thank Jacob Adenbaum, Miguel Almunia, Facundo Alvaredo, Leticia Arroyo-Abad, Luc Behaghel, Federico Curci, Denis Cogneau, Garance Genicot, María Hernández de Benito, Gabriel Leite, Thomas Piketty, José Vicente (Sevi) Rodríguez Mora, Pedro Salas Rojo, Diego Sánchez-Ancochea, Paolo Santini, Emiliano Travieso, and Juan Vargas, as well as participants at the Second World Inequality Conference, the ILC Uneven Ground Conference, and seminars at PSE, the WIL and CUNEF Universidad for their helpful comments. We are also grateful to Indu Chhatwani, Shalmali Ghaisas, Alberto López Yagüe, and Gustavo Sepúlveda for excellent research assistance. The authors acknowledge financial support from the International Land Coalition (ILC), the World Inequality Lab (WIL), and the Deutsches Institut für Entwicklungspolitik (DIE). Bauluz acknowledges financial support from the Spanish Ministry of Science and Innovation (PID2022-141254NA-I00) and Madrid's Autonomous Government (2023-T1/PH-HUM-29231). An earlier version of this paper circulated under the title "Global Land Inequality."

[†] CUNEF Universidad, and World Inequality Lab, luis.bauluz@cunef.edu.

[‡] Aix-Marseille School of Economics, Copenhagen Business School, and World Inequality Lab, yg.eco@cbs.dk

[§] University of Zagreb, and World Inequality Lab, fnovokmet@efzg.hr

[¶] Roma Tre University, and World Inequality Lab, d.sanchez.ordonez@psemail.eu

1 Introduction

Agricultural land is central to the economic lives of the world’s poor, with three out of four of the poorest billion individuals working in agriculture (FAO, 2016). For these households, land is not merely a factor of production but also their principal store of wealth, collateral, and a determinant of occupational choice and economic security (Binswanger and Deininger, 1997; Deininger and Feder, 2001). Its distribution has therefore long been seen as a fundamental driver of development, shaping institutions, human capital, political power, and the persistence of inequality (Engerman and Sokoloff, 2002; Banerjee and Iyer, 2005; Galor et al., 2009; Bartels et al., 2024). Hence, properly measuring land wealth distribution across households is central to the developing world.

Existing cross-country evidence on land inequality relies almost exclusively on agricultural censuses, which record the area of operational holdings and underpin much of the literature linking land inequality to growth, institutions, and development (Alesina and Rodrik, 1994; Deininger and Squire, 1998; Frankema, 2010; Baten and Hippe, 2018). Yet, censuses are ill-suited to measuring land as an economic asset: they record land under operation rather than land *ownership*, area rather than *value*, and exclude the *landless*, who account for about one-third of agricultural households. Furthermore, census data are aggregated and not linked to households, leaving unexplored how land relates to households’ economic outcomes and poverty.

We address these limitations by constructing a new harmonized database on land ownership, land value, income, wealth and credit at the household level, covering seventeen developing countries across Africa, Asia, and Latin America that together account for roughly two-thirds of the developing world by population and GDP. Drawing on nationally representative household surveys with detailed land modules, we shift the unit of analysis from farms to households and from land area to land value, and include agricultural households that own no land. It allows us to shed light on features of land inequality that census-based analyses cannot capture.

We begin by establishing that land *value*, rather than land *area*, is the economically relevant margin. Among landowning households, a 1 percent increase in land value per hectare is associated with a larger income gain than a 1 percent increase in land area (elasticities of 0.20 versus 0.12 in a joint specification). The same pattern holds more strongly for wealth, and to a lesser extent for credit. Value likely reflects differences in soil quality, irrigation, or proximity to markets that hectares do not, motivating the measurement of land inequality by value.

The central insight of the paper is that measuring land by value rather than area has implications for land inequality and poverty. First, value-based measures imply higher levels of land inequality on average but sharply compress its variation across countries. The average Gini rises from 0.65 for land area among owners to 0.68 for land value among owners and to 0.77 once landless agricultural households are included. At 0.77, the land Gini is substantially above the

corresponding average income Gini of 0.60 and close to the average total-wealth Gini of 0.73.¹ At the same time, the cross-country standard deviation of the land Gini falls by about 35 percent, from 0.132 for area to 0.085 for value including the landless, compared with 0.110 for income inequality. Land wealth is therefore highly concentrated within countries, but countries are more similar to one another in its concentration than area-based measures suggest.

An inverse value–size relationship helps explain why cross-country dispersion falls even as average inequality rises. Larger holdings have lower values per hectare in all seventeen countries, so valuing land produces the largest downward adjustments where land-area dispersion is greatest. Countries that appear most unequal in hectares therefore look less exceptional when land is measured by value. This pattern is only partly explained by the classic inverse farm size–productivity relationship (Sen, 1962; Berry and Cline, 1979; Vollrath, 2007). We provide evidence of the link between land-use structure and the large cross-country differences in area dispersion: animal-oriented production and greater reliance on pasture and grazing are associated with high area dispersion. Area-based country variations, frequently used in the literature (Alesina and Rodrik, 1994; Deininger and Squire, 1998; Baten and Hippe, 2018), therefore partly capture differences in land-use structure, rather than only differences in agrarian wealth concentration.

Second, measuring land by value also has implications for poverty measurements. We show that a value-based measure is better at identifying the poor than an area-based measure: poverty rate among landowning households at the bottom 20% of the land value distribution is 7.8 percentage points higher than among the bottom 20% of the land area distribution. We provide a measure of land-based asset poverty: the share of agricultural households with land wealth below US\$1,000 (PPP), roughly the scale of transfers shown to lift households out of poverty in “big-push” programs (Balboni et al., 2022; Banerjee et al., 2021; Blattman et al., 2020; Barker et al., 2024).

We show that such deprivation is widespread but highly uneven: the incidence is 13.1 percent of agricultural households in Vietnam and 15.3 percent in China, compared with 68.6 percent in Malawi. The aggregate shortfall below the threshold amounts to less than 5 percent of existing land value in most countries, and under 1 percent in several. High concentration is compatible with widespread deprivation at the bottom and a modest aggregate shortfall. The Gini, headcount, and shortfall jointly characterize concentration, incidence, and depth within a common land-wealth framework. Throughout the paper, we show that including the landless households reinforces these results.

We contribute to three strands of literature. First, we complement the literature measur-

¹The 0.73 figure is the unweighted average across the four surveys with direct household wealth modules. Matching WID.world estimates of overall wealth inequality to each survey year across all seventeen countries yields a comparable average of 0.76; see also Chancel et al., eds 2021 and UBS 2024.

ing land inequality across countries with agricultural census data (Deininger and Squire, 1998; Frankema, 2010; Lowder et al., 2025) by measuring land through ownership and value rather than the area of operational holdings, and by including landless households. Our estimates indicate that area-based measures of land concentration, widely used in work on long-run development (Alesina and Rodrik, 1994; Deininger and Squire, 1998; Galor et al., 2009; Ramcharan, 2010; Cinnirella and Hornung, 2016; Goñi, 2023), understate the level of land inequality and overstate its cross-country variation. Second, we complement studies of household wealth and its distribution in developing countries (Davies et al., 2011; Chancel et al., eds, 2021), which have lacked household-level data on land, the principal asset of rural households: our estimates show this asset to be as concentrated as household wealth itself. Lastly, we contribute to the literature on poverty traps and asset-transfer programs (Banerjee and Newman, 1993; Galor and Zeira, 1993; Balboni et al., 2022; Barker et al., 2024) by quantifying land poverty in the developing world.

The rest of the paper is structured as follows. Section 2 describes the data and concepts. Section 3 examines the relationship between landownership, land value, and household economic outcomes. Section 4 presents the implications of our new measures of land inequality and studies cross-country patterns of land-based asset poverty. Section 5 concludes.

2 Data and Concepts

This section presents the data and concepts underlying our analysis. We construct a harmonized household-level database covering land area, land value, income, wealth, and credit across seventeen developing countries. The micro-level structure of our database lets us trace how land value and land area relate to income, wealth, and access to credit (Section 3); construct measures of land wealth inequality based on ownership and value rather than area alone, and including the landless; and quantify land-based asset poverty (Section 4).

2.1 Data Sources and Database Construction

Our data consists of nationally representative household surveys with detailed modules on land ownership, agricultural activity, income, and borrowing. We harmonize information on land, loan, and wealth variables across surveys using consistent definitions. We rely on the most recent pre-2020 survey wave that has data on both land area and land value, as well as income.² The final sample, reported in Table 1, covers seventeen developing countries across Africa, Asia, and Latin America, including the largest developing economies and covering roughly two-thirds of the developing world by population or GDP. All estimates use survey sampling weights to ensure national

²Only two out of the seventeen countries have a survey wave during or post-COVID.

representativeness.

Table 1: Database Summary: Survey Sources and Key Characteristics

Country	Survey type	Year	Obs.	Agric.	Share	Variable availability				
				HH share (%)	landless (%)	Area	Value	Loans	Wealth	Income
South Asia										
India	National	2019	116,461	44.9	42.3	✓	✓		✓	
Bangladesh	National	2011	6,503	68.9	32.4	✓	✓	✓		✓
Pakistan	National	2010	16,345	38.7	32.3	✓	✓	✓		✓
Indonesia	National	2014	15,921	35.9	27.9	✓	✓	✓		✓
China–Vietnam										
China	National	2012	13,218	63.4	5.6	✓	✓			✓
Vietnam	National	2014	9,399	64.2	10.9	✓	✓	✓		✓
Latin America										
Ecuador	National	2014	28,970	28.2	35.7	✓	✓	✓		✓
Guatemala	LSMS	2000	7,276	47.5	44.6	✓	✓	✓		✓
Brazil	LSMS	1997	4,940	17.9	43.9	✓	✓	✓		✓
Panama	LSMS	2008	7,045	33.7	27.7	✓	✓			✓
Nicaragua	National	2014	6,851	29.6	39.7	✓	✓			✓
Sub-Saharan Africa										
Ethiopia	LSMS	2018	6,894	63.1	35.3	✓	✓	✓	✓	✓
Gambia	National	2015	13,282	36.5	32.7	✓	✓	✓		✓
Malawi	LSMS	2016	12,447	78.5	25.4	✓	✓	✓		✓
Subsample 2016	LSMS	2016	2,508	81.3	19.2	✓	✓		✓	
Niger	LSMS	2014	3,617	82.1	15.4	✓	✓			✓
Nigeria	LSMS	2015	4,697	54.3	22.1	✓	✓			✓
Tanzania	LSMS	2014	3,352	70.0	20.6	✓	✓	✓		✓
Subsample 2019	LSMS	2019	1,184	93.6	25.5	✓	✓		✓	
Total / Average			280,910	46.3	34.3					

Notes: *Agric. HH share* is the weighted share of agricultural households among all households. *Share landless* is the weighted share of landless households among agricultural households. Checkmarks indicate availability of the corresponding variable. No data cleaning is applied: no winsorization, no dropping of zero or missing values. Malawi subsample refers to the IHPS panel; Tanzania subsample refers to the NPS-SDD panel. Both are used for wealth analysis. The Living Standards Measurement Study (LSMS) is a nationally representative survey conducted by the World Bank. *National* refers to nationally representative household surveys conducted by the country's own statistical agency or affiliated research institutions.

We rely on a cross-sectional analysis as few countries have more than one wave with land value data, and land inequality tends to be highly persistent over time. As a robustness test, we identify and estimate Gini coefficients using additional survey waves and the Demographic and Health Surveys (DHS), which are nationally representative surveys that contain standardized questions on land ownership and plot size across several countries, but no information on value. These alternative estimates provide validity to our main estimate and confirm that land inequality is stable across waves, as shown in Appendix Figures A3 and A4.

2.2 Definitions and Measurement

Our unit of analysis is the *agricultural household*, defined to include both landowning and landless households. Any household that owns agricultural land is classified as a landowner. A household is classified as landless if it does not own agricultural land but reports at least one member employed in agriculture. This definition captures the full population whose livelihoods are directly linked to agriculture and provides an appropriate denominator for analyzing inequality in agrarian economies. In robustness checks, we consider more restrictive definitions of landlessness that only include households whose head reports agriculture as the primary activity or, when this information is unavailable, households in which more than half of the members work in agriculture. For some exercises, we further distinguish between all landless agricultural households and the subset of landless households that are income poor, defined as having per-capita daily income below US\$2.15 PPP.

A key challenge in measuring land ownership inequality across countries is defining ownership consistently in environments with heterogeneous property-rights institutions. Our working definition of ownership is a *proprietor* - a household holding use, control, and exclusion rights – who can also *bequeath*. Given the data limitations across seventeen countries, we rely on the most direct evidence each survey offers to determine land tenure: i) whether the household declares owning land when the question is directly asked in the survey, and otherwise, ii) whether a household reports having acquired a plot of land through transfer of rights, either purchase or inheritance. This classification indirectly proxies for the standard “bundle of rights” approach to land ownership in the development literature (Schlager and Ostrom 1992; Deininger 2003). Although imperfect, our measure of ownership is reliable as it aligns with estimates from the DHS, which measures ownership systematically across countries (Online Appendix Figure A3). Appendix A.1.1 discusses our ownership definition.

This broader definition accommodates contexts like China, where households can use the land, control its production, pass it on, while not holding sales rights. Our definition of ownership captures socially or legally recognized use rights that are de facto exclusive and heritable. Consistent with this definition, land owned by corporations or communal land that cannot be exclusively controlled or transferred by the household is excluded from our baseline ownership definition, in line with standard practice for measuring household wealth (Alvaredo et al. 2021).³ Although operation and ownership are distinct concepts, family farms operate roughly 75% of the world’s agricultural land (Lowder et al. 2016).⁴ Appendix Table A1 summarizes the ownership questions and

³For example, this definition excludes land “granted by village leaders”, a common land tenure set-up in Sub-Saharan countries.

⁴Available FAO data for seven countries in our sample confirm that about 90% of land is operated by households (Appendix A.2).

definitions used in each country. Our ownership definition is both robust and externally validated: results are unchanged when all potentially owned land— that is, all land held by the household other than rented, leased, or sharecropped-in land— is included (Appendix Table B9), and our harmonized measure closely replicates estimates based on the DHS, which elicits ownership directly and identically across countries (Figures A3).

For landowning households, we measure agricultural land along two dimensions. Land area is the sum of the areas of all plots owned by the household. Direct household reports of current market value are available in thirteen of the seventeen countries. The remaining four use alternative valuation methods: China provides a household-level valuation constructed by the survey; Ethiopia reports the household’s assessed annual rental value; and Brazil and Nicaragua impute land values as owned area times regional land prices, drawn from an external official source in Brazil and supplied within the survey in Nicaragua (Appendix A.1.2 and Table A2). We verify that our main regression and inequality results are robust to excluding all four countries (Appendix Table B11 and Appendix Figure C9). Measuring land in value rather than area terms is essential because land is heterogeneous: differences in soil quality, irrigation, location, and market access are reflected in prices but not in hectares. Self-reported current market values are standard for valuing real estate in household wealth surveys. The widely used SCF and HFCS, for example, apply this approach to residences, other real estate, and farmland (Board of Governors of the Federal Reserve System 2023; Household Finance and Consumption Network 2023).

Beyond land, we measure three household outcomes: income, wealth, and credit. Each is harmonized across surveys following the protocols in Appendices A.1.3–A.1.5. Total household income sums agricultural income (own-farm net income, farm wages, and land rents), non-agricultural income (business profits and wages), and capital and property income (interest, dividends, and non-agricultural rents). Our benchmark measure excludes transfers such as pensions, remittances, and government cash assistance, so that income reflects the household’s own earning capacity, the returns to its labor and assets, rather than redistributive flows that partially smooth income without altering underlying asset constraints. We hold transfers in a separate variable, and adding them back leaves the main estimates essentially unchanged (Appendix B.1.5, column 4).

Household wealth is measured in four surveys with a detailed asset module, and it is the current value of all assets owned, grouped into the dwelling, the agricultural plot(s), and other non-land assets such as financial holdings, livestock, and business capital (Appendix A.1.4). For the eleven surveys that record borrowing, household credit is *market credit*: the outstanding loan principal summed across all loans, excluding subsidized lending from state development banks, NGOs, and microfinance institutions (Appendix A.1.5). All monetary variables are expressed in 2018 USD PPP terms, and income and wealth are winsorized at the 99th percentile. Given that variable availability differs across surveys (Table 1), the income, wealth, and credit analyses draw

on sixteen, four, and eleven surveys, respectively. Differences in country composition across outcomes do not drive our results: the main estimates are unchanged when each analysis is restricted to common survey subsamples (Appendix B.1.7)

2.3 Inequality Measurement Concepts

Cross-country studies of land inequality rely on agricultural censuses, which report the size distribution of *operational holdings*—units of production under single management, irrespective of ownership (e.g., Deininger and Squire 1998; Vollrath 2007; Frankema 2010). Censuses measure area accurately across ownership types,⁵ but are ill-suited to measuring land as an economic asset. First, they record land operated rather than owned, which differs especially where land is leased or managed by tenants, leading to underestimation of land concentration when a single owner controls multiple operational holdings. Second, it measures area, ignoring differences in value. Area-based estimates overlook differences arising from soil quality, irrigation, and location, thereby understating inequality in land wealth. Finally, censuses exclude landless households, who account for 34% of agricultural households on average (Table 1). Since a substantial share of the agricultural population depends on agriculture for their livelihoods, without owning land, failing to account for them biases inequality measures downwards. Household surveys allow us to address each limitation, yielding three increasingly comprehensive measures: land area among landowners, land value among landowners, and land value among all agricultural households, including the landless.

2.4 Data Validation

The shift from censuses to surveys raises a natural question: do household surveys provide reliable measures of land inequality? We address this concern by benchmarking our baseline survey-based estimates against two independent sources in Appendix A.3: agricultural censuses and the Demographic and Health Surveys (DHS), which include standardized land modules across many countries but do not collect land value data. Land-area inequality from our household surveys aligns closely with DHS-based estimates across countries, lending credibility to our measure. Furthermore, survey-based Gini coefficients also correlate strongly with census-based measures ($R^2 = 0.94$). This relationship is not mechanical as censuses record operational holdings while surveys record ownership. This is a key finding, showing that censuses provide a reliable proxy for land-area ownership inequality. The correspondence weakens markedly when we compare census-based area inequality with survey-based value inequality ($R^2 = 0.56$), or when comparing within-survey area-based against value-based inequality estimates ($R^2 = 0.63$), suggesting that value captures di-

⁵Coverage typically extends beyond household farms to government and corporate land, though it varies across countries and over time (Lowder et al., 2016).

mensions of land heterogeneity, such as quality, irrigation, location, that area does not. If the goal is to measure land wealth distribution, then census-based estimates can be misleading.

3 Relevance of land value

This section documents how land ownership relates to household economic outcomes across developing countries. We use our harmonized household survey microdata to describe agricultural households, focusing on land as the central productive asset. We study three key dimensions of welfare and development potential—income, wealth, and credit—and analyze their relationship with land ownership. We estimate a pooled OLS regression with country-survey fixed effects, exploiting within-country variation across households:

$$Y_{ij} = \alpha + \beta_1 \text{LogLandArea}_{ij} + \beta_2 \text{LogLandValuePerHectare}_{ij} + \lambda_j + \epsilon_{ij} \quad (1)$$

where Y_{ij} denotes income, wealth, or credit of household i in country j , each measured in logs.⁶ For each outcome, we estimate three specifications: conditional on owning land, (i) the effect of log land area (ii) the effect of log land value; and (iii) the effect of land area and value per hectare jointly (eq. 1).

Table 2 reports the regression results. Land area alone carries a modest elasticity (column 1, 0.033) and explains almost no within-country variation in income (within- R^2 of 0.001). Adding the price per hectare in column (3) raises the within- R^2 around twentyfold to 0.034, with a price elasticity (0.202) exceeding the area elasticity (0.122), suggesting that land predicts income through its worth rather than its physical size. Consistent with this, land value alone attains nearly the same fit (column 2, within- R^2 of 0.029), thus capturing the more relevant dimension of land.⁷

Columns (4)–(6) turn to household wealth. We measure wealth as non-plot assets, so that the outcome does not mechanically include the land whose value we study. If anything, the results are stronger than for income: the within- R^2 rises from 0.002 with area alone to 0.127 once the price per hectare is included, and the value specification in column (5) again captures essentially all of this fit. Households with more valuable land, not more extensive land, hold more non-plot wealth. This pattern is stronger when including land wealth.⁸

Borrowings provide further evidence for the central role of land value, as it is a key constraint

⁶We address the treatment of zeros in Supplementary Appendix B.1.2.

⁷Country-specific results are reported in Supplementary Appendix B.1.4.

⁸This pattern holds for other definitions of wealth, also when we include land wealth, as shown in Appendix Table B14.

Table 2: Income, Non-Plot Wealth, and Loans Regressions

	<u>Income</u>			<u>Non-Plot Wealth</u>			<u>Loans</u>		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Log Land Area	0.033*	–	0.122***	0.050	–	0.297***	0.186***	–	0.285***
	(0.019)		(0.019)	(0.044)		(0.058)	(0.029)		(0.033)
Log Land Value	–	0.154***	–	–	0.378***	–	–	0.250***	–
		(0.016)			(0.072)			(0.028)	
Log Value/Area	–	–	0.202***	–	–	0.420***	–	–	0.213***
			(0.018)			(0.080)			(0.027)
Survey FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Clusters	441	441	441	174	174	174	295	295	295
Surveys	16	16	16	4	4	4	11	11	11
Sample	Owners	Owners	Owners	Owners	Owners	Owners	Borrowers	Borrowers	Borrowers
Observations	50,031	50,031	50,031	34,913	34,913	34,913	8,303	8,303	8,303
Adj. R^2	0.568	0.580	0.582	0.316	0.401	0.408	0.627	0.638	0.640
Within R^2	0.0014	0.0285	0.0339	0.0018	0.1266	0.1360	0.0326	0.0615	0.0650
Mean dep. var.	6.089	6.089	6.089	9.593	9.593	9.593	5.940	5.940	5.940

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Each outcome is shown under three specifications: on log land area alone, on log land value alone, and on log land area jointly with log value per hectare ($\ln(\text{Value}/\text{Area}) = \ln V - \ln A$). In the joint specification, the Log Land Area coefficient is the effect of area holding price per hectare fixed, and Log Value/Area is the price-per-hectare effect. Income columns pool the income surveys (owners); wealth columns use the four wealth surveys with Non-Plot Assets (dwelling + other, excluding the plot) as the outcome; loan columns use income surveys with reported loans (owner-borrowers). All regressions include survey fixed effects and probability weights renormalized so that each survey carries equal total weight; standard errors clustered at the survey-strata level.

on investment and income growth in rural economies, particularly where land serves as collateral (e.g., [Besley 1995](#); [Carter and Olinto 2003](#); [Deininger and Feder 2009](#)). Conditional on borrowing, both land area and land value positively predict loan size. Value again outperforms area alone, albeit to a lesser extent than for income and wealth: the within- R^2 nearly doubles when the price per hectare is added. This result is complementary to the wealth findings, since a way to accumulate wealth is through collateral-based credit for investment.

Across all three outcomes, a key finding emerges: the *value* of land, beyond *area*, is the more economically relevant margin. These results have direct implications for measuring land inequality: if land value has greater predictive power, then the distribution of land value is the object that matters for understanding economic inequality in rural economies.

The results are robust to multiple specification checks as reported in Appendix B. We show that results hold under population-weighting, alternative income transformations and consumption measures, that they are not driven by cross-country composition effects, and remain robust to including non-owned land, including transfers in income, and varying winsorization thresholds. Furthermore, income non-response is largely unrelated to land area and value. Results are similarly

robust across alternative wealth definitions.

4 Land Wealth Inequality

We construct consistent estimates of land inequality using value-based measures and incorporating landless households. We start by documenting new stylized facts on land inequality across developing countries, explaining why cross-country differences in land inequality depend on how land is measured, and finally quantifying land wealth poverty at the bottom of the distribution.

4.1 From Land Area to Land Wealth

We construct three increasingly comprehensive measures of land inequality: (i) land area among landowners, (ii) land value among landowners, and (iii) land value including landless agricultural households. These measures move from a purely physical concept of land to a value-based measure of land wealth, and finally to a measure that also incorporates the landless.

Figure 1, Panel A reports country-level Gini coefficients under each concept. The solid blue lines show the change in inequality when the focus shifts from land area to land value among landowners. The dashed red lines show the additional effect of including landless agricultural households, who are assigned zero land value; estimates including only income-poor landless households are similar (Appendix Figure C8).


The main pattern is the relatively high concentration of land wealth. Averaging across countries, Panel B shows that the Gini coefficient is 0.65 for land area among landowners and 0.68 for land value among landowners. Once landless households are included, the Gini rises to 0.77, compared with an average income Gini of 0.60 across the sixteen surveys with income data.⁹ This level is comparable to reported estimates of overall wealth inequality in developing countries.¹⁰ Agricultural land is thus much more unequally distributed than area-based measures suggest and, once measured as wealth and inclusive of the landless, is about as concentrated as overall household wealth.

The transition from area to value varies across countries. In several South Asian and Sub-Saharan African countries, including Bangladesh, Pakistan, Ethiopia, and Malawi, the Gini is higher for land value than for land area, indicating that price dispersion amplifies inequality. By contrast, in several Latin American countries, as well as in China, Vietnam, and Indonesia, the


⁹Whiskers in Panels B and C are 95 percent survey-design bootstrap confidence intervals, capturing sampling uncertainty within the underlying surveys.

¹⁰See, for example, the World Inequality Report 2022 (Chancel et al., eds 2021) and the UBS Global Wealth Report 2024 (UBS 2024). Using the four surveys where wealth can be measured directly, the average Gini of total household wealth is 0.73, remarkably close to the land-value Gini of 0.77; averaging WID.world overall-wealth Gini estimates matched to each survey year across all seventeen countries gives 0.76.


Figure 1: Land Inequality Across Measures



(a) Gini index across measurement concepts by country



(b) Cross-country mean Gini by concept



(c) Cross-country standard deviation of the Gini by concept

Notes: Panel (a) reports country-level Gini coefficients for land inequality under a sequence of measurement concepts. The blue segment shows the change in the Gini when moving from land *area* among landowners to land *value* among landowners. The red segment shows the change when moving from land value among landowners to land value including landless households (i.e., assigning zero land value to the landless). Panels (b) and (c) report the unweighted cross-country mean and standard deviation. Land statistics use all seventeen surveys; income statistics use sixteen, excluding India. Whiskers show 95 percent confidence intervals; see Appendix C.1.

shift is small or negative. In these settings, large holdings tend to have lower values per hectare, consistent with extensive holdings being located on lower-value land, a mechanism we examine in Section 4.2.

Assigning zero land value to landless households mechanically raises the Gini in every country, but the magnitude depends jointly on the landless share and inequality among landowners. All else equal, greater landlessness raises the adjustment, while an already high owners-only Gini leaves less room for it to rise. This heterogeneity underscores that analyses restricted to landowners overlook an important source of inequality: the divide between agricultural households that own land and those that own none.

Beyond levels, the cross-country dispersion of inequality also changes markedly with the concept used. Panel C shows that the standard deviation of the Gini falls from 0.132 for land area among owners to 0.092 for land value among owners and to 0.085 when landless households are included, a 35.8 percent reduction (95 percent design-bootstrap confidence interval: 30.0–42.3 percent). The corresponding income standard deviation, available for sixteen countries, is 0.110. The mean and dispersion results are virtually unchanged when India is excluded and are robust to trimming 0.5 percent from each tail or omitting any one country (Appendix Table C15).

This compression matters because a large literature uses cross-country variation in land-area inequality to study growth, institutions, and human-capital formation (Alesina and Rodrik, 1994; Birdsall and Londoño, 1997; Deininger and Squire, 1998; Deininger and Olinto, 2000; Frankema, 2010; Baten and Juif, 2014; Baten and Hippe, 2018). Under the economically relevant measure, roughly one-third of the cross-country dispersion disappears: countries look much more similar in land wealth inequality than in land-area inequality, and even more similar than they do in income inequality.¹¹

4.2 The Inverse Value-Size Relationship

We next explain why cross-country differences in land inequality shrink when land is measured by value rather than area. We focus on landowners because the landless contribute little to cross-country dispersion: the compression from area to value operates almost entirely among owners.

The key fact is simple: larger landholdings tend to have lower value per hectare, so a hectare held in a large estate does not translate into the same land wealth as a hectare in a smaller, higher-value plot. Because the Gini cannot be decomposed into area and price components, we use log variances to study the mechanism; Appendix Figure C14 shows that the same cross-country pattern holds for the Gini.

Panel (a) of Figure 2 shows the relationship of interest. The area-to-value adjustment, $\text{Var}(\log V) - \text{Var}(\log A)$, is more negative where land-area inequality, $\text{Var}(\log A)$, is highest. Countries that look most unequal in hectares therefore look less exceptional once land is valued. The reason follows from an identity. Let $V = A \times P$, where V is land value, A is land area, and P

¹¹Restricting the landless to households below the US\$2.15 PPP poverty line yields similar results: the mean Gini falls only from 0.77 to 0.74, while its cross-country standard deviation is virtually unchanged (Appendix Figure C8).

is value per hectare. Then:

$$\text{Var}(\log V) - \text{Var}(\log A) = \text{Var}(\log P) + 2 \text{Cov}(\log A, \log P)$$

The first term, $\text{Var}(\log P)$, captures dispersion in value per hectare, which raises value inequality relative to area inequality. The second term captures the size–value relationship. If larger holdings have lower value per hectare, this covariance is negative and compresses value inequality.


Panel (b) decomposes the change separately for “descending” countries, where $\text{Var}(\log V) - \text{Var}(\log A) < 0$, and “ascending” countries, where it is positive. In descending countries, the negative size–value covariance more than offsets dispersion in value per hectare; in the ascending countries, it does not. Since $\text{Cov}(\log A, \log P) = \beta \text{Var}(\log A)$, where β is the elasticity of value per hectare with respect to size, the covariance can be more negative either because holdings are more dispersed or because the inverse value–size gradient is steeper. Appendix Figure C11b and Table C16 show that the first margin is more important. The elasticity is negative across nearly all countries, but not systematically more negative where land-area dispersion is higher; the difference between the two groups is driven mainly by $\text{Var}(\log A)$.

What makes area inequality high in some countries? Our harmonized household-level data allow us to relate landholding dispersion to within-country differences in production patterns, which aggregate census data cannot. We classify landowners by the share of crop income in total farm income, using the 14 countries where crop and animal income can be separated and restricting the sample to households with positive crop or animal income. Appendix C.6 reports the group definitions and sample shares.

Panel (c) shows that squared deviations of log land area from country means rise monotonically with animal orientation, and Appendix Figure C13, Panel (a), shows the same ordering for within-group variances. A continuous specification provides complementary evidence: crop share is strongly negatively associated with squared deviations from the country mean but essentially unrelated to mean log holding size, and the association is concentrated in the upper tail rather than around the median (Appendix Table C18).

Taken together, these results are consistent with a land-use channel: animal-oriented production is associated with a thicker upper tail of landholdings and greater $\text{Var}(\log A)$. The pattern appears to reflect differences in holding-size dispersion rather than systematically different value gradients, since value–size elasticities are negative and similar across crop-income groups (Appendix Figure C13, Panel (b)). A related pattern appears in country-level FAO data matched to survey years: countries with lower arable shares of agricultural land—more oriented toward pasture and grazing—have higher land-area dispersion and a more negative size–value covariance (Appendix Figure C12).

Figure 2: **The inverse value–size relationship and land inequality**



Notes: Panel (a) plots $\text{Var}(\log V) - \text{Var}(\log A)$, against land-area inequality $\text{Var}(\log A)$, where V denotes land value and A land area. Panel (b) decomposes this change into $\text{Var}(\log P) + 2 \text{Cov}(\log P, \log A)$, separately for countries where the change is negative and positive. Panel (c) plots coefficients from a regression of the squared deviation of log area from the country mean on crop-income group indicators and country fixed effects; pure-crop households are the omitted group. Standard errors are clustered by country. The Panel (c) sample covers landowners with positive crop or animal income in the 14 countries with crop- and animal-income data. The result is robust to replacing the group indicators with a continuous measure of crop share (Appendix Table C18).

This inverse value–size relationship is related to, but not fully explained by, the classic inverse farm size–productivity relationship (e.g., [Vollrath, 2007](#)). Land values reflect current productivity but also capitalize location, tenure security, and expected future returns. Controlling for realized farm income per hectare makes the value–size elasticity less negative, but it remains neg-

ative in most countries (Appendix Table C17), suggesting that productivity explains only part of the association.

Taken together, these results do not imply that area-based measures are uninformative. They show, however, that land-area inequality is not a pure measure of agrarian wealth concentration: part of its cross-country variation appears to reflect production structure and extensive land use. With seventeen countries, we cannot re-estimate the cross-country relationships in the literature using value-based measures, but the evidence calls for caution in interpreting area-based country rankings as rankings of land wealth inequality.

4.3 Land Value at the Bottom of the Distribution

The same shift from land area to land value also matters for poverty, another central dimension of economic performance in developing countries, where land is the main productive asset of rural households.¹² Hectares do not capture differences in land quality, location, or productive potential, and may therefore misidentify households with little land wealth. We first compare area- and value-based rankings at the bottom without an absolute threshold. We then follow the logic of consumption-poverty measurement, using a common land-value benchmark to quantify the incidence and depth of land poverty.


Panel (a) of Figure 3 compares income poverty among landowners in the bottom 20 percent of the land distribution when households are ranked within each survey by value and by area. The two groups have exactly the same weighted size, and income plays no role in their construction. The income-poverty rate is 59.5 percent among the bottom quintile by value, compared with 51.7 percent among the bottom quintile by area, a difference of 7.8 percentage points. Moreover, only 46.9 percent of the two groups overlap; among households selected by only one measure, the poverty-rate gap is 15.8 percentage points. Value therefore does not simply relabel the same lower tail: it identifies a substantially different group that is poorer in income terms.

To translate this ranking result into an absolute measure, we set a benchmark at US\$1,000 PPP per household, near the midpoint of the one-off asset transfers studied in “big-push” programs, which range from roughly US\$500 to US\$1,400 PPP (Balboni et al. 2022; Banerjee et al. 2021; Blattman et al. 2020; Barker et al. 2024). We treat this as an economically meaningful order of magnitude, not a structural threshold, and report the same exercise for US\$500 and US\$1,500 PPP in Appendix C.10.¹³ Because it is defined in value rather than area, the benchmark does not impose a common physical-size threshold across heterogeneous land markets, unlike the widely


¹²Land accounts for roughly 70–80 percent of the productive assets of agricultural households in our sample (Online Appendix Figure C15, Panel B).

¹³For Ethiopia, we capitalize reported annual rental values using a benchmark factor of 10; Online Appendix Table C21 reports sensitivity to factors of 20 and 6.67. This constant rescaling leaves household rankings, Gini coefficients, and the regressions in Section 3 unchanged.

Figure 3: Land poverty in value terms



(a) Income poverty at the bottom: value versus area



(b) Incidence and value shortfall, by country

Notes: Panel (a) reports income-poverty rates among the bottom 20 percent of landowners ranked within each survey by land value and land area. The final two bars restrict attention to the non-overlapping portions of these groups: households selected by value but not area, and vice versa. Surveys receive equal weight. Brackets report paired differences and 95 percent confidence intervals based on survey-level estimates. Panel (b) reports, for the US\$1,000 PPP benchmark: (i) the share of agricultural households below the benchmark and (ii) the aggregate land value needed to bring all land-poor households up to it, expressed as a share of aggregate land value (“shortfall”). Landless households are assigned zero land value and therefore fall below the benchmark. Bars are stacked by household type. Income poverty is defined as per-capita daily income below US\$2.15 PPP. Panel (a) excludes India because income is unavailable; Panel (b) covers all seventeen countries. For Ethiopia, Panel (b) converts reported annual rental values into capitalized land values using the benchmark factor of 10.

used two-hectare cutoff for classifying small holdings (Lowder et al. 2016; Lowder et al. 2025). It also selects landowners with a substantially higher income-poverty rate.¹⁴

We classify as land poor all agricultural households whose land wealth falls below the benchmark, assigning zero land value to the landless—a purely asset-based criterion applied identically to owners and non-owners. Online Appendix Figure C19 reports an alternative that counts the landless as land poor only when also income poor. Although defined solely on land wealth, the resulting groups differ systematically in income poverty: the rate is highest among below-benchmark owners, lower among the landless, and lowest among above-benchmark owners (Online Appendix Figure C16).

Panel (b) applies this classification across countries and reports two measures analogous to the incidence and depth of consumption poverty. The headcount is the share of agricultural households below the benchmark, split between below-benchmark owners and landless households. The aggregate land-value shortfall is the sum of the gaps between the benchmark and observed land value among land-poor households, expressed as a share of aggregate land value. This is an accounting measure, not a policy simulation.

The incidence of land poverty varies substantially, from 13.1 percent in Vietnam and 15.3 percent in China to 68.6 percent in Malawi. Landless households account for most cases in most countries, while below-benchmark owners are especially important in China, Malawi and Tanzania. Relative inequality and absolute land poverty capture distinct features of the land-wealth distribution: similar land-value Ginis coexist with very different headcounts. Their association is modestly positive under the baseline ($R^2 = 0.18$), partly because landlessness raises both measures (Online Appendix Figure C17).

The shortfall gives a complementary picture. In most countries it is below 5 percent of aggregate land value, and in several below 1 percent, though it is larger in some low-income agrarian economies. A large mass of households at the bottom can therefore coexist with a modest aggregate shortfall when land wealth is concentrated far up the distribution.

Together, the value-based Gini, the headcount, and the shortfall characterize the same distribution from three angles. The Gini summarizes overall concentration; the headcount shows how many agricultural households hold very little land wealth, and whether deprivation falls on low-value owners or the landless; and the shortfall measures its aggregate depth relative to the existing stock. These statistics show that measuring land by value changes both who is identified as land poor and how land poverty is quantified.

¹⁴Among landowners selected by each rule, the adjusted income-poverty rate is 62.9 percent under the US\$1,000 benchmark and 54.3 percent under the two-hectare cutoff. The difference is 8.5 percentage points (95 percent confidence interval: 3.6–13.4; Online Appendix Table C20).

5 Conclusion

Using harmonized surveys from seventeen developing countries, this paper shows that measuring land by value rather than area changes what we learn about land inequality and poverty. Land value is more closely related to household economic outcomes than land area, and value-based measures reveal highly concentrated land wealth—comparable to overall household wealth—while substantially reducing apparent cross-country differences. An inverse value–size relationship helps explain this compression, while suggestive evidence indicates that area-based rankings partly reflect land-use structure. At the bottom of the distribution, value-based rankings identify poorer households than area-based rankings, and the headcount and shortfall quantify the incidence and depth of land poverty. Including landless agricultural households raises both measured inequality and land-poverty incidence.

More broadly, because land area and land wealth are conceptually and empirically distinct, analyses linking land inequality to development are incomplete without accounting for land wealth. This distinction also matters for targeting land, credit, and rural-development programs. Extending comparable household-level data across countries and over time would allow future work to revisit these links beyond hectare-based measures.

References

- Alesina, Alberto and Dani Rodrik**, “Distributive politics and economic growth,” *The quarterly journal of economics*, 1994, 109 (2), 465–490.
- Alvaredo, Facundo, Anthony B. Atkinson, Luis Bauluz, Thomas Blanchet, Lucas Chancel, Matthew Fisher-Post, Ignacio Flores, Bertrand Garbinti, Jonathan Goupille-Lebret, Clara Martínez-Toledano, Marc Morgan, Theresa Neef, Thomas Piketty, Anne-Sophie Robiliard, Emmanuel Saez, Li Yang, and Gabriel Zucman**, “Distributional National Accounts (DINA) guidelines: Methods and Concepts used in the World Inequality Database,” *WID Working Paper 2016/2*, 2021.
- Balboni, Clare, Oriana Bandiera, Robin Burgess, Maitreesh Ghatak, and Anton Heil**, “Why do people stay poor?,” *The Quarterly Journal of Economics*, 2022, 137 (2), 785–844.
- Banerjee, Abhijit and Lakshmi Iyer**, “History, institutions, and economic performance: The legacy of colonial land tenure systems in India,” *American economic review*, 2005, 95 (4), 1190–1213.
- , **Esther Duflo, and Garima Sharma**, “Long-term effects of the targeting the ultra poor program,” *American Economic Review: Insights*, 2021, 3 (4), 471–486.
- Banerjee, Abhijit V. and Andrew F. Newman**, “Occupational Choice and the Process of Development,” *Journal of Political Economy*, 1993, 101 (2), 274–298.
- Barker, Nathan, Dean Karlan, Christopher Udry, and Kelsey Wright**, “The Fading Treatment Effects of a Multifaceted Asset-Transfer Program in Ethiopia,” *American Economic Review: Insights*, 2024, 6 (2), 277–294.
- Bartels, Charlotte, Simon Jäger, and Natalie Obergruber**, “Long-term effects of equal sharing: Evidence from inheritance rules for land,” *The Economic Journal*, 2024, 134 (664), 3137–3172.
- Baten, Joerg and Dácil Juif**, “A Story of Large Landowners and Math Skills: Inequality and Human Capital Formation in Long-Run Development, 1820–2000,” *Journal of Comparative Economics*, 2014, 42 (2), 375–401.
- **and Ralph Hippe**, “Geography, Land Inequality and Regional Numeracy in Europe in Historical Perspective,” *Journal of Economic Growth*, 2018, 23 (1), 79–109.
- Berry, Albert and William Cline**, *Agrarian Structure and Productivity in Developing Countries*, John Hopkins University Press, 1979.
- Besley, Timothy**, “Property Rights and Investment Incentives: Theory and Evidence from Ghana,” *Journal of Political Economy*, 1995, 103 (5), 903–937.
- Binswanger, Hans P and Klaus Deininger**, “Explaining agricultural and agrarian policies in developing countries,” *Journal of economic literature*, 1997, 35 (4), 1958–2005.
- Birdsall, Nancy and Juan Luis Londoño**, “Asset inequality matters: an assessment of the World Bank’s approach to poverty reduction,” *The American Economic Review*, 1997, 87 (2), 32–37.

- Blattman, Christopher, Nathan Fiala, and Sebastian Martinez**, “The long-term impacts of grants on poverty: Nine-year evidence from Uganda’s youth opportunities program,” *American Economic Review: Insights*, 2020, 2 (3), 287–304.
- Board of Governors of the Federal Reserve System**, “Codebook for the 2022 Survey of Consumer Finances,” Survey documentation, Board of Governors of the Federal Reserve System October 2023.
- Carter, Michael R. and Pedro Olinto**, “Getting Institutions “Right” for Whom? Credit Constraints and the Impact of Property Rights on the Quantity and Composition of Investment,” *American Journal of Agricultural Economics*, 2003, 85 (1), 173–186.
- Chancel, Lucas, Thomas Piketty, Emmanuel Saez, and Gabriel Zucman, eds**, *World Inequality Report 2022* 2021. Report on global income and wealth inequality based on the World Inequality Database (WID.world).
- Cinnirella, Francesco and Erik Hornung**, “Landownership Concentration and the Expansion of Education,” *Journal of Development Economics*, 2016, 121, 135–152.
- Davies, James B, Susanna Sandström, Anthony Shorrocks, and Edward N Wolff**, “The level and distribution of global household wealth,” *The Economic Journal*, 2011, 121 (551), 223–254.
- Deininger, Klaus**, *Land policies for growth and poverty reduction*, World Bank Publications, 2003.
- **and Gershon Feder**, “Land Institutions and Land Markets,” in Bruce L. Gardner and Gordon C. Rausser, eds., *Handbook of Agricultural Economics*, Vol. 1, Elsevier, 2001, pp. 288–331.
- **and —**, “Land registration, governance, and development: Evidence and implications for policy,” *The World Bank Research Observer*, 2009, 24 (2), 233–266.
- **and Lyn Squire**, “New ways of looking at old issues: inequality and growth,” *Journal of development economics*, 1998, 57 (2), 259–287.
- **and Pedro Olinto**, “Asset distribution,” *Inequality, and Growth “*, World Bank Policy Research WP, 2000, (2375).
- Engerman, Stanley L. and Kenneth L. Sokoloff**, “Factor Endowments, Inequality, and Paths of Development Among New World Economies,” *Economía*, 2002, 3 (1), 41–109.
- FAO**, *The State of Food and Agriculture: Climate Change, Agriculture and Food Security*, Rome, Italy: Food and Agriculture Organization of the United Nations, 2016.
- Food and Agriculture Organization of the United Nations**, *A System of Integrated Agricultural Censuses and Surveys, Volume 1: World Programme for the Census of Agriculture 2010* FAO Statistical Development Series, Rome: Food and Agriculture Organization of the United Nations, 2005.
- Frankema, Ewout**, “The colonial roots of land inequality: geography, factor endowments, or institutions?,” *The Economic History Review*, 2010, 63 (2), 418–451.

- Galor, Oded and Joseph Zeira**, “Income distribution and macroeconomics,” *The review of economic studies*, 1993, 60 (1), 35–52.
- , **Omer Moav, and Dietrich Vollrath**, “Inequality in landownership, the emergence of human-capital promoting institutions, and the great divergence,” *The Review of economic studies*, 2009, 76 (1), 143–179.
- Goñi, Marc**, “Landed Elites and Education Provision in England: Evidence from School Boards, 1871–99,” *Journal of Economic Growth*, 2023, 28 (1).
- Household Finance and Consumption Network**, “HFCS Blueprint Questionnaire: 2021 Wave,” Survey documentation, European Central Bank July 2023.
- Lowder, Sarah, Jakob Skoet, and Terri Raney**, “The number, size, and distribution of farms, smallholder farms, and family farms worldwide,” *World Development*, 2016, 87, 16–29.
- Lowder, Sarah K, Garima Bhalla, and Benjamin Davis**, “Decreasing farm sizes and the viability of smallholder farmers: Implications for resilient and inclusive rural transformation,” *Global Food Security*, 2025, 45, 100854.
- Ramcharan, Rodney**, “Inequality and Redistribution: Evidence from U.S. Counties and States, 1890–1930,” *Review of Economics and Statistics*, 2010, 92 (4), 729–744.
- Ravallion, Martin and Dominique Van de Walle**, *Land in transition: Reform and poverty in rural Vietnam*, World Bank Publications, 2008.
- Schlager, Edella and Elinor Ostrom**, “Property-Rights Regimes and Natural Resources: A Conceptual Analysis,” *Land Economics*, 1992, 68 (3), 249–262.
- Sen, Amartya Kumar**, “An aspect of Indian agriculture.,” 1962.
- UBS**, “Global Wealth Report 2024,” Report, UBS Group AG 2024. Accessed 2026-02-18.
- Vendryes, Thomas**, “Land rights in rural China since 1978. Reforms, successes, and shortcomings,” *China Perspectives*, 2010, 2010 (2010/4).
- Vollrath, Dietrich**, “Land distribution and international agricultural productivity,” *American Journal of Agricultural Economics*, 2007, 89 (1), 202–216.

Online Appendix for

Land Inequality in the Developing World

Luis Bauluz

CUNEF Universidad

World Inequality Lab

Yajna Govind

AMSE, Aix-Marseille University,

Copenhagen Business School, WIL

Filip Novokmet

University of Zagreb

World Inequality Lab

Daniel Sánchez-Ordóñez

Roma Tre University

World Inequality Lab

Outline

- Appendix [A](#) presents additional information on variable definitions and comparisons with alternative data sources, supporting Section [2](#).
- Appendix [B](#) presents various robustness checks related to Section [3](#).
- Appendix [C](#) presents additional analyses and robustness checks related to Section [4](#).

A Appendix to Section 2: Data and Concepts

A.1 Definitions

A.1.1 Ownership

Comparing land ownership across countries with very different property-rights institutions requires being explicit about which rights a household actually holds. We follow the standard “bundle of rights” approach (Schlager and Ostrom 1992; Deininger 2003), under which property is a cumulative ladder of rights rather than an all-or-nothing status. In ascending order, a household may hold the right to *use* the land (to farm, graze, or otherwise exploit it), to *control* production (to make management decisions, improve the land, and keep its profits), to *bequeath* it within the family (for example through inheritance), to *fully transfer* it to others (by sale, lease, or use as collateral), and to hold a formal *title* that secures these rights against expropriation. In the terminology of Schlager and Ostrom (1992), a household holding use, control, and exclusion rights but not the right of sale is a *proprietor*; adding the right of alienation (sale) makes it an *owner*; a registered title adds formal-legal recognition above that.

Our working definition of ownership is a *proprietor who can also bequeath*: a household that holds use and control rights and can pass the land on within the family. We keep the word *ownership* for this bundle throughout the paper, while being clear that it is a proxy—closer to proprietorship-plus-bequest than to the sale-based *owner* of Schlager and Ostrom (1992). We adopt it for two reasons. First, it is the richest bundle of rights that is present, and measurable, in all seventeen countries: a landowner in Guatemala and a landowner in China alike can use their land, control its production, and pass it on, so inequality measured under this definition is comparable across the whole sample. Second, no stricter bundle is available everywhere: the right to freely sell the land, for example, does not exist for agricultural land in China, so a definition that also required the right of sale—or a formal legal title—could not be applied to every country, and comparability across the full sample would be lost. What matters economically is the security and transferability of land rights, not their formal legal form, so secure and heritable rights can function as ownership even where households lack freehold title (Deininger 2003). We therefore stay at this baseline throughout and do not estimate inequality under narrower definitions.

In practice we apply the most direct evidence each survey offers. Where the survey asks the household directly whether it owns land, we take the answer at face value. Where it does not, we infer ownership from how the plot was acquired, counting as owners those households whose land came through a permanent transfer of the kind that confers these rights—in most surveys *purchase* or *inheritance*. Rented, leased, and sharecropped-in land—a use right without control of the asset or any right to transfer it—is never counted as owned. Table A1 lists, for each country, the exact question and the full set of acquisition categories we treat as ownership.

A natural question is how this definition applies to countries where agricultural land is not privately owned at all. The cases readers most often have in mind are China and Vietnam, where land is formally owned by the state or the village collective rather than by households, so one might doubt whether their farmers should count as landowners. Under our definition they clearly do, because the rights that matter for our analysis rest with the household. Chinese families hold secure, long-term, heritable use rights: they farm the land, decide how it is used, and pass it on

within the family, even though they cannot sell it.¹⁵ Vietnamese households hold an even fuller bundle—since the 1993 Land Law they may also transfer, lease, and mortgage their land.¹⁶ Both therefore clear our baseline, and the contrast shows why we set that baseline at use, control, and bequest rather than at the right of sale: a sale-based rule would arbitrarily exclude Chinese farmers, who plainly hold and control heritable land, while admitting their otherwise-similar Vietnamese neighbours. A parallel logic applies where customary tenure predominates, notably in much of Sub-Saharan Africa: secure and heritable rights frequently exist without any formal title (Deininger 2003), so we again count as owners the households reporting purchase or inheritance, while land that is merely allocated for use is treated separately, as we turn to next.¹⁷

That separate category is land the household holds under customary or otherwise locally specific tenure regimes that carry fewer rights than our baseline requires—or, at least, fewer than we can establish with confidence from the survey. Such arrangements take many forms across our countries: land granted by village or local leaders (the Gambia; “distributed by the community” in Nigeria; “allocated by the village council” in Tanzania), squatted or occupied land (the *posseiro* in Brazil; land held “by possessory right” or through “invasion” in Panama), and communal or unregistered holdings (Guatemala), among others. What unites them is that the household controls the land’s use and cannot readily be evicted, yet holds no clear right to transfer or bequeath the plot, so the bundle falls short of ownership as we define it. We label this *non-owned land*; our baseline excludes it, and the relevant categories for each survey are listed in the final column of Table A1 (“N/A” where a survey has none).

Excluding non-owned land is consequential, because the ownership boundary is not neutral with respect to the outcomes we study. By restricting attention to land held under secure, transferable tenure, our baseline may disproportionately select households with stronger property rights; and because tenure security is itself an incentive to invest in land (Deininger 2003), these households may also have invested more in their plots, raising both their land values and their farm incomes. The estimated association between land value and income could then partly reflect this selection on secure, high-investment tenure rather than a relationship that holds more broadly across landholding households. This is precisely why we isolate non-owned land as a category: it lets us test the concern directly, by adding it back. We re-estimate our main income–land specification after reclassifying as owned every plot the household does not rent, lease, or sharecrop in—regardless of tenure regime—which brings into the sample exactly those households whose tenure is weaker and whose land-related investment is likely lower. Were the baseline relationship an artifact of selecting on secure, high-investment tenure, this broadening should attenuate the estimated elasticities. Instead, the coefficients on log land area and log value per hectare are essentially unchanged—0.118 and 0.192, against 0.122 and 0.202 in the baseline (Table B9, column 3)—indicating that our findings do not hinge on the precise ownership definition: the price

¹⁵Rural land in China is collectively owned; under the Household Responsibility System households hold long-term (currently thirty-year), renewable contracts that confer secure use and control rights and can be passed on within the household, but the land itself cannot be sold to other parties—so China meets our baseline but not a marketable definition (Vendryes 2010).

¹⁶The 1993 Land Law issued Vietnamese households land-use certificates conferring the rights to exchange, transfer, lease, inherit, and mortgage their land, even though all land remains formally state-owned (Ravallion and Van de Walle 2008).

¹⁷Customary systems are heterogeneous, so this proxy is necessarily approximate: we claim only that the purchased- or inherited-land subsample holds our baseline bundle, not that all customary land does.

per hectare, not physical area, carries the relationship whatever the tenure boundary.

Table A1: Ownership questions, ownership definitions, and non-owned land categories by country

Country	Ownership question	Ownership definition	Non-owned land
Bangladesh	How was the plot acquired?	Purchased/bought; Inherited (wife's or husband's family)	N/A
Brazil	In relation to this area, the household is: [owner / tenant / squatter / other] (<i>Em relação a esta área é</i>)	Owner (<i>proprietário</i>)	Squatter / occupant (<i>posseiro ou ocupante</i>)
China	How much collectively distributed land does your family have?	Collectively distributed land (heritable use rights)	N/A
Ecuador	How was the plot acquired?	Purchased; Inherited; Awarded; Usufruct	N/A
Ethiopia	How did your household acquire this parcel?	Inherited; Purchased	Granted by local leaders; Borrowed for free; Occupied without permission; Other
Gambia	How did the household acquire this parcel of land?	Inherited; Purchased	Use right granted by local authorities; Traded for another parcel; Other
Guatemala	What ownership documents do you have?	Registered title deed; Title deed in process; Inheritance (pending); Public deed	Communal/cooperative land; Owner without registered title; No title
India	Land owned by the household (NSS–AIDIS asset schedule, Blocks 5.1/5.2)	Owned	N/A
Indonesia	Do you have land for farming?	Yes (owns farm land)	N/A
Malawi	How did your household acquire this plot?	Inherited by the death of a family member; Bride price; Purchased; Allocated by family member	Granted by local leaders; Borrowed for free; Occupied without permission; Gift from non-household member; Other
Nicaragua	Reported in the module on owned farms/plots (<i>fincas o parcelas propias</i>)	Owned (owned-farms module)	N/A
Niger	Mode d'acquisition de cette parcelle / Way of acquiring this plot	Purchased; Inherited	Gift; Possession after clearing (<i>prise de possession simple</i>)
Nigeria	How was the plot acquired?; Have you or any household member acquired a legal title?	Outright purchase; Family inheritance; Holds legal title	Used free of charge; Distributed by community
Pakistan	Does the household possess personal agricultural land?	Owns personal agricultural land	N/A
Panama	How did you acquire this farm/plot? (<i>cómo adquirió esta finca o parcela</i>)	Purchase; Inheritance; Allocation/concession	Possessory right (<i>derecho posesorio</i>); Invasion/squatting; Other

Table A1 (continued)

Country	Ownership question	Ownership definition	Non-owned land
Tanzania	How was this plot acquired?	Inheritance; Gift; Purchased	Allocated by village council; Used free of charge
Vietnam	Why do you have this land parcel?	Inherited/given; Purchased; State-allocated; Reclaimed	N/A

A.1.2 Land Value

Table A2 reports the land-value measure used for each country. Our target concept is the same everywhere: the household’s estimate of what the plot would fetch *if it were sold today*—its market value. That exact question is not always asked, and where it is missing we adopt the closest available measure. In China the value is supplied *pre-constructed* by the survey producers rather than elicited from the household. In Ethiopia only an annual *rental* value is available. We use this measure directly in scale-invariant exercises, including household rankings, Gini coefficients, and log regressions, and capitalize it for the absolute land-poverty calculations in Section 4.3. In Nicaragua, there is no direct value question, but the survey itself provides regional per-*manzana* land prices, from which we impute value as owned area $\times$ price—so the imputation stays entirely within the survey. Brazil follows the same area- $\times$ -price logic but draws prices from an external official source—the Federal Revenue Service’s *Valor da Terra Nua* (bare-land value) table—resolved by both region and land type: because the Brazilian survey records land use in detail (crop vs. pasture), we apply type- and region-specific per-hectare prices, adding within-region heterogeneity to the imputed values.

A.1.3 Income Aggregate

This section describes the harmonized protocol used to construct total household income from survey microdata across all countries in our sample. Total household income is defined as:

$$\text{Total Income} = \text{Agricultural Income} + \text{Non-Agricultural Income} + \text{Non-labor Income}$$

where each component is constructed as follows:

- *Agricultural Income* = Own Farm Income + Wage Farm Income + Land Rent Income. Own farm income is net income from crop, livestock, forestry, and fishing production operated by the household—whether on owned, rented-in, or communal land—with rent paid for land deducted as a production cost. Wage farm income is earnings from agricultural wage work. Land rent income is payments received for renting out owned agricultural land.
- *Non-Agricultural Income* = Non-Agricultural Business Income + Non-Agricultural Wage Income. Business income is net self-employment or employer income in non-agricultural sectors. Wage income is earnings from non-agricultural wage employment.
- *Non-labor Income* has two parts: capital and property income (interest, dividends, and rent from non-agricultural assets) and transfers (pensions, remittances, government cash assis-

Table A2: Land value questions and valuation basis by country

Country	Land-value question (variable)	Valuation basis
Bangladesh	If you wanted to buy the plot, what amount would you need?	Self-reported value
Brazil	No survey question; owned area $\times$ official per-hectare bare-land values (<i>Valor da Terra Nua</i>), matched by municipality and land type (crop vs. pasture)	Imputed (external prices)
China	Land value	Provided by Survey
Ecuador	Estimated sale value (<i>valor estimado de venta</i>)	Self-reported value
Ethiopia	How much could you rent [parcel] for?	Annual rental value
Gambia	How much would the parcel fetch if sold today?	Self-reported value
Guatemala	Sale value (<i>valor de venta</i>)	Self-reported value
India	Value of owned land, NSS–AIDIS asset schedule	Self-reported value
Indonesia	Value of farm land (value-bracket midpoints where missing)	Self-reported value
Malawi	If you sold this [garden] today, how much could you sell it for?	Self-reported value
Nicaragua	Regional per- <i>manzana</i> land price provided by the survey $\times$ owned area	Imputed (within survey)
Niger	Current sale price of the parcel if you had to sell, in FCFA	Self-reported value
Nigeria	If the [plot] were to be sold today, how much could it be sold for?	Self-reported value
Pakistan	If the owner wanted to sell, what is the expected price?	Self-reported value
Panama	Estimated value of the farm (<i>valor estimado de la finca</i>)	Self-reported value
Tanzania	What would the value of this plot be if it were sold today?	Self-reported value
Vietnam	Current market value of this plot, excluding the house (province-average price where “don’t know”)	Self-reported value

Source (Brazil). Land values are imputed from the Brazilian Federal Revenue Service’s *Valor da Terra Nua* (VTN) bare-land-value table: <https://www.gov.br/receitafederal/pt-br/assuntos/orientacao-tributaria/auditoria-fiscal/arquivos/pdf-planilha-final-para-publicacao2020.pdf/view>.

tance, scholarships, and social benefits). Our benchmark income retains only the former, as explained below.

Several harmonization rules apply across all countries. Survey responses coded as “don’t know” or similar are recoded to missing. All incomes are annualized. Business incomes are net of costs where revenue data are available; where missing, independent income from the labor module is substituted. Negative own-farm or business incomes are set to zero. One-off receipts (asset sales, lotteries, inheritances) are excluded. If all components are missing for a household, total income is assigned missing; if at least one component is observed, missing components are set to zero. When an activity filter confirms the household did not engage in a given activity, the corresponding

missing income is recoded to zero. Only cash income is included where the distinction between cash and in-kind is available.

Our benchmark definition of income includes labor and capital income but excludes benefits and transfers, so that income reflects the household's own earning capacity—the returns to its labor and assets—rather than redistributive flows. Within each survey we isolate a transfer component—pensions, remittances, government cash assistance, scholarships, and any social or war-related benefits—and remove it, leaving labor income together with capital and property income (interest, dividends, and rent on non-agricultural assets). This choice matters because transfers partially smooth income among asset-poor households without altering their underlying productive-asset constraints; including them would narrow precisely the differences across landholding groups that we study. We hold transfers in a separate variable so they can be added back: doing so recovers a broader income concept and leaves the main estimates essentially unchanged (Appendix B.1.5, column 4). The split is stable across the sample—what we classify as transfers is almost always pensions, remittances, and public cash assistance (and, in some surveys, scholarships or social benefits), while what remains in non-labor income is the household's capital income from interest, dividends, and non-agricultural rents.

Worked Example: Guatemala (2000) Table A3 illustrates the income aggregate construction for Guatemala, based on the *Encuesta Nacional sobre Condiciones de Vida* (2000). The table reports the full, transfer-inclusive construction. The benchmark used in the analysis removes the transfer components of Non-labor Income—pensions, remittances, cash assistance, and scholarships—while keeping dividends and non-farm rent as capital and property income; the transfer-inclusive total in the last row is the aggregate used when transfers are added back in Appendix B.1.5 (column 4).

Table A3: Income aggregate components: Guatemala (2000)

Component	Description	Survey File
Own Farm Income	Net crop income + net forestry income + net live-stock income – rent of land	ECV2#A16; ECV3#A16
Wage Farm Income	Gross cash salary (incl. bonuses/tips) from main job + second job + other jobs (last 12 months)	ECV11P10
Land Rent Income	Total money received for renting out land in the last 12 months	ECV21A16; ECV22A16
Agricultural Income	= Own Farm + Wage Farm + Land Rent	
Own Business Income	Net independent income from main + second + other non-agricultural jobs	ECV11P10
Wage Income	Gross cash salary from main + second + other non-agricultural jobs	ECV11P10
Non-Agric. Income	= Own Business + Wage	
Non-labor Income	Pensions + dividends + remittances + cash assistance + non-farm rent + scholarships	ECV16I13; ECV11P10
Total HH Income (incl. transfers)	= Agricultural + Non-Agricultural + Non-labor	

A.1.4 Wealth

For the four surveys that carry a detailed household asset module—Ethiopia (ESS 2018), Malawi (IHPS 2016), Tanzania (NPS 2019), and India (AIDIS 2019)—we also build a measure of total household wealth as the current value of all assets the household owns. We sort each survey’s asset items into three groups: the *dwelling*, the household’s *agricultural plot(s)*, and *other (non-land) assets*. The dwelling and the plot are valued at the respondent’s own estimate of current value; the *other* group collects everything the survey measures beyond land and housing, so its exact composition depends on how wide each survey casts its asset net (Table A4). The plot valuation applies the same concept as the land-value measure described above—the household’s own assessment of current value—but comes from each survey’s asset module, which in Ethiopia, Malawi, and Tanzania is a separate instrument from the land-value question; only in India do both originate from the same module (Appendix C.8). All figures are PPP-adjusted, consistent with the income and land-value measures elsewhere in the paper. We winsorize each wealth component at the 99th percentile within survey before summing them, mirroring our treatment of income.

Our primary measure, used in the wealth columns (4)–(6) of Table 2, is *Non-Plot Assets*: dwelling + other, i.e. total wealth net of land value. Because the value of the agricultural plot is itself one of our regressors, we benchmark on the measure that excludes the plot, so the outcome cannot mechanically include land value. The appendix reports two alternatives that re-partition the same asset pool: *Total Assets* (dwelling + plot + other) and *Productive Assets* (plot + other, dropping the dwelling). Reporting all three guards against the concern that the wealth–land relationship is mechanical, and the estimated elasticities are stable across the three definitions (Appendix B.1.9).

Table A4: Household wealth components by survey

Country vey)	(sur-	Dwelling	Agric. plot	Other (non-land) assets
Ethiopia (ESS4)	2018	Dwelling value	Plot value	Financial account balances; mobile phones; livestock
Malawi (IHPS)	2016	Dwelling value	Plot (garden) value	Financial account balances
Tanzania (NPS)	2019	Dwelling value	Plot value	Financial account balances; non-farm business capital
India (AIDIS, NSS-77)	2019	Buildings and constructions	Land value	Financial assets (shares, deposits); livestock; transport equipment; agri- cultural machinery; non-farm business equipment

Notes. The dwelling and plot are valued at the household's self-reported current value. **Non-Plot Assets** (the main-table measure) = dwelling + other (total wealth net of land value); **Total Assets** = dwelling + plot + other; **Productive Assets** = plot + other (excl. dwelling). All monetary values are PPP-adjusted. The alternative measures are reported in Appendix B.1.9.

As an external level benchmark for overall wealth concentration, we match WID.world estimates of the net-wealth Gini to each survey's year for all seventeen countries and take the unweighted cross-country mean, 0.76. WID wealth distributions for several of these countries are extrapolated from regional aggregates rather than estimated from national micro-data; we therefore use this figure only as a level reference for the concentration of land wealth and do not interpret its cross-country dispersion.

A.1.5 Loans

Eleven of our surveys record household borrowing, which we use as a third outcome, *market credit* (the loan columns of Table 2 and the credit robustness in Appendix B.1.7). For each household we sum the outstanding principal across all of its loans; borrowing is never added to household income. In building this total, our definition excludes credit from lenders whose loans function more as subsidized transfers than as ordinary market credit—chiefly state development banks, non-governmental organizations (NGOs), and microfinance institutions. Such lending is typically directed and offered below market rates, so it reflects program eligibility rather than a household’s creditworthiness or its demand for capital. We apply this rule wherever the survey identifies the lender; where a survey records the source only coarsely, or not at all, no exclusion is possible and we keep the full loan amount. Table A5 lists the lender types removed for each survey.

Table A5: Loan sources excluded by country

Country	Loan sources excluded from the household total
Bangladesh	NGOs; state agricultural development banks (Bangladesh Krishi Bank; Rajshahi Krishi Development Bank)
Brazil	None
Ecuador	Public financial institutions
Ethiopia	Microfinance institutions; NGOs
Gambia	Microfinance institutions; NGOs
Guatemala	NGOs
Indonesia	None
Malawi	NGOs
Pakistan	None
Tanzania	Microfinance institutions; NGOs
Vietnam	State development banks (Bank for Social Policies; Bank for Agriculture and Rural Development)

Notes. The excluded sources are lenders whose credit is directed or subsidized, and thus closer to a transfer than to market credit. “None” means no source-based exclusion is applied—the survey records the lender only coarsely, or not at all. All remaining sources are retained and summed to the household level; the loan total is PPP-adjusted and enters the analysis only as an outcome, never as a component of income.

A.1.6 Alternative Definition of Landless Households

This subsection assesses the robustness of our results to an alternative definition of landless agricultural households. Our benchmark definition classifies as landless those households that report owning no agricultural land. The alternative definition relies on an agricultural classification based on labor participation, thereby modifying the reference population and, mechanically, the measured share of landless households. We show that this change has negligible implications for our main results. In Bangladesh, however, we are unable to construct the alternative definition due to data limitations, and the analysis therefore excludes this country.


Figure A1 presents a direct comparison of the two definitions. Panel (a) reports the landless share of agricultural households across countries using side-by-side bars, while Panel (b) displays a scatter plot contrasting the benchmark and alternative measures. The two definitions produce remarkably similar levels of landlessness. Differences are quantitatively small and do not alter the cross-country ranking. Countries with high landlessness under the benchmark definition (e.g., Guatemala, Brazil, and Nicaragua) remain so under the alternative definition, while countries with low landlessness (e.g., China and Vietnam) continue to exhibit very low shares. The scatter plot shows a tight clustering around the 45-degree line, indicating an almost one-to-one relationship between the two measures and a very high correlation between them.

Given this close correspondence, redefining the agricultural population has limited quantitative impact on our core results. While the alternative definition mechanically changes the reference group, the relative position of countries and the magnitude of landlessness remain broadly unchanged. This reinforces the conclusion that our findings are not driven by the particular way in which landless households are defined.


A.2 Legal Status of the agricultural holder

In seven of the seventeen countries in our study, the FAO Agriculture Census of 2010 provides information on the legal status of the land operator. We classify land as household-operated when its legal status is reported under any of the following three categories: Civil, Individual(s), or Family Group/HHSDs. According to [Food and Agriculture Organization of the United Nations \(2005\)](#), Civil encompasses cases in which land is operated “singly, by an individual or household; jointly, by two or more individuals or households.” The FAO Census makes explicit that countries are encouraged to distinguish in their census tables between the household sector, captured by these definitions, and the non-household sector, which includes corporations and cooperatives. The share of total land operated by the household sector is as follows for the seven countries with available information: Brazil, 85%; Ecuador, 88%; Guatemala, 86%; India, 99%; Nicaragua, 99%; Panama, 91%; and Vietnam, 80%.

Figure A1: **Alternative Definitions of Landless Households**



(a) Share of Landless Households: Benchmark vs. Alternative Definition



(b) Scatter Plot: Benchmark vs. Alternative Definition

Notes: Panel (a) reports the landless share of agricultural households under the benchmark definition and the alternative definitions of landless. Differences in levels are small and do not alter the cross-country ranking. Panel (b) plots the two measures against each other. Observations cluster tightly around the 45-degree line, indicating a very high correlation between definitions and confirming that the choice of definition has a negligible impact on cross-country comparisons.

A.3 Census–Survey and Alternative Survey Comparisons

This appendix compares inequality measures derived from agricultural censuses and household surveys. Agricultural censuses report the distribution of *operational holdings*, whereas household surveys allow us to compute inequality among *landowner households*. We focus on the Gini index of land area and land value, and assess the extent to which census-based measures approximate survey-based estimates.

Appendix Figure A2a plots the Gini coefficient of land area from the census (operational holdings) against the corresponding land-area Gini among landowners from household surveys. The two measures are strongly correlated across countries, with an R^2 of 0.94. This suggests that census data provide a reliable proxy for land-area inequality among landowners.


Appendix Figure A2b compares census-based land-area inequality (operational holdings) with survey-based land-value inequality among landowners. The correlation is substantially weaker ($R^2 = 0.56$), indicating that land area is an imperfect proxy for the distribution of land wealth.

Finally, Appendix Figure A2c plots survey-based land-area inequality against survey-based land-value inequality among landowners. The correlation remains moderate ($R^2 = 0.63$), confirming that differences between area and value inequality arise even within the same data source.


Taken together, these results show that while census data capture land-area inequality reasonably well, they are less informative about land-value inequality. Since land value is more closely related to wealth, relying exclusively on land-area measures may misrepresent cross-country differences in land inequality.

The following Appendix Figure A3 provides an additional validation by comparing land-area Gini coefficients from our benchmark household surveys with estimates obtained from the Demographic and Health Surveys (DHS) and from alternative waves of the same surveys. For each country, the figure plots Gini coefficients across available sources and years. The DHS-based estimates closely match the benchmark survey values, reinforcing the conclusion that household surveys provide reliable measures of land-area inequality. The stability across survey waves is also consistent with the high persistence of land inequality documented in the literature. Appendix Figure A4 provides the same time stability across survey years for the value estimates.


Figure A2: Census and survey measures of land inequality



(a) Census vs. Survey: Land-area Gini




(b) Census (area) vs. Survey (value): Gini




(c) Within-survey: Land-area vs. land-value Gini

Notes: Census-based Gini indices are computed over *operational holdings*, while survey-based Gini indices are computed over *landowner households*. Panel (a) plots census and survey land-area Gini coefficients. Panel (b) plots the census-based land-area Gini against the survey-based land-value Gini. Panel (c) compares, within surveys, land-area and land-value Gini indices among landowners. Each dot corresponds to a country; the fitted line is from an OLS regression. When reported in the text, the R^2 values are 0.94 in Panel (a), 0.56 in Panel (b), and 0.63 in Panel (c).


Figure A3: Land area inequality: alternative country-specific surveys and DHS




(a) Bangladesh




(b) Ethiopia




(c) Gambia




(d) Guatemala




(e) Indonesia




(f) India




(g) Malawi




(h) Niger



(i) Nigeria




(j) Pakistan




(k) Tanzania


Figure A4: Land value inequality: alternative country-specific surveys




(a) India




(b) Malawi



(c) Niger



(d) Nigeria



(e) Tanzania

B Appendix to Section 3: Relevance of land value

B.1 Robustness

B.1.1 Population-Weighted Results

The pooled regressions in the main text renormalize each survey's probability weights to sum to one, so that every survey carries the same total weight. Table B6 instead uses the surveys' original weights, under which each country enters in proportion to its number of households and China alone accounts for roughly three quarters of the weighted income sample. The pattern of the main table is unchanged: value remains the stronger margin for income and non-plot wealth, the loan estimates are close to the baseline, and the income coefficients are larger under population weighting.

Table B6: Robustness: Population-Weighted Regressions

	<u>Income</u>			<u>Non-Plot Wealth</u>			<u>Loans</u>		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Log Land Area	0.153*** (0.028)	–	0.386*** (0.032)	0.139*** (0.025)	–	0.319*** (0.024)	0.204*** (0.032)	–	0.300*** (0.036)
Log Land Value	–	0.427*** (0.032)	–	–	0.362*** (0.027)	–	–	0.264*** (0.030)	–
Log Value/Area	–	–	0.445*** (0.034)	–	–	0.396*** (0.035)	–	–	0.221*** (0.029)
Survey FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Clusters	441	441	441	174	174	174	295	295	295
Surveys	16	16	16	4	4	4	11	11	11
Sample	Owners	Owners	Owners	Owners	Owners	Owners	Borrowers	Borrowers	Borrowers
Observations	50,031	50,031	50,031	34,913	34,913	34,913	8,303	8,303	8,303
Adj. R^2	0.384	0.478	0.479	0.327	0.435	0.438	0.661	0.671	0.673
Within R^2	0.0137	0.1643	0.1661	0.0209	0.1774	0.1823	0.0382	0.0664	0.0704
Mean dep. var.	6.089	6.089	6.089	9.593	9.593	9.593	5.940	5.940	5.940

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Reproduces the main specification (Table 2) using the surveys' original probability weights instead of the per-survey renormalized weights: each country enters in proportion to its number of households, and China alone accounts for roughly three quarters of the total weight in the income columns. Income columns pool the income surveys (owners); wealth columns use the four wealth surveys with Non-Plot Assets as the outcome; loan columns use income surveys with reported loans (owner-borrowers). All regressions include survey fixed effects; standard errors clustered at the survey-strata level.

B.1.2 Alternative Treatment of Zero Income

Table B7 addresses the treatment of zero-income households. The baseline drops households with zero or negative equivalized income before taking logarithms; if these households differ systematically in their landholdings, excluding them could affect the estimates. We therefore retain them using two transformations defined at zero, $\ln(y + 1)$ (columns 1–3) and the inverse hyperbolic sine

(columns 4–6), keeping the three main-text specifications and all other restrictions unchanged. The joint estimates are close to the baseline: 0.160 (area) and 0.230 (value per hectare) under $\ln(y + 1)$, and 0.145 and 0.215 under the inverse hyperbolic sine, against 0.122 and 0.202 in Table 2, with the value-per-hectare coefficient slightly larger. Adding value still raises the within- R^2 substantially, though less than in the baseline—as expected, since zero-income households add noise that landholdings cannot explain. Appendix B.1.8 examines whether non-response is related to landholdings.

Table B7: Robustness: Alternative Treatment of Zero Income

	$\ln(Y + 1)$			$\text{asinh}(Y)$		
	(1)	(2)	(3)	(4)	(5)	(6)
Log Land Area	0.062** (0.029)	–	0.160*** (0.028)	0.053* (0.031)	–	0.145*** (0.031)
Log Land Value	–	0.187*** (0.023)	–	–	0.172*** (0.026)	–
Log Value/Area	–	–	0.230*** (0.024)	–	–	0.215*** (0.027)
Survey FE	Yes	Yes	Yes	Yes	Yes	Yes
Clusters	444	444	444	444	444	444
Surveys	16	16	16	16	16	16
Sample	Owners	Owners	Owners	Owners	Owners	Owners
Observations	58,083	58,083	58,083	58,215	58,215	58,215
Adj. R^2	0.367	0.375	0.376	0.338	0.344	0.345
Within R^2	0.0018	0.0153	0.0168	0.0010	0.0105	0.0117

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Agricultural landowners across all 16 income surveys. Columns (1)–(3) use $\ln(y + 1)$ and columns (4)–(6) the inverse hyperbolic sine $\text{asinh}(y)$, both retaining zero-income households, with the area / value / value-per-hectare specifications. $\ln(y + 1)$ is defined only for $y > -1$, so the $\ln(y + 1)$ columns drop a small number (about 130) of negative net-of-transfer income households that the asinh columns retain. Survey FE, probability weights renormalized so that each survey carries equal total weight, standard errors clustered at the survey–strata level.

A related concern is income measurement itself: several surveys, particularly in Sub-Saharan Africa, were not designed to construct income aggregates, and self-reported income may capture only part of household earnings. Table B8 therefore replaces income with the consumption aggregates built by the survey providers, available for eight surveys. It reports the three main-text specifications for income and for log equivalized consumption, first dropping zero-income households as in the baseline (columns 1–6), then retaining them (columns 7–12). Consumption responds mainly to the price of land: on its own, the area coefficient is slightly negative and insignificant, while in the joint specification the value-per-hectare coefficient is positive, precisely estimated, and more than twice the area coefficient, with or without zero-income households. Because consumption is measured independently of self-reported income, the value margin is not an artifact of noisy income data. The income estimates in this eight-survey subsample are also close to the full-sample results in Table 2, though the area-alone coefficient turns slightly negative and insignificant.

Table B8: Robustness: Income and Consumption Returns to Land

	Drop Zero-Income Households						Retain Zero-Income Households					
	ln(Income)			ln(Consumption)			ln($Y + 1$)			ln(Consumption)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Log Land Area	-0.030 (0.024)	–	0.100*** (0.027)	-0.024 (0.016)	–	0.053*** (0.012)	-0.039 (0.037)	–	0.093*** (0.036)	-0.025 (0.016)	–	0.050*** (0.012)
Log Land Value	–	0.159*** (0.026)	–	–	0.092*** (0.010)	–	–	0.159*** (0.028)	–	–	0.090*** (0.010)	–
Log Value/Area	–	–	0.211*** (0.030)	–	–	0.126*** (0.015)	–	–	0.215*** (0.034)	–	–	0.123*** (0.016)
Survey FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Clusters	179	179	179	179	179	179	180	180	180	180	180	180
Surveys	8	8	8	8	8	8	8	8	8	8	8	8
Sample	Owners	Owners	Owners	Owners	Owners	Owners	Owners	Owners	Owners	Owners	Owners	Owners
Observations	19,738	19,738	19,738	19,738	19,738	19,738	23,722	23,722	23,722	23,722	23,722	23,722
Adj. R^2	0.253	0.268	0.276	0.826	0.832	0.836	0.235	0.240	0.243	0.822	0.828	0.832
Within R^2	0.0009	0.0206	0.0308	0.0033	0.0384	0.0623	0.0005	0.0072	0.0113	0.0035	0.0367	0.0606

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Agricultural landowners in the 8 surveys with consumption data (China 2012, Ethiopia 2018, Guatemala 2000, Malawi 2016, Niger 2014, Nigeria 2015, Panama 2008, Tanzania 2014). Columns (1)–(6) drop zero-income households as in the baseline; columns (7)–(12) retain them, with income entering as $\ln(y + 1)$. Consumption is equalized identically to income (OECD scale; China: per capita). Survey FE, probability weights renormalized so that each survey carries equal total weight, standard errors clustered at the survey–strata level.

B.1.3 Including Non-Owned Land

Ownership is difficult to measure consistently across surveys: our harmonized definition is a proxy built from each survey’s ownership questions (Appendix A.1.1), and it may draw the ownership boundary imperfectly. Table B9 therefore applies a stringent alternative that does not rely on that boundary at all: it re-estimates the main table counting all land held by the household—everything other than rented, leased, or sharecropped-in plots—adding non-owned plots to both area and value and counting their holders as owners. The income and loan estimates are very close to the main table, and the wealth estimates show the same pattern, with somewhat smaller value coefficients.

B.1.4 Cross-Country Heterogeneity

Figure B5 estimates the joint specification separately for each of the 16 income surveys. Land markets and agricultural structure differ widely across countries, so the pooled estimates could in principle be driven by a subset of surveys. The value-per-hectare coefficient is positive and statistically significant in 13 of the 16 surveys; the exceptions are Gambia and Ethiopia, where it is close to zero, and Niger, where it is negative and imprecise. The area coefficient is of comparable size in most countries; since more hectares at a constant price per hectare also mean proportionally more land value, the area coefficient partly reflects value, and the effect of hectares per se is the gap between the two coefficients, which is small in most surveys. The cross-country estimates are in line with the pooled results.

Table B9: Robustness: Including Non-Owned Land

	<u>Income</u>			<u>Non-Plot Wealth</u>			<u>Loans</u>		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Log Land Area	0.025 (0.018)	–	0.118*** (0.019)	0.104* (0.059)	–	0.272*** (0.054)	0.179*** (0.028)	–	0.285*** (0.032)
Log Land Value	–	0.149*** (0.016)	–	–	0.296*** (0.046)	–	–	0.254*** (0.027)	–
Log Value/Area	–	–	0.192*** (0.017)	–	–	0.304*** (0.052)	–	–	0.220*** (0.027)
Survey FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Clusters	443	442	442	174	174	174	299	298	298
Surveys	16	16	16	4	4	4	11	11	11
Sample	Owners	Owners	Owners	Owners	Owners	Owners	Borrowers	Borrowers	Borrowers
Observations	55,749	53,807	53,807	36,192	35,978	35,978	9,250	8,829	8,829
Adj. R^2	0.566	0.586	0.588	0.319	0.391	0.392	0.615	0.631	0.632
Within R^2	0.0008	0.0282	0.0331	0.0083	0.1092	0.1099	0.0309	0.0638	0.0667
Mean dep. var.	6.077	6.079	6.079	9.518	9.535	9.535	5.862	5.879	5.879

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Reproduces the main specification (Table 2) counting all land held by the household other than rented, leased, or sharecropped-in plots: non-owned plots are added to both land area and land value, and their holders are counted as owners. Income columns pool the 16 income surveys; wealth columns apply the same reclassification in the four wealth surveys, with Non-Plot Assets as the outcome; loan columns use income surveys with reported loans (owner-borrowers). All regressions include survey fixed effects and probability weights renormalized so that each survey carries equal total weight; standard errors clustered at the survey–strata level.


B.1.5 Main Specification Variants

Table B10 re-estimates the joint specification under four alternative data choices: retaining observations with missing land flags (column 1), applying no winsorization to income (column 2), winsorizing at the 95th rather than the 99th percentile (column 3), and using total income including transfers (column 4). The estimates are close to the baseline in all four columns.

B.1.6 Excluding Countries without Directly Self-Reported Land Values

Our benchmark measure of land value is the household’s own report of the market value of its plots. Four countries are the exception: in China (2012) the value is supplied pre-constructed by the survey producers, in Ethiopia (2018) only an annual rental value is available, and in Brazil (1997) and Nicaragua (2014) values are constructed from regional land prices—taken from an external source in Brazil and from within the survey in Nicaragua (Appendix A.1.2 and Table A2). Table B11 re-estimates the main regressions excluding these four countries, leaving 12 income, 9 loan, and 3 wealth surveys. The income and loan estimates are close to the baseline, except that area alone loses significance in the income regression; the wealth estimates, now based on the three remaining wealth surveys, show the same pattern as the main table, with larger value coefficients.

Figure B5: Land Returns by Country: Main Specification



Notes: Each pair of points shows the OLS coefficients from jointly regressing $\ln(\text{income})$ on log land area and log value per hectare for agricultural landowners within a single survey, using probability weights and the survey design. Horizontal bars: 95% confidence intervals. Surveys sorted by the log value-per-hectare coefficient (descending). Zero-income households dropped. Open triangle ($\triangle$) = $\ln(\text{Value}/\text{Area})$; open diamond ($\diamond$) = $\ln(\text{Area})$.

B.1.7 Subsample Robustness for Wealth and Credit Results

Table B12 addresses the credit results, which in the main table are estimated among borrowers. Columns (1)–(3) estimate the extensive margin: an indicator for any outstanding loan, regressed on the three specifications among landowners in the 11 loan surveys. All coefficients are essentially zero—land holdings do not predict whether a household borrows. Columns (4)–(6) estimate the income regressions on all landowners in these surveys, and columns (7)–(9) on borrowers only; the estimates are close across the two samples.

Table B13 addresses the difference in country coverage between the income and wealth samples. Columns (1)–(3) estimate the income regressions on the three income surveys that overlap with the wealth countries (Ethiopia 2018, Malawi 2016, Tanzania 2014); columns (4)–(6) estimate the non-plot wealth regressions on three wealth surveys, excluding India 2019. The pattern is the same as in the full samples: area alone is insignificant in both panels, while value per hectare remains significant—at the 1 percent level for income, with smaller coefficients than in the full sample, and at the 1 percent level for wealth, with estimates close to the baseline.

B.1.8 Non-Response Analysis

Figure B6 examines whether the households dropped from the baseline—those reporting zero or missing income—differ systematically in their landholdings. Panel A reports a pooled linear

Table B10: Robustness: Main Specification Variants

	Missing retained (1)	No winsor. (2)	Winsor. 95th (3)	Incl. transfers (4)
Log Land Area	0.122*** (0.019)	0.132*** (0.020)	0.103*** (0.017)	0.104*** (0.017)
Log Value/Area	0.202*** (0.018)	0.211*** (0.018)	0.184*** (0.017)	0.193*** (0.016)
Survey FE	Yes	Yes	Yes	Yes
Clusters	441	442	439	442
Surveys	16	16	16	16
Sample	Owners	Owners	Owners	Owners
Observations	50,031	50,372	48,541	53,692
Adj. R^2	0.582	0.573	0.598	0.613
Within R^2	0.0339	0.0361	0.0298	0.0341


Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Dependent variable: log income for agricultural landowners, on log land area and log value per hectare. Column (1) retains observations with missing land/value flags; (2) applies no winsorization; (3) winsorizes income at the 95th percentile; (4) uses total income including transfers. All regressions include survey fixed effects and probability weights renormalized so that each survey carries equal total weight; standard errors clustered at the survey–strata level.

probability model of a non-response indicator on within-country standardized log land area and log value per hectare: both coefficients are small—below half a percentage point per standard deviation—and statistically insignificant, so there is no strong evidence that non-response varies with landholdings. Panel B repeats the regression country by country; the grey band marks effects below one percentage point per standard deviation, which we treat as economically negligible, since a differential of less than one percentage point per standard deviation alters the composition of the estimation sample by too little to move the income estimates. Panel C regresses the log of the survey providers’ own income aggregates, which impute income for non-reporting households, on log land area and log value per hectare in China and Panama, without dropping zero-income observations.

With no strong evidence that non-response varies with landholdings in the pooled sample, dropping zero- and missing-income households is unlikely to select on land area or value, and the baseline estimates are unlikely to be biased by their exclusion. Table B7 is consistent with this: retaining zero-income households, the value-per-hectare coefficient remains larger than the area coefficient and significant at the 1 percent level.

Panel C examines whether the excluded households are genuinely zero-income. In China and Panama—two countries where the country-level selection estimates in Panel B are among the largest—the surveys provide internally constructed income aggregates that impute income for non-reporting households. Regressing these aggregates on log land area and log value per hectare, without dropping zero-income observations, yields positive and significant coefficients on both.

Figure B6: Non-Response and Standardized Land Returns



Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Dependent variable: indicator for zero or missing income among agricultural landowners (Panels A–B); log survey-imputed income (Panel C). Panel A: survey fixed effects and probability-weighted OLS; standard errors clustered at survey–strata level; regressors standardized within each country. Panel B: survey-weighted OLS estimated separately for each country; regressors standardized within each country’s owner sample; 95% confidence intervals shown. Open triangle ($\triangle$) = std. ln(Value/Area); open diamond ($\diamond$) = std. ln(Area). Countries sorted by ln(Value/Area) coefficient (descending). Panel C: OLS with probability weights and robust standard errors; no survey fixed effects (single-country regressions). Outcome is log of the survey designers’ own income aggregate, which imputes income for non-reporting households. Zero-income households are not dropped.

Table B11: Income, Non-Plot Wealth, and Loans Regressions — Excluding Countries without Directly Self-Reported Land Values

	Income			Non-Plot Wealth			Loans		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Log Land Area	0.009 (0.020)	–	0.119*** (0.020)	0.020 (0.054)	–	0.357*** (0.084)	0.194*** (0.030)	–	0.307*** (0.032)
Log Land Value	–	0.152*** (0.015)	–	–	0.478*** (0.086)	–	–	0.268*** (0.027)	–
Log Value/Area	–	–	0.186*** (0.017)	–	–	0.531*** (0.086)	–	–	0.228*** (0.028)
Survey FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Clusters	372	372	372	164	164	164	277	277	277
Surveys	12	12	12	3	3	3	9	9	9
Sample	Owners	Owners	Owners	Owners	Owners	Owners	Borrowers	Borrowers	Borrowers
Observations	41,037	41,037	41,037	33,650	33,650	33,650	8,157	8,157	8,157
Adj. R^2	0.622	0.631	0.632	0.396	0.530	0.543	0.637	0.649	0.651
Within R^2	0.0001	0.0226	0.0269	0.0003	0.2217	0.2435	0.0353	0.0671	0.0714
Mean dep. var.	5.999	5.999	5.999	9.660	9.660	9.660	5.958	5.958	5.958

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Reproduces the main specification after excluding the four countries whose land value is not directly self-reported but based on the closest substitute: China (2012), Ethiopia (2018), Brazil (1997), and Nicaragua (2014). Income columns use the remaining 12 income surveys; loan columns the remaining 9 loan surveys. The wealth columns use the three remaining wealth surveys (Malawi IHPS 2016, Tanzania 2019, India 2019), with Non-Plot Assets as the outcome. All regressions include survey fixed effects and probability weights renormalized so that each survey carries equal total weight; standard errors clustered at the survey–strata level.

B.1.9 Alternative Wealth Measures

Table B14 checks the wealth result across three definitions of household wealth and two outcome transformations. *Non-Plot Assets* (columns 5–6) is the benchmark used in the main table: dwelling plus other (financial and non-farm) assets, excluding the agricultural plot, so that the land-value regressor cannot enter the outcome mechanically. *Total Assets* (columns 1–2) adds the plot back in, and *Productive Assets* (columns 3–4) excludes the dwelling instead. For each definition, the first column uses $\ln(\text{wealth})$, restricting to households with strictly positive wealth, and the second uses the inverse hyperbolic sine, which retains zero-wealth households. The value-per-hectare coefficient is positive and significant at the 1% level in all six columns; it is largest for *Total Assets*, which includes the plot’s own value in the outcome, and smallest for *Productive Assets*. In the plot-inclusive measures the plot component is the asset module’s own valuation rather than the land value on the right-hand side, so outcome and regressor share the same variable only in India.

Because wealth components are measured differently across surveys, particularly the non-farm and financial asset categories, Figure B7 estimates the specification separately for each of the four wealth surveys. The value-per-hectare coefficient is positive and significant in Tanzania, Malawi, and India, and the area coefficient in all four. Ethiopia is the exception, with a value-per-hectare coefficient close to zero and insignificant; it also stands out in the income estimates

Table B12: Loan Robustness

	Any Loans			Income (Full)			Income (Borrowers)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Log Land Area	0.001 (0.003)	–	0.002 (0.003)	0.050*** (0.019)	–	0.113*** (0.020)	0.102*** (0.021)	–	0.163*** (0.023)
Log Land Value	–	0.002 (0.003)	–	–	0.133*** (0.017)	–	–	0.148*** (0.021)	–
Log Value/Area	–	–	0.002 (0.003)	–	–	0.162*** (0.019)	–	–	0.132*** (0.024)
Survey FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Clusters	349	349	349	349	349	349	295	295	295
Surveys	11	11	11	11	11	11	11	11	11
Sample	Owners	Owners	Owners	Owners	Owners	Owners	Borrowers	Borrowers	Borrowers
Observations	39,267	39,267	39,267	39,267	39,267	39,267	8,303	8,303	8,303
Adj. R^2	0.087	0.087	0.087	0.639	0.646	0.647	0.628	0.633	0.633
Within R^2	0.0000	0.0001	0.0001	0.0034	0.0239	0.0259	0.0118	0.0260	0.0268
Mean dep. var.	0.211	0.211	0.211	5.939	5.939	5.939	6.352	6.352	6.352

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Columns (1)–(3): extensive margin, dependent variable an indicator for any positive outstanding loan (owners in the 11 loan surveys). Columns (4)–(6): log income for all owners in the 11 loan surveys. Columns (7)–(9): log income restricted to owner-borrowers (positive outstanding loans). All regressions include survey fixed effects and probability weights renormalized so that each survey carries equal total weight; standard errors clustered at the survey–strata level.

Table B13: Wealth Robustness: Subsamples

	Income (Subsample)			Non-Plot Wealth (Subsample)		
	(1)	(2)	(3)	(4)	(5)	(6)
Log Land Area	0.024 (0.041)	–	0.094** (0.044)	0.019 (0.054)	–	0.288*** (0.074)
Log Land Value	–	0.115*** (0.036)	–	–	0.382*** (0.090)	–
Log Value/Area	–	–	0.127*** (0.041)	–	–	0.423*** (0.096)
Survey FE	Yes	Yes	Yes	Yes	Yes	Yes
Clusters	93	93	93	87	87	87
Surveys	3	3	3	3	3	3
Sample	Owners	Owners	Owners	Owners	Owners	Owners
Observations	8,677	8,677	8,677	3,295	3,295	3,295
Adj. R^2	0.042	0.051	0.052	0.131	0.230	0.239
Within R^2	0.0004	0.0105	0.0111	0.0002	0.1150	0.1257
Mean dep. var.	5.522	5.522	5.522	7.541	7.541	7.541

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Columns (1)–(3): log income for the three income surveys matching the wealth countries (Ethiopia 2018, Malawi 2016, Tanzania 2014). Columns (4)–(6): Non-Plot Assets for the three wealth surveys, excluding India (2019). Reproduces the earlier wealth robustness with the new area / value / value-per-hectare specifications. All regressions include survey fixed effects and probability weights renormalized so that each survey carries equal total weight; standard errors clustered at the survey–strata level.

Table B14: Robustness: Alternative Wealth Measures

	Total Assets		Productive Assets		Non-Plot Assets	
			(excl. Dwelling)		(excl. Plot)	
	ln	asinh	ln	asinh	ln	asinh
	(1)	(2)	(3)	(4)	(5)	(6)
Log Land Area	0.470*** (0.075)	0.441*** (0.077)	0.583*** (0.084)	0.678*** (0.106)	0.302*** (0.059)	0.275*** (0.068)
Log Value/Area	0.481*** (0.084)	0.469*** (0.089)	0.338*** (0.086)	0.358*** (0.107)	0.429*** (0.079)	0.426*** (0.087)
Survey FE	Yes	Yes	Yes	Yes	Yes	Yes
Clusters	174	174	174	174	174	174
Surveys	4	4	4	4	4	4
Sample	Owners	Owners	Owners	Owners	Owners	Owners
Observations	35,074	35,376	34,724	35,376	34,913	35,376
Adj. R^2	0.609	0.384	0.629	0.384	0.424	0.266
Within R^2	0.2439	0.0869	0.2227	0.0785	0.1445	0.0504
Mean dep. var.	10.893	11.486	10.568	11.053	9.593	10.151

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Agricultural landowners in the four wealth surveys (Ethiopia 2018, India 2019, Malawi IHPS 2016, Tanzania 2019), on log land area and log value per hectare. **Total Assets** = dwelling + plot + other; **Productive Assets** = plot + other (excl. dwelling); **Non-Plot Assets** = dwelling + other (excl. plot). Within each measure, ln restricts to strictly positive wealth; asinh retains zero-wealth households. Survey FE, probability weights renormalized so that each survey carries equal total weight, standard errors clustered at the survey-strata level.

(Figure B5). Overall, the country-level estimates are in line with the pooled results.


C Appendix to Section 4: Land Wealth Inequality

C.1 Sampling uncertainty and robustness of Figure 1

The confidence intervals in Panels (b) and (c) of Figure 1 are obtained from 999 survey-design bootstrap replications. Within each country, primary sampling units are resampled with replacement within sampling strata, retaining the original survey weights. For Indonesia and Nicaragua, for which primary sampling-unit identifiers are unavailable, households are instead resampled within strata; observations without usable primary sampling-unit identifiers in the other surveys are treated similarly. Singleton strata are held fixed across replications. All inequality concepts and paired changes are recalculated jointly in each replication, and the reported 95 percent confidence intervals correspond to the 2.5th and 97.5th percentiles of the resulting bootstrap distributions. These intervals capture survey-sampling uncertainty conditional on the countries included in the analysis.

Table C15 examines the robustness of the aggregate results in Figure 1. On the common

Figure B7: Effects of Land and Value on Wealth by Country



Notes: Each point shows the OLS coefficient from regressing $\ln(\text{non-plot wealth})$ on log land area and log value per hectare within a single wealth survey, using probability weights and the survey design. Non-plot wealth is dwelling value plus other assets, excluding the agricultural plot. Horizontal bars: 95% confidence intervals. Countries sorted by $\ln(\text{Value/Area})$ coefficient (descending). Open triangle ($\triangle$) = $\ln(\text{Value/Area})$; open diamond ($\diamond$) = $\ln(\text{Area})$.

sixteen-country sample, the mean Gini's change by at most 0.004 and the reduction in cross-country dispersion is 35.3 percent (95 percent design-bootstrap confidence interval: 29.3–42.0 percent). Trimming 0.5 percent from each tail preserves the ordering of the inequality concepts and produces a 37.9 percent reduction in dispersion. Across the leave-one-out estimates, this reduction ranges from 31.4 to 41.2 percent. At the country level, the 95 percent design-bootstrap confidence interval for the paired area-to-value change excludes zero in 12 of the 17 countries: land-value inequality is significantly higher in nine and significantly lower in three. Complete country estimates and confidence intervals are provided in the replication files.

C.2 Land inequality excluding non-income-poor landless households

This appendix replicates Figure 1 using a narrower definition of landless households. Instead of including all landless agricultural households, we include only those whose per-capita daily income is below US\$2.15 PPP. Figure C8 reports the results. Under this definition, the average Gini remains high, at 0.74, compared with 0.77 in the baseline specification that includes all landless households. The cross-country standard deviation is 0.082, roughly 40 percent below the land-area standard deviation of 0.135 in this sixteen-country sample and below the 0.11 standard deviation of income inequality. India is excluded because income data are not available.

Table C15: **Robustness of the Figure 1 aggregates**

Statistic	Baseline	Common 16	Trim 0.5%	LOO range
Mean Gini: area, owners	0.651	0.655	0.617	[0.635, 0.663]
Mean Gini: value, owners	0.679	0.679	0.643	[0.666, 0.690]
Mean Gini: value and landless	0.767	0.766	0.737	[0.758, 0.777]
SD: area, owners	0.132	0.135	0.129	[0.119, 0.137]
SD: value, owners	0.092	0.095	0.083	[0.076, 0.095]
SD: value and landless	0.085	0.088	0.080	[0.078, 0.088]
Reduction in SD	35.8	35.3	37.9	[31.4, 41.2]

Notes: Baseline land estimates use all 17 countries, whereas income inequality is available for 16. Common 16 excludes India from the land aggregates. Trim 0.5% removes 0.5 percent from each tail of the relevant distribution. The leave-one-out range reports the minimum and maximum obtained by omitting each country in turn. The 95% design-bootstrap confidence interval for the baseline reduction in the standard deviation is 30.0–42.3%; for the common sixteen-country sample, it is 29.3–42.0%.

C.3 Excluding Countries without Directly Self-Reported Market Values

Our benchmark measure of land value is the household’s own report of the current market value of its plots. Four countries use alternative valuation methods: in China (2012), the survey provides a household-level valuation; in Ethiopia (2018), households report rental rather than market values; and in Brazil (1997) and Nicaragua (2014), land values are imputed from regional land prices—drawn from an external official source in Brazil and supplied within the survey in Nicaragua (Appendix A.1.2 and Table A2). We replicate the main inequality results after excluding all four countries. Figure C9 reports the cross-country mean and standard deviation of the Gini under the measurement concepts used in Figure 1, together with the area-to-value relationship and its decomposition from Panels (a) and (b) of Figure 2. The results remain close to those for the full sample.

C.4 Elasticities, Area Inequality, and the Covariance Decomposition


This appendix provides additional evidence on the inverse value–size relationship and clarifies how it shapes the covariance term in Section 4. The analysis is motivated by the variance identity

$$\text{Var}(\log V) - \text{Var}(\log A) = \text{Var}(\log P) + 2 \text{Cov}(\log A, \log P),$$


which shows that differences between land-value and land-area inequality arise from two components: price dispersion, $\text{Var}(\log P)$, and the covariance between land size and land prices. The covariance term can itself be written as $\text{Cov}(\log A, \log P) = \beta \text{Var}(\log A)$, where β is the elasticity of land prices per hectare with respect to size.

Elasticities across countries. Figure C10 reports the elasticity of land prices per hectare with respect to landholding size (β) for each country. In all cases, the elasticity is negative, indicating that larger landholdings command lower prices per hectare. While its magnitude varies across


Figure C8: **Land Inequality Across Measures: Income-Poor Landless Definition**



(a) Gini index across measurement concepts by country



(b) Cross-country mean Gini by concept




(c) Cross-country standard deviation of the Gini by concept


Notes: Panel (a) reports country-level Gini coefficients for land inequality under a sequence of measurement concepts. The blue segment shows the change in the Gini when moving from land *area* among landowners to land *value* among landowners. The red segment shows the change when moving from land value among landowners to land value including income-poor landless households, who are assigned zero land value. Panels (b) and (c) summarize the cross-country distribution of Gini coefficients by concept: panel (b) plots the unweighted cross-country mean, and panel (c) plots the unweighted cross-country standard deviation. Income-poor landless households are defined as landless agricultural households with per-capita daily income below US\$2.15 PPP. India is excluded because income data are not available.

countries, the inverse price gradient is a consistent feature of the data.


Figure C9: Inequality results excluding countries without directly self-reported market values




(a) Cross-country mean Gini by concept



(b) Cross-country standard deviation by concept



(c) Area inequality and the change from area to value



(d) Decomposing the change in inequality

Notes: Brazil, China, Ethiopia, and Nicaragua are excluded. Panels (a) and (b) report the unweighted cross-country mean and standard deviation of country-level Gini coefficients under each measurement concept. Land statistics are computed across 13 countries; income statistics exclude India and are computed across 12 countries. Panel (c) plots $\text{Var}(\log V) - \text{Var}(\log A)$ against $\text{Var}(\log A)$ among landowners. Panel (d) decomposes the change into $\text{Var}(\log P) + 2\text{Cov}(\log P, \log A)$, separately for countries where the change is negative and positive. Panels (c) and (d) use the same 13-country sample. V , A , and P denote land value, land area, and value per hectare, respectively.


The covariance term and land-area inequality. Panel (a) of Figure C11 plots the covariance between land size and prices per hectare against initial land-area inequality, measured by $\text{Var}(\log A)$. The covariance becomes more negative as land-area dispersion increases. Thus, the countries with

the most unequal distributions of land area also experience the largest compression in dispersion when landholdings are measured by value rather than area.

Elasticity and initial land-area inequality. Panel (b) plots country-level elasticities against $\text{Var}(\log A)$. Although elasticities are uniformly negative, there is no strong systematic relationship between the steepness of the inverse gradient and area inequality. Countries with highly unequal land distributions do not exhibit systematically more negative elasticities. Cross-country differences in the covariance term are therefore unlikely to be driven primarily by variation in the slope of the inverse value–size relationship.

Price dispersion and land-area inequality. Panel (c) plots $\text{Var}(\log P)$ against $\text{Var}(\log A)$. While price dispersion is positively related to land-area inequality, the relationship is comparatively weak and considerably noisier than the association between the covariance term and land-area inequality shown in panel (a). Variation in price dispersion alone therefore cannot account for the strong equalizing effect observed when moving from area- to value-based inequality measures.


Figure C10: Elasticity of land prices with respect to land size




Notes: The figure reports country-level estimates of the elasticity of log land prices per hectare with respect to log land size. Negative estimates indicate that larger landholdings have lower prices per hectare.

Decomposing the covariance term. Given that $\text{Cov}(\log A, \log P) = \beta \text{Var}(\log A)$, cross-country differences in the covariance reflect differences in both elasticities and land-area dispersion. Table C16 implements an exact Shapley decomposition of the between-group difference in the covariance effect shown in Figure 2b in the main paper. The results show that the more negative covariance in the descending group is driven primarily by higher dispersion in landholdings, with a secondary contribution from slightly more negative elasticities. Differences in the cross-country covariance between β and $\text{Var}(\log A)$ partially offset these mean effects. Overall, dispersion in land size accounts for the majority of the difference in the covariance effect across groups.


Figure C11: **The covariance decomposition and land-area inequality**



(a) Covariance and land-area inequality



(b) Elasticity and land-area inequality



(c) Price dispersion and land-area inequality

Notes: Panel (a) plots the covariance between log land area and log land prices per hectare against land-area inequality, measured by $\text{Var}(\log A)$. Panel (b) plots the elasticity of log land prices per hectare with respect to log land size against $\text{Var}(\log A)$. Panel (c) plots price dispersion, $\text{Var}(\log P)$, against $\text{Var}(\log A)$. Each observation represents a country.

Is the inverse value–size relationship driven by productivity differences? An important remaining question is whether the negative elasticities documented above simply reflect the well-known inverse relationship between farm size and productivity. If smaller farms are more pro-

Table C16: Shapley decomposition of the covariance effect

	Descending	Ascending
Mean elasticity (β)	-0.5864	-0.3839
Mean $\text{Var}(\log A)$	4.2810	1.5992
Mean $2\text{Cov}(\log A, \log P)$	-3.9983	-1.1718
<i>Between-group difference (Ascending – Descending)</i>		
Total difference	2.8266	
Elasticity component	1.1908	
Area inequality component	2.6024	
$\text{Cov}(\beta, \text{Var}A)$ component	-0.9666	

Notes: Countries are grouped into *descending* (where $\text{Var}(\log V) - \text{Var}(\log A) < 0$) and *ascending* (where the change is positive). β denotes the elasticity of log land prices with respect to log land size, and $\text{Var}(\log A)$ is initial land-area inequality. The covariance term entering the variance identity can be written as $2\beta \text{Var}(\log A)$. The lower panel reports an exact Shapley decomposition of the between-group difference in the mean covariance effect. The elasticity component captures differences in mean elasticities holding dispersion fixed at the mid-point; the area inequality component captures differences in mean land-area dispersion holding elasticities fixed; and the $\text{Cov}(\beta, \text{Var}A)$ component reflects cross-country comovement between elasticities and dispersion within groups.

ductive, they may command higher land values per hectare, mechanically generating a negative elasticity of land prices with respect to land area.

Table C17 evaluates this channel in two steps. First, Column 2 restricts the sample to households with strictly positive agricultural income (Drop 0), ensuring comparability with specifications that incorporate productivity controls. The resulting elasticities are very similar to those obtained in Column 1 using the full sample, indicating that the sample restriction itself does not materially affect the inverse value–size relationship.

Second, Column 3 adds realized farm income per hectare as a proxy for current agricultural productivity.¹⁸ Controlling for measured productivity makes the inverse elasticity less negative, but only partially. The unweighted mean elasticity changes from -0.4584 in Column 2 to -0.3708 in Column 3, a reduction in magnitude of roughly 19 percent. In the vast majority of countries, the coefficient remains negative and economically meaningful.

We do not interpret this approximately 19 percent attenuation as a decomposition of the share of the relationship explained by productivity. Realized farm income per hectare is an imperfect and potentially endogenous proxy for productivity, may be affected by measurement error and transitory shocks, and shares land area in its denominator with value per hectare. Nevertheless, the

¹⁸Agricultural income is defined as farm income plus farm wage income plus land-rent income, as in Appendix A.1.3. Farm income can, in turn, be decomposed into crop and animal income. Our results are nearly identical when we use agricultural income or crop income instead of farm income.

Table C17: Country-level elasticity robustness checks

Country	Elasticity estimates				Sample size (N)	
	Baseline	Drop 0	Control prod.	SE	Baseline	Drop 0
Bangladesh	-0.1230	-0.1109	-0.0937	0.0173	3049	1623
Brazil	-0.0593	0.0746	0.1158	0.0410	452	283
China	-0.5196	-0.5987	-0.1581	0.0079	8218	5787
Ecuador	-0.6348	-0.5989	-0.4446	0.0113	8812	5614
Ethiopia	-0.4600	-0.5022	-0.5207	0.0411	1535	935
Gambia	-0.4647	-0.4740	-0.4412	0.0195	7306	4060
Guatemala	-0.5769	-0.5951	-0.4851	0.0296	1931	944
India	-0.4072	.	.	.	43602	0
Indonesia	-0.7077	-0.6913	-0.5252	0.0195	3901	2723
Malawi	-0.4245	-0.3747	-0.3353	0.0216	7035	3231
Nicaragua	-0.0017	-0.0187	0.0128	0.0288	631	548
Niger	-0.5879	-0.6919	-0.7714	0.0776	1098	235
Nigeria	-0.9539	-0.9191	-0.8860	0.0322	2220	893
Pakistan	-0.2834	-0.2533	-0.0858	0.0248	3698	2078
Panama	-0.5846	-0.4913	-0.3583	0.0243	2080	705
Tanzania	-0.6757	-0.6211	-0.5321	0.0329	1616	1100
Vietnam	-0.4826	-0.4699	-0.4235	0.0204	5664	4045
Unweighted mean	-0.4675	-0.4585	-0.3708	0.0281		

Notes: Entries report country-level elasticities of log land prices per hectare with respect to log landholding size. Baseline estimates use all landowning households. Drop 0 restricts the sample to households with strictly positive agricultural income. Control prod. additionally controls for agricultural productivity, measured as agricultural income per hectare. All specifications are estimated using household sampling weights. SE reports the standard error of the elasticity in the productivity-controlled specification. The controlled regression is estimated on the Drop-0 sample (so its number of observations equals the Drop-0 column by construction). The final row reports the unweighted average across countries.


results suggest that measured productivity captures only part of the inverse value–size association.

C.5 Land use composition and land-area dispersion

This appendix provides suggestive evidence on whether cross-country differences in agricultural land use are correlated with the dispersion of landholdings. The motivation is the variance identity in Section 4.2, where the covariance term depends both on the elasticity of land prices with respect to land size and on the dispersion of land sizes. Using FAO country-level data aligned with the survey dates, we construct the share of arable land in total agricultural land.

Figure C12 shows that countries with a higher arable share tend to have lower land-area inequality, measured by $\text{Var}(\log A)$, and a less negative covariance between land size and land prices, $\text{Cov}(\log A, \log P)$. By contrast, arable share exhibits no systematic relationship with the elasticity of land prices with respect to land size (not shown). Taken together, these patterns suggest that land-use structure is primarily associated with cross-country differences in land-area dispersion, and thus with the magnitude of the covariance effect, rather than with differences in the slope of the inverse value–size gradient itself.

Figure C12: Land use composition and the covariance mechanism



Notes: The horizontal axis reports the FAO share of arable land in total agricultural land. Panel (a) plots land-area inequality $\text{Var}(\log A)$. Panel (b) plots the covariance between land size and land prices, $\text{Cov}(\log A, \log P)$. Each dot is a country.


C.6 Crop versus Animal Income Mix and Land Area Dispersion

This appendix provides household-level evidence on the association between agricultural land use and the dispersion of landholdings documented in Appendix C.5. We exploit within-country variation in production patterns by classifying landowning households according to the share of net crop income in total farm income, which serves as a micro-level analog of the FAO arable share.

We group households into four categories: pure animal (crop share = 0; 24.7% of the sample), animal-dominant (crop share in $(0, 0.5]$; 8.3%), crop-dominant (crop share in $(0.5, 1)$; 14.5%), and pure crop (crop share = 1; 52.6%). The sample contains 33,457 landowning households with positive farm income across 14 countries.

Panel (a) of Figure C13 plots the variance of log land area, $\text{Var}(\log A)$, by crop-income group. Dispersion declines sharply with crop intensity: $\text{Var}(\log A)$ falls from 3.28 among pure-animal households to 1.49 among pure-crop households, a decline of more than 50%. By contrast, average log land area varies little and shows no monotonic relationship with crop intensity. This contrast indicates that the crop–animal distinction operates through who occupies the upper tail of the farm-size distribution, rather than through shifts in typical farm size.

Figure C13: Crop–animal income mix and the land-inequality mechanism



Notes: The sample includes landowning households in 14 countries with positive crop or animal income. Households are grouped by the share of net crop income in total farm income: pure animal (crop share = 0), animal-dominant (crop share in $(0, 0.5]$), crop-dominant (crop share in $(0.5, 1)$), and pure crop (crop share = 1). Panel (a) plots the weighted variance of $\log A$ within each group. Panel (b) plots the coefficient on $\log A$ from separate weighted OLS regressions of $\log P$ on $\log A$ within each group, controlling for country fixed effects; whiskers denote 95% confidence intervals based on standard errors clustered by country.

Table C18 provides complementary evidence using continuous crop-share specifications. Regressing the squared deviation of log area from its country mean, $(\log A - \overline{\log A_c})^2$, on crop share and country fixed effects yields a large negative coefficient. By contrast, crop share is essentially unrelated to mean log holding size. The effect is concentrated in the upper tail: higher crop share lowers the probability of being in the top decile of the within-country land-area distribution, but has no meaningful effect on the probability of being above the median. A formal equality test rejects equality between the top-decile and median effects. These results indicate that animal-oriented production is associated with very large landholdings at the top of the distribution, rather than with a general rightward shift of the land-size distribution.

Panel (b) of Figure C13 reports the elasticity of land value per hectare with respect to land size, estimated separately within each crop-income group. The elasticities are negative in all groups, with overlapping confidence intervals and no monotonic pattern by crop orientation. The inverse value–size gradient is therefore present across land-use groups and does not vary

Table C18: Land use, land-area dispersion, and the upper tail

Outcome	Estimate	Interpretation
$\log A$	0.017 $p = 0.89$	Crop share does not shift mean land size
$(\log A - \overline{\log A_c})^2$	-0.492 $p < 0.01$	Crop share reduces within-country dispersion
Top decile of $\log A$	-0.0435 $p = 0.02$	Crop share lowers upper-tail probability
Above median of $\log A$	0.0171 $p = 0.67$	No comparable median effect
Equality test	$F(1, 13) = 12.53$ $p < 0.01$	Upper-tail effect differs from median effect

Notes: The sample consists of landowning households with positive crop or animal income in the 14 countries for which crop and animal income can be separated ($N = 33,457$). All specifications include country fixed effects and use survey weights. The first two rows report weighted OLS coefficients on crop-income share, with standard errors clustered by country. The top-decile and above-median rows report average marginal effects from logit models, with standard errors clustered by country. The equality test is from a stacked weighted linear probability model with country fixed effects interacted by equation and standard errors clustered by country. Crop-group shares in this sample are: pure animal 24.7%, animal-dominant 8.3%, crop-dominant 14.5%, and pure crop 52.6%.

systematically with the crop–animal mix. This supports the interpretation that land-use composition shapes land inequality primarily through dispersion in landholding sizes rather than through changes in the slope of the value gradient. This is also in line with the Shapley decomposition in Table C16, which attributes cross-country differences in the covariance term mainly to differences in $\text{Var}(\log A)$ rather than to differences in the elasticity.

Overall, the household-level evidence links animal-oriented production to greater land-size dispersion concentrated in the upper tail. This complements the country-level FAO evidence in Appendix C.5, where lower arable shares are associated with higher $\text{Var}(\log A)$ and a more negative covariance term. Because average crop-income shares are essentially uncorrelated with country-level $\text{Var}(\log A)$ ($r = -0.02$), we view the two exercises as complementary suggestive evidence rather than a direct aggregation mechanism.


C.7 Gini-Based Area-to-Value Adjustment

Figure C14 provides a Gini-based counterpart to the log-variance evidence in Section 4.2. Countries with higher area-based land inequality experience larger downward adjustments when land is measured by value rather than area. This supports the interpretation that the compression documented in Section 4.1 is not specific to the log-variance metric used for the decomposition.

C.8 Land as a Share of Household Wealth

Figure C15 reports the survey-weighted share of household wealth held in agricultural plots in the four wealth surveys. The left panel expresses plot wealth as a share of total household wealth


Figure C14: **Area-to-value adjustment in Gini space**



Notes: The figure plots the country-level change in measured land inequality when moving from land area to land value, defined as the land-value Gini minus the land-area Gini among landowners, against the land-area Gini among landowners. Each dot is a country. The fitted line is from an unweighted OLS regression. The figure provides a Gini-based counterpart to the log-variance adjustment used in Section 4.2.

(dwelling, plots, and other assets); the right panel as a share of productive assets, excluding the dwelling. The numerator is the wealth module's own plot valuation rather than the self-reported land transaction value: in the three African surveys the two measures come from separate survey instruments with little overlap, and using the transaction value frequently produces shares above one; India is the only survey where both originate from the same module and are directly comparable. The sample comprises agricultural landowners in the four wealth surveys and is broader than the wealth-regression sample, which additionally applies the land-value consistency screens and requires positive non-plot wealth. Plots account for a substantial share of total wealth and for the majority of productive assets in every country except Ethiopia. Ethiopia stands out as an outlier here, with a markedly lower plot-wealth share than the other countries; this likely reflects its status as the most important cattle-grazing country in Africa—and among the top five in the world—such that a large share of household wealth, on average 44%, is held in livestock rather than in land.

Figure C15: Plot Wealth as a Share of Total Wealth, by Country



Notes: Each bar shows the survey-weighted mean share of household wealth attributable to plot wealth (from the wealth survey module), with 95% confidence intervals. Countries sorted ascending by share within each panel. *Left panel*: plot wealth as a fraction of total wealth (dwelling + plots + other assets). *Right panel*: plot wealth as a fraction of plot wealth + other assets (excluding the dwelling). Sample: agricultural landowners in the four wealth surveys.

C.9 Land Value versus Land Area at the Bottom

Figure 3, Panel (a), shows that ranking landowners by value identifies a poorer lower-tail group than ranking them by area, while holding the weighted size of the two groups fixed. This appendix examines the robustness and cross-survey consistency of that result.

Table C19 shows similar results when the lower tail is defined as the bottom 10 or 25 percent. The bottom-quintile difference also rises to 9.7 percentage points when we retain only the twelve surveys with directly household-reported land values. The result therefore does not depend on the size of the lower-tail group or on the alternative valuation methods used in China, Ethiopia, Brazil, and Nicaragua.

Table C19: Income Poverty among Lower-Tail Groups

Definition	All bottom-group households			Non-overlapping households			
	Value	Area	Difference	Value only	Area only	Difference	Overlap
Baseline, bottom 10%	59.9	51.0	8.9 [3.9, 13.8]	61.0	47.7	13.2 [6.0, 20.5]	35.9
Baseline, bottom 20%	59.5	51.7	7.8 [3.8, 11.7]	63.4	47.6	15.8 [7.4, 24.2]	46.9
Baseline, bottom 25%	59.4	52.6	6.7 [3.3, 10.2]	63.3	49.7	13.6 [6.7, 20.5]	51.2
Direct-value sample, bottom 20%	60.7	51.0	9.7 [5.1, 14.3]	65.0	48.1	16.9 [8.3, 25.6]	41.6

Notes: Rates and overlap are percentages; differences are percentage points. Groups are constructed separately within each survey, and value and area groups have exactly the same weighted size. Surveys receive equal weight. Difference confidence intervals are calculated from survey-specific paired differences using a t distribution with survey-level degrees of freedom. The direct-value row excludes China, Ethiopia, Brazil, and Nicaragua. Bottom 20% is the main specification; bottom 10% and bottom 25% are prespecified robustness checks. The leave-one-survey-out range for the baseline bottom-20% full-group difference is 6.8 to 8.5 p.p.

Overall, land value changes who is identified at the bottom and selects households that are poorer in income terms. This provides a threshold-free validation of value as a more economically

informative measure of land endowments than area.

C.10 Income Poverty and Land Poverty: Alternative Specifications

This appendix provides supporting evidence for the land-poverty analysis in Section 4.3. We first compare the US\$1,000 PPP value benchmark with the commonly used two-hectare cutoff. We then report income-poverty levels across landholding groups and examine robustness to alternative land-value thresholds. Income poverty is defined using the international US\$2.15 PPP per-day line. All regressions use household survey weights, with standard errors clustered at the survey level unless otherwise indicated.

Value benchmark versus the two-hectare cutoff. Table C20 compares the income-poverty rates among landowners selected by the US\$1,000 PPP value benchmark and by the two-hectare cutoff. The adjusted income-poverty rate is 62.9 percent under the value benchmark and 54.3 percent under the area cutoff, a difference of 8.5 percentage points (95 percent confidence interval: 3.6–13.4). The value benchmark selects 11.9 percent of landowners, compared with 66.1 percent under the two-hectare cutoff. Because the two rules select groups of different sizes, this comparison complements the equal-sized lower-tail comparison in Figure 3, Panel (a).

Table C20: **Income poverty among landowners selected by value and area thresholds**

	Adjusted income-poverty rate (%)	Share of landowners selected (%)
Land value below US\$1,000 PPP	62.9	11.9
Land area below two hectares	54.3	66.1
Difference: value minus area	8.5	
95% confidence interval	[3.6, 13.4]	

Notes: The sample consists of agricultural landowners in the 16 countries with income data. Income poverty is defined as per-capita daily income below US\$2.15 PPP. Adjusted poverty rates are marginal predictions from a linear probability model estimated on stacked rule-specific samples, with country fixed effects. Household survey weights are normalized within each country–rule cell so that every country receives equal weight under both rules. Standard errors are clustered by country. Households satisfying both thresholds appear in both rule-specific samples. The reported difference is the value-threshold rate minus the area-threshold rate; its p -value is 0.002. The final column reports the equal-country average of the within-country share of landowners selected by each rule.

Pooled poverty levels. We begin with a linear probability model without survey fixed effects:

$$y_{is} = \alpha + \beta \mathbf{1}\{\text{land-poor owner}_{is}\} + \gamma \mathbf{1}\{\text{non-land-poor owner}_{is}\} + u_{is},$$

where landless households are the omitted group. The coefficient α is therefore the poverty rate among landless households, while $\alpha + \beta$ and $\alpha + \gamma$ give the rates among land-poor and non-poor landowners, respectively.

Panel (a) of Figure C16 reports these levels at the benchmark US\$1,000 PPP threshold. The pattern is non-monotonic: land-poor landowners have the highest income-poverty rate, followed by landless households, while non-poor landowners have the lowest rate. The corresponding rates are approximately 60, 55, and 45 percent.

Within-survey differences. To absorb differences in average poverty rates across surveys, we estimate

$$y_{is} = \alpha_s + \beta \mathbf{1}\{\text{land-poor owner}_{is}\} + \gamma \mathbf{1}\{\text{non-land-poor owner}_{is}\} + u_{is},$$

where α_s denotes survey fixed effects. The coefficients β and γ measure the within-survey income-poverty gaps between each owner group and landless households. Panel (b) of Figure C16 reports these estimates at the US\$1,000 PPP benchmark. Land-poor landowners are significantly more likely to be income poor than the landless, whereas non-poor landowners are significantly less likely. Thus, although landless households have no agricultural land wealth, their income-poverty rate lies between those of below- and above-benchmark owners.

Alternative land-value thresholds. Panels (c) and (d) of Figure C16 re-estimate the fixed-effects specification using land-value thresholds of US\$500 and US\$1,500 PPP, respectively. Both panels report the within-survey differences for land-poor and non-poor landowners relative to landless households. The qualitative ordering is unchanged: land-poor landowners are more likely to be income poor than the landless, whereas non-poor landowners are less likely to be income poor. The relationship between land poverty and income poverty is therefore robust to reasonable variation in the land-value benchmark.

C.11 Relative Land Inequality and Land-Poverty Incidence


Figure C17 compares the Gini coefficient of land value across all agricultural households, assigning zero land value to landless households, with the incidence of land poverty across countries. Panel (a) uses the baseline definition in the main text: all landless agricultural households are classified as land poor, together with owners below US\$1,000 PPP. The association is modestly positive ($R^2 = 0.18$), partly because a higher landless share raises both the Gini and the headcount. Panel (b) instead counts landless households as land poor only when they are also income poor. The fitted relationship is then essentially flat ($R^2 = 0.00$). Under either definition, similar levels of relative land-value inequality coexist with substantially different levels of absolute land deprivation.

C.12 Robustness: Land-Poverty Incidence and Value Shortfall

This subsection assesses the robustness of the land-poverty headcount and aggregate land-value shortfall reported in Section 4.3 along three dimensions: the land-value threshold, the treatment of landless households, and the capitalization factor applied to Ethiopia's annual rental values.

Figure C18, Panels (a) and (b), retains the baseline definition of land poverty—under which all landless agricultural households are classified as land poor—but varies the land-value threshold to US\$500 and US\$1,500 PPP. Figure C19 retains the benchmark US\$1,000 PPP threshold but

Figure C16: Income poverty by landholding group: within-survey gaps and robustness to alternative thresholds




Notes: Panel (a) reports predicted income-poverty rates for landless households, land-poor landowners, and non-poor landowners from a pooled three-group specification without survey fixed effects at the benchmark US\$1,000 PPP threshold. Panel (b) reports within-survey differences in income poverty for land-poor and non-poor landowners relative to landless households at the same threshold. Panels (c) and (d) report the corresponding within-survey differences using thresholds of US\$500 and US\$1,500 PPP, respectively. The landless are the omitted group in Panels (b)–(d). All regressions use household survey weights and cluster standard errors at the survey level. Whiskers denote 95% confidence intervals. Income poverty is defined as per-capita daily income below US\$2.15 PPP.

uses the income-screened definition, under which landless households are classified as land poor only when they are also income poor.

In each figure, the bars report the land-poverty headcount and the aggregate land-value shortfall, with both measures decomposed between landowners and landless households. The headcount measures the share of agricultural households classified as land poor, while the shortfall measures the value needed to bring these households up to the relevant threshold, expressed as a share of aggregate land value.

As expected, raising the land-value threshold increases both the incidence of land poverty and the aggregate value shortfall, while lowering it reduces both measures. Screening landless households by income at the US\$1,000 PPP threshold reduces their contribution to the headcount and shortfall. Across these alternatives, however, the main conclusion is unchanged: land poverty

Figure C17: Land-value inequality and land-poverty incidence



(a) Baseline: all landless households classified as land poor

(b) Income-screened landless definition


Notes: Both panels plot the Gini coefficient of land value across all agricultural households, assigning zero land value to landless households. In Panel (a), land-poor households are defined as landowners holding land worth less than US\$1,000 PPP together with all landless agricultural households. In Panel (b), landless households are classified as land poor only if their per-capita daily income is below US\$2.15 PPP. Each dot represents a country, and the fitted lines are from unweighted OLS regressions. To hold the country sample fixed across definitions, both panels use the same sixteen countries; India is excluded because income is unavailable for Panel (b). The R^2 is reported in each panel.

is widespread in some countries, while the aggregate land-value shortfall remains small relative to total land wealth.


Capitalization of Ethiopia's annual rental values. Ethiopia reports annual rental values rather than observed sale values. Our benchmark capitalizes these rents using a factor of 10, corresponding to a 10% capitalization rate. We regard this as a transparent central assumption rather than a uniquely correct rate, particularly because rural land rights in Ethiopia are comparatively illiquid and subject to restrictions on transfer. Table C21 therefore also reports results using a 5% rate, corresponding to a factor of 20, and a 15% rate, corresponding to a factor of 6.67. At the US\$1,000 PPP threshold, the benchmark incidence and shortfall are 39.33% and 1.457%, respectively. They are 37.14% and 0.711% under the factor of 20, and 41.68% and 2.244% under the factor of 6.67.

Across the alternative thresholds, treatments of landless households, and Ethiopia capitalization factors, the main conclusion is unchanged: land poverty is widespread in some countries, while the aggregate land-value shortfall generally remains small relative to total land wealth.

Figure C18: **Land-poverty incidence and value shortfall: alternative thresholds**



(a) Threshold = US\$500 PPP




(b) Threshold = US\$1,500 PPP

Notes: Each panel reports, by country, (i) the share of agricultural households below the relevant land-value threshold and (ii) the aggregate land value needed to bring all land-poor households up to that threshold, expressed as a share of aggregate land value (“Value shortfall”). Bars are stacked by household type (landowners versus landless households). All landless agricultural households are assigned zero land value and therefore fall below each threshold. Panels (a) and (b) report results for land-value thresholds of US\$500 and US\$1,500 PPP, respectively. Both panels cover all seventeen countries.

C.13 Landless Households

Table C22 regresses log income, log non-plot wealth, and log outstanding loans on a landowner indicator, with survey fixed effects. The income and wealth regressions include all agricultural

Figure C19: **Land-poverty incidence and value shortfall: income-screened landless definition**



Notes: The figure reports, by country, (i) the share of agricultural households classified as land poor and (ii) the aggregate land value needed to bring all land-poor households up to the US\$1,000 PPP threshold, expressed as a share of aggregate land value (“Value shortfall”). Bars are stacked by household type (landowners versus landless households). Under this income-screened definition, landowners are classified as land poor when their land value is below US\$1,000 PPP, while landless households are classified as land poor only if their per-capita daily income is below US\$2.15 PPP. India is excluded because income is unavailable.

Table C21: **Ethiopia: Sensitivity of land poverty to the capitalization factor**

Capitalization rate	Factor	Annual-rent cutoff	Incidence (%)	Value shortfall (%)
10% (benchmark)	10.00	100.00	39.33	1.457
5%	20.00	50.00	37.14	0.711
15%	6.67	149.93	41.68	2.244

Notes: The sample consists of benchmark agricultural households in Ethiopia. Capitalized land value equals the reported annual rental value multiplied by the indicated factor. The land-poverty threshold is US\$1,000 PPP, and all landless households are assigned zero land value and classified as land poor. Incidence is the weighted share of agricultural households below the threshold. The value shortfall is the aggregate amount required to bring land-poor households to the threshold, expressed as a share of aggregate capitalized land value. The factor 6.67 approximates the reciprocal of a 15% capitalization rate.

households (landowners vs. landless); the loan regression is estimated among borrowers, so its coefficient compares loan sizes rather than the probability of borrowing. Among borrowers, owners hold significantly larger loans, but the premium to ownership is statistically insignificant for both income and non-plot wealth. Figure C20 examines whether this mean effect masks variation across the conditional income distribution by estimating quantile regressions of log equivalized income on the landowner indicator at the 10th, 25th, 50th, 75th, and 90th percentiles. All regressions use the same sample restrictions as the main table (agricultural households, zero-income households

dropped, income winsorized at the 99th percentile).

Table C22: Landowner Premium: Income, Non-Plot Wealth, and Loans


	Income (1)	Non-Plot Wealth (2)	Loans (3)
Landowner	0.027 (0.039)	0.196 (0.142)	0.361*** (0.086)
Survey FE	Yes	Yes	Yes
Clusters	447	176	307
Surveys	16	4	11
Sample	Agric	Agric	Borrowers
Observations	67,408	60,171	11,703
Adj. R^2	0.563	0.269	0.584
Within R^2	0.0001	0.0026	0.0130

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Each column regresses the log outcome on an indicator for owning agricultural land, with survey fixed effects and probability weights renormalized so that each survey carries equal total weight; standard errors clustered at the survey-strata level. Columns (1)–(2) include all agricultural households (landowners vs. landless). Column (3)’s outcome is log outstanding loans, so it is estimated among borrowers: the coefficient compares loan size between owner and landless borrowers, not the probability of borrowing. These are the landowner-dummy regressions removed from the main table.

The ownership coefficient increases monotonically across the income distribution: it is negative and significant at the 10th percentile, negative but insignificant at the 25th percentile, positive and significant at the median, and increases in magnitude while remaining significant at the 75th and 90th percentiles. The small and statistically insignificant mean effect in Table C22 therefore masks substantial heterogeneity: land ownership is associated with *lower* income at the bottom of the distribution and *higher* income at the top. This pattern is consistent with the interpretation that land ownership per se does not guarantee higher income; what matters is the *value* of the land owned, which is concentrated among higher-income agricultural households.


The pooled premium also averages over very different countries. Figure C21 estimates the owner premium separately for each of the 16 income surveys, for two outcomes: log equivalized income (maroon circles) and the share of agricultural income in total household income (navy triangles), with 95 percent confidence intervals. Both coefficients vary widely across countries. In some surveys the landless earn as much as or more than owners and derive far less of their income from agriculture; in others the ranking reverses. These offsetting country patterns contribute to the small pooled income premium. A lower agricultural income share reflects more income from non-agricultural sources—that is, a more diversified income mix. Suggestively, the countries toward the top of the figure—where ownership is associated with a markedly more agricultural income mix—tend to be economies in which agriculture accounts for a large share of value added, while those toward the bottom are less agriculture-dependent; it is thus possible that income diversification is more limited in the agriculture-dependent economies, and that the position of the landless partly reflects how households sort between agricultural and non-agricultural activities in more and less agrarian economies.

Figure C20: Quantile regression coefficients: landowner indicator on log income



Notes: Each point shows the coefficient on the landowner dummy from a quantile regression of log equivalized income at the indicated percentile, with survey fixed effects (as indicator dummies). Vertical bars show 90% confidence intervals. The regressions are unweighted: estimating quantile regressions with probability weights and high-dimensional fixed effects is not feasible. Sample: agricultural households; same cleaning as the main table.

Figure C21: Landowner–Landless Gap by Country: Income and Agricultural Share



Notes: Each point shows the survey-weighted OLS coefficient on the landowner dummy from a country-level regression of the outcome on a landowner indicator. Maroon circles: $\ln(\text{Income})$, equalized using the OECD-modified scale (China: per capita). Navy triangles: share of agricultural income in total household income (non-equalized). Horizontal bars show 95% confidence intervals. Sample: agricultural households; same cleaning as the main table.