Discussion Paper Series

IZA DP No. 18856

August 2026

Human Capital and the Uneven Geography of Poverty: Evidence from Peru

Gustavo Canavire-Bacarreza
World Bank and IZA@LISER

Luciana de la Flor Giuffra
World Bank

Carlos Rodriguez-Castelan
World Bank and IZA@LISER

Hugo Vaca Pereira Rocha
University of Wisconsin - La Crosse

IZA@LISER NETWORK
ISSN: 2365-9793

IZA DP No. 18856 August 2026

Human Capital and the Uneven Geography of Poverty: Evidence from Peru*

Abstract

Regional convergence theory predicts that poorer places may catch up with richer ones as productivity and human capital accumulate, yet convergence remains uneven in many developing countries. Peru experienced substantial reductions in poverty between 2007 and 2017, but many districts that started among the poorest remained there a decade later. This paper examines which initial district characteristics are associated with upward movement in the poverty distribution. We find that the composition of local human capital is more informative than municipal fiscal indicators. Higher secondary educational attainment is strongly associated with upward mobility, while tertiary attainment exhibits a nonlinear relationship with persistent poverty among initially poor districts. We also find that secondary education in neighboring districts is associated with favorable re-ranking, although evidence on pre-period migration indicates that this relationship reflects both local complementarities and skill-selective sorting rather than pure spillovers. These findings highlight the importance of territorial human capital in shaping local convergence despite broad national poverty reduction.

JEL classification

I32, O15, O40, R12

Keywords

poverty convergence, movements in the relative ranking, human capital, migration sorting, spatial human capital, Peru

Corresponding author

Gustavo Canavire-Bacarreza
gcanavire@worldbank.org

* We thank Bledi Celiku, Daniel Barco, Hugo Ñopo and Alejandro Torrez for their comments to earlier versions of the paper. The opinions expressed in this paper are those of the authors and do not necessarily reflect the views of the World Bank, its Board of Directors, or the countries they represent.

1 Introduction

Regional convergence theory posits that poorer territories may gradually catch up with richer ones as diminishing returns to capital, technological diffusion, and improvements in human capital reduce productivity gaps over time. In the neoclassical framework, regions with lower initial levels of income or development are expected to grow faster because capital is relatively more productive where it is scarce (Solow, 1956; Barro and Sala-i Martin, 1992). However, subsequent theoretical and empirical research has shown that convergence is neither automatic nor uniform across space. Differences in institutional quality, infrastructure, human capital accumulation, market access, and local state capacity can generate persistent divergence or the clustering of territories into distinct re-ranking groups, in which areas with similar structural characteristics move toward different long-run equilibria (Quah, 1996; Azariadis and Drazen, 1990). As a result, aggregate national growth may coexist with enduring territorial disparities, particularly in developing countries where structural inequalities across regions remain pronounced.

Peru provides a clear illustration of this dynamic, between 2007 and 2017, the average district-level poverty rate in Peru fell from 58 to 28 percent, and 95 percent of its 1,734 districts reduced poverty in absolute terms. Yet substantial territorial disparities persisted: more than two-thirds of districts that began the period in the poorest quintile remained in the bottom two quintiles a decade later, and a considerable subset stayed in the very poorest quintile throughout. Peru therefore experienced broad poverty reduction without widespread upward movement in the relative poverty distribution across territories. This divergence between aggregate progress and persistent local hierarchies reflects a broader question in development economics: under what conditions does economic growth translate into territorial convergence rather than simply parallel improvements across regions? Existing evidence shows that the elasticity of poverty reduction to growth varies substantially across countries and over time (Ravallion, 1997, 2001; Bourguignon, 2003; Barro, 2000; Kraay, 2006; Ravallion, 2012), suggesting that the capacity of poorer territories to catch up depends not only on aggregate growth itself, but also on local structural characteristics. Peru's experience resembles the persistence of poverty clubs documented across countries by Marrero et al. (2025) and the heterogeneous municipal convergence patterns identified for Mexico by López-Calva et al. (2022).

This paper examines district-level poverty dynamics in Peru between 2007 and 2017. We classify districts into five relative-poverty re-ranking groups according to how their position in the national poverty distribution shifts over the decade: Low-poverty, Improving, Stable, Worsening, and Lagging. The grouping is based on quintile transitions and draws on the distribution-dynamics tradition of Quah (1993, 1996); the labels summarize relative-rank persistence and re-ranking. We then examine how baseline human capital, labor-market structure, municipal fiscal capacity, central-government spending, and neighboring-district characteristics relate to membership in these groups. The analysis is descriptive and predictive rather than causal: its aim is to identify the structural characteristics associated with upward movement in the relative poverty distribution, as distinct from absolute poverty reduction alone.

Our analysis yields three main findings. First, local human-capital composition emerges as a central correlate of movements in the relative ranking, with secondary education the most robust margin. A higher baseline secondary-education share is broadly associated with upward movement in the poverty distribution and steepens the conditional convergence slope, a pattern that holds in both the convergence regressions and an ordered re-ranking specification. Tertiary education and the absence of formal schooling matter more selectively: among the initially poorest (fifth-quintile) districts, the probability of remaining in the Lagging group falls sharply at low levels of tertiary attainment before flattening, while the no-education share is positively associated with persistent lagging over the same range. These gradients are descriptive and do not identify a mechanism: the same pattern is consistent with local complementarities and threshold externalities (Azariadis and Drazen, 1990; Gennaioli et al., 2013), but also with the selective migration of educated residents toward already-advantaged districts, which we document below. We therefore read the education gradients as one of several interpretations the cross-sectional evidence cannot adjudicate.

Second, spatial context also matters, although selectively. Among neighboring characteristics, secondary education is the only robust predictor of favorable re-ranking outcomes. By contrast, neighboring consumption levels, informality, and agricultural employment display limited predictive power for re-ranking-group membership. Moreover, the spatial association is concentrated primarily along upward re-ranking margins: neighboring secondary education is positively associated with Low-poverty and Improving status, but displays little relationship with Lagging or Worsening outcomes. These results point toward a territorial human-capital environment shaped by spatial complementarities, local feedback effects, and skill-selective sorting, consistent with the framework developed by Desmet et al. (2026). Because education-selective migration can generate similar spatial gradients, we interpret the neighboring-education coefficient as a spatial human-capital association rather than as direct evidence of pure spillovers. More broadly, the findings are consistent with geographic poverty-trap mechanisms (Jalan and Ravallion, 1998, 2002), while cautioning against a purely spillover-based interpretation of neighboring education (Rey and Montouri, 1999; Conley and Ligon, 2002).

Third, aggregate municipal fiscal variables provide limited explanatory power once education and labor-market composition are taken into account. Municipal spending, investment, and budget execution are not consistently associated with upward movement. We exclude central-government social-protection spending from the substantive interpretation. Major programs such as Juntos and Pensión 65 are geographically allocated using the same poverty maps that define our outcome classification (World Bank, 2023), making any in-sample association mechanical rather than behavioral; moreover, the administrative records do not isolate program-specific transfers at the district level. Overall, the results suggest that local and neighboring human capital are substantially more informative for explaining movements in the relative ranking than the aggregate municipal fiscal-capacity measures available in the data.

The paper contributes to the within-country poverty-convergence literature (López-Calva et al., 2022; Miranti, 2021; Najam and Gibson, 2022; Crespo Cuaresma et al., 2022) in four ways. First, it adapts the convergence-clubs perspective of Marrero et al. (2025) to a within-country setting, using district-level poverty-quintile transitions to describe relative re-ranking while holding national institutions and macroeconomic conditions approximately fixed. Second, it places the convergence question at the local administrative margin, where decentralization, public-service delivery, and human-capital accumulation are directly observed, consistent with evidence that the geography of poverty is predominantly subnational (Kim et al., 2016). Third, it contrasts human-capital composition with municipal fiscal capacity and shows that education profiles are more consistently associated with movements in the relative ranking than aggregate local fiscal indicators, complementing the evidence for Mexico in López-Calva et al. (2022). Finally, it shows that spatial associations operate primarily through neighboring human capital and that substantial heterogeneity remains even within the Lagging group. Districts that stayed near the bottom of the relative poverty distribution but entered the decade with stronger educational endowments experienced larger absolute poverty reductions, underscoring that the re-ranking groups are descriptive categories rather than internally homogeneous re-ranking regimes.

The remainder of the paper proceeds as follows. Section 2 describes Peru's decentralization context and situates the analysis within the poverty-convergence literature. Section 3 presents the data, outcome measures, and the relative-poverty re-ranking groups. Section 4 documents absolute poverty reduction alongside persistent relative rankings and presents the convergence benchmark that anchors the paper. Section 5 treats the re-ranking groups as a descriptive complement, examining how baseline human capital and fiscal capacity relate to group membership. Section 6 extends this to spatial context, using neighboring-district characteristics and pre-period migration to interpret why conditional catch-up remains incomplete. Section 7 presents robustness checks, including ordered-logit re-ranking, within-group dispersion, and within-Lagging heterogeneity. Section 8 concludes.

2 Context and Setting

2.1 Peru's poverty decline

Peru is an ideal setting to examine incomplete convergence because a decade of broad-based growth coexists with a decentralized fiscal architecture and a marked geographic gradient in human capital, leaving the same national aggregate consistent with very different local trajectories. Peru grew at an average real GDP rate of 6.6 percent between 2004 and 2013—nearly double the regional average of 3.6 percent. Real GDP per capita roughly doubled between 2004 and 2019 (World Bank, World Development Indicators; constant 2015 US$), and growth disproportionately benefited the bottom 40 percent of the income distribution. Poverty rates fell from 58.7 percent in 2004 to 20.2 percent in 2019. Roughly 85 percent of the reduction reflects economic growth; the remainder reflects redistribution (World Bank, 2023).

Peru launched a decentralization reform in 2001, transferring competencies and resources to regional and municipal governments. Implementation has lagged. Municipalities hold competencies across sectors but face no obligation to coordinate, plan, or set priorities, and development plans lack effective enforcement instruments. Territorial zoning—the formal mechanism through which municipalities direct land use and local investment—remains underdeveloped (Cuenca et al., 2022).

Human capital shifted alongside these macroeconomic and institutional changes. The working-age share of the population rose 8 percentage points between 2007 and 2017, and educational attainment expanded across all levels. Schooling quality, by contrast, lagged: by 2019, only 24 percent of fourth-grade primary students and 9 percent of second-grade secondary students reached satisfactory scores on national standardized tests (World Bank, 2023). Our analysis therefore measures the local human-capital stock through attainment shares of the working-age population, which are observed at the district level for all 1,734 districts at the 2007 baseline.

2.2 Why convergence may remain incomplete

The four candidate channels flagged in Section 1—initial poverty, human-capital composition, fiscal capacity, and spatial context—originate in two connected literatures. The empirical convergence literature documents which aggregates converge, which do not, and along which dimensions groups emerge. Ravallion (2012) established that poverty does not converge even when mean incomes do, tracing the result to a vicious circle in which high initial poverty slows growth and dampens its poverty-reducing elasticity. Crespo Cuaresma et al. (2022) show that absolute poverty convergence is robust under a refined specification of the growth–poverty link. Marrero et al. (2025) test for club convergence directly, applying the Phillips and Sul (2007) panel-clustering framework[1] to four absolute-poverty measures across 100 emerging and developing countries from 1981 to 2019; they reject both absolute and conditional convergence in favour of club convergence, identifying four clubs in the headcount-poverty series with the highest-poverty club—which contains roughly half of the world's poor—exhibiting near-stationary poverty rates and the hallmarks of a poverty trap.

2.3 Candidate predictors at the district margin

At the sub-national level, the canonical cross-country evidence comes from Gennaioli et al. (2013), who document across 1,569 regions in 110 countries that human capital is the dominant correlate of regional GDP per capita and productivity, dwarfing geographic, institutional, and cultural channels. Country-specific sub-national studies are broadly consistent

[1] The Phillips and Sul (2007) log-t convergence test identifies clusters from each unit's relative transition path and requires sufficiently long time-series variation to estimate that path and to test the null of common transitional behavior. The Peruvian small-area-estimation poverty maps cover two end-period observations (2007 and 2017), which does not provide the longitudinal variation the test requires. We therefore construct re-ranking groups directly from observed quintile transitions rather than from the panel-clustering procedure of Phillips and Sul (2007).

with this prior. López-Calva et al. (2022) document municipal poverty convergence in Mexico and credit federal transfers, but do not separate central-government targeting from municipal spending. Miranti (2021) documents poverty convergence alongside persistent spatial polarization in Indonesian districts without decomposing the structural drivers or testing for threshold nonlinearities. Najam and Gibson (2022) find that spatial spillovers accelerate convergence in Pakistan, though the specific channel remains unidentified. For Peru, Canavire-Bacarreza et al. (2018) document substantial heterogeneity in the sectoral growth–poverty relationship across the 25 *departamentos*, finding that manufacturing growth reduces poverty and service-sector growth strengthens the middle class; the district margin—where Peru's fiscal decentralization actually operates—has not been examined. Our paper ap- plies the distribution-dynamics perspective of Marrero et al. (2025) at the within-country margin: by holding country-level growth and macro institutions roughly fixed, we ask which baseline conditions distinguish catch-up districts from districts that improved in absolute terms.

Theoretical work motivates three priors about which baseline conditions should predict catch-up when convergence is incomplete. First, threshold externalities in human capital (Azariadis and Drazen, 1990) imply that below a critical stock of education, initially poor regions remain trapped; education-cost heterogeneity sharpens this prediction (Desmet et al., 2026). Second, O-Ring-style complementarities (Kremer, 1993) imply that skills interact within the local economy, so that the composition of education—not just its level—predicts outcomes. Third, standard spatial-dependence models (Rey and Montouri, 1999; Conley and Ligon, 2002) predict spillovers operating through neighbors' consumption or employment. Yet a spatial human-capital gradient need not reflect spillovers at all: if educated workers sort into already-advantaged districts, selective migration can reproduce the same neighbor–outcome association without any local complementarity. Whether the spatial channel runs through human capital—and, if so, whether it reflects complementarities or skill-selective sorting—is therefore an empirical question, which we revisit using predetermined internal migration.

These priors map directly into the empirical work. Section 4 tests whether catch-up is conditional and whether the convergence slope varies with baseline education; Section 5 examines whether nonlinearities concentrate among districts that start in the poorest quintile; and Section 6 asks whether the spatial human-capital association implies some of the mechanisms discussed.

3 Data, outcomes, and re-ranking-group construction

3.1 Data sources and baseline covariates

Studying which districts catch up requires a unit at which catch-up is observable. Peru's 1,734 districts are small enough to make within-country heterogeneity visible, yet covered consistently enough by the census to provide comparable covariates at both endpoints of the decade. We combine three complementary sources. Household surveys measure consumption and poverty rigorously but lack district-level resolution; the population censuses fill this gap with granular human-capital and labor-market characteristics for every district (*distrito*); and administrative records from the Ministry of Economy and Finance (MEF) supply the fiscal dimension.

All covariates are measured at the 2007 baseline; the outcomes-—relative-poverty re-ranking-group membership and the ordered quintile-transition index—are constructed from poverty maps for 2007 and 2017. Temporal precedence rules out reverse causality from outcomes to covariates, though we do not claim strict causal identification. Poverty and consumption maps for 2021 are available but not analyzed here.

The main analysis uses two outcome variables—per-capita consumption and the head-count poverty rate—measured at the 2007 baseline and the 2017 end point. Consumption is estimated from INEI's National Census using INEI's regional per-capita consumption model. Poverty rates in every year are measured against the national poverty line—the minimum per-capita cost of an adequate basket of goods and services—which varies by geographic region and by rural or urban classification.

3.2 Poverty outcomes and quintile transitions

To motivate the re-ranking-group classification used throughout the paper, we first describe district-level poverty dynamics in the spirit of Quah (1993, 1996). Let $P_{ij} = \Pr(Q_{2017} = j \mid Q_{2007} = i)$ be the empirical transition probability between poverty quintiles, where Q_1 is the least-poor quintile and Q_5 the most-poor. Table 1 reports the matrix. Two features are central to the analysis. First, the chain is strongly persistent at both tails: $P_{11} = 0.720$ and $P_{54} + P_{55} = 0.717$. The latter quantity is our Lagging-quintile transition rate expressed as a one-step Markov persistence probability. Second, the middle of the distribution re-ranks much more than the tails, with diagonal persistence in Q_2–Q_4 between 0.29 and 0.45.The asymmetric tail-persistence pattern echoes the bimodal distribution dynamics emphasised by Quah (1996), here documented as relative-rank persistence within Peru rather than as a long-run-stationary Markov-chain result. Marrero et al. (2025) document an analogous pattern at the cross-country level.

Table 1: District poverty-quintile transitions, 2007–2017

	Quintile in 2017					
Quintile in 2007	**Q1**	**Q2**	**Q3**	**Q4**	**Q5**	N_i
Q1, least poor	**0.720**	0.199	0.061	0.012	0.009	347
Q2	0.170	**0.447**	0.205	0.084	0.095	347
Q3	0.063	0.182	**0.288**	0.199	0.268	347
Q4	0.040	0.124	0.219	**0.305**	0.311	347
Q5, most poor	0.006	0.049	0.228	**0.402**	**0.315**	346
Diagonal persistence, P_{ii}	0.720	0.447	0.288	0.305	0.315	

Notes: Cells report $P_{ij} = \Pr(Q_{2017} = j \mid Q_{2007} = i)$, the empirical probability of moving from poverty quintile i in 2007 to quintile j in 2017. Q1 is the least-poor quintile; Q5 is the most-poor. Shaded diagonal cells indicate quintile persistence. Darker shaded cells identify the Lagging-group transition rate, $P_{54} + P_{55} = 0.717$. Because quintiles are equal-sized by construction, the long-run marginal distribution is not reported. Sources: 2007 and 2017 poverty maps, INEI and ENAHO.

3.3 Relative-poverty re-ranking groups

To identify which baseline characteristics are associated with stronger or weaker upward movement, we classify districts into five relative-poverty re-ranking groups based on quintile transitions, following the distribution-dynamics approach of Quah (1993, 1996). Each district is assigned a poverty quintile in 2007 and 2017:

$$Q_{2007,i} = \text{quintile}(\text{poverty}_{2007,i}) \tag{1}$$

$$Q_{2017,i} = \text{quintile}(\text{poverty}_{2017,i}) \tag{2}$$

Five groups are defined: (1) *Low-poverty*: $Q_{2007} = 1$ and $Q_{2017} = 1$; (2) *Lagging*: $Q_{2007} = 5$ and $Q_{2017} \in \{4, 5\}$; (3) *Improving*: $Q_{2017} < Q_{2007}$ and not already classified as Low-poverty or Lagging (so a Q5→Q4 transition is Lagging, not Improving); (4) *Worsening* (in relative rank): $Q_{2017} > Q_{2007}$; (5) *Stable*: all remaining. The Worsening label refers exclusively to the change in quintile position, not to the absolute poverty change: most Worsening districts still reduced poverty in absolute terms, although 82 districts of the 1,734 (distributed across groups) did experience small absolute increases. This classification is robust to measurement error at the tails because it uses ordinal ranks rather than cardinal poverty levels. For the ordered-logit robustness check, we also compute a nine-category ordered quintile-transition measure, $\text{mobility}_i = Q_{2007,i} - Q_{2017,i} \in \{-4, -3, \ldots, +4\}$, where positive values indicate upward re-ranking. We emphasize that these are descriptive re-ranking groups summarizing relative-rank movement, not convergence clubs in the formal sense. Throughout, "mobility" and "re-ranking" refer to changes in a district's relative position within the national poverty distribution, as captured by the quintile transition matrix described above.

The quintile transition matrix shows that 14 percent of districts (248) started in Q5 in 2007 and were still in Q4 or Q5 in 2017—our Lagging group—while 22 percent (375) moved to a strictly lower quintile—our Improving group. Q5→Q4 transitions are counted in Lagging, not Improving. We probe the empirical consequences of this rule in Appendix

Table 18, where the multinomial logit is re-estimated with Q5→Q4 reclassified as Improving; the headline re-ranking coefficients are robust on the Improving margin, while on the Lagging margin only initial poverty retains its sign and significance and the education coefficients are weak and unstable—consistent with our demotion of the group apparatus. Geographically, the Sierra is mixed (26.7 percent lagging, 27.7 percent improving), the Selva is dominated by relative worsening (62.6 percent), and the Costa concentrates in the Low-poverty group (30.7 percent), with only 4.3 percent of its districts lagging.

Figure 1 maps these five groups across Peru's 1,734 districts and shows how Lagging districts concentrate in the southern Sierra and parts of the Selva, while Low-poverty districts cluster along the Costa.

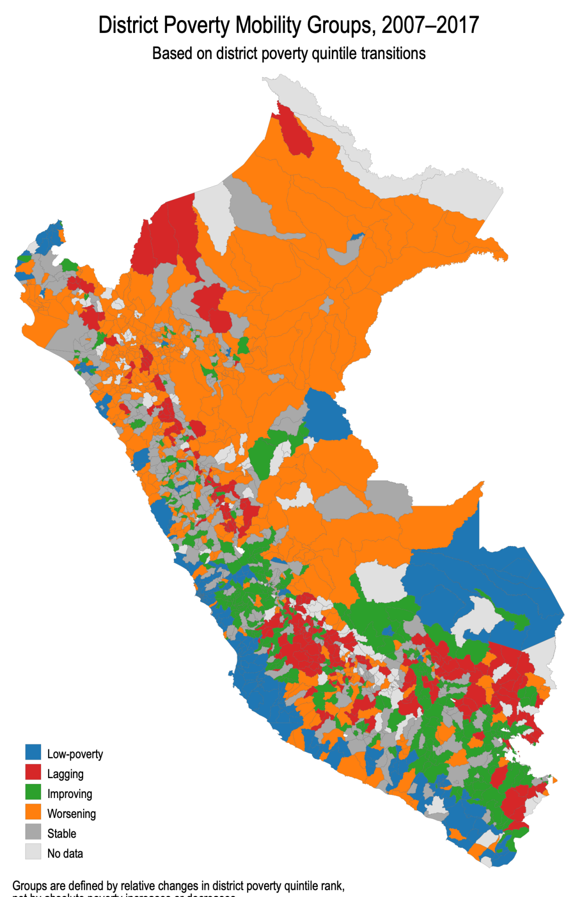

Figure 1: District Poverty re-ranking groups, 2007–2017. *Notes:* Groups are based on district poverty quintile transitions. Categories indicate relative changes in district poverty quintile rank, not absolute poverty increases or decreases.

3.4 Baseline differences across re-ranking groups

Human capital and labor. The census provides the core socioeconomic covariates: economically active population, employment levels, agricultural-employment share, and informality, approximated by the share of workers in firms with fewer than five employees. Education decomposes into four attainment categories for the working-age population (PET, *población en edad de trabajar*)—no formal education, primary, secondary, and tertiary—with tertiary computed as the residual after all other categories are subtracted. We verify this construction by confirming the absence of negative residuals (a non-trivial check on the residual category) and through internal consistency checks against known district profiles sampled across all macro-regions.

Fiscal variables, physical capital, and spatial structure. MEF administrative data cover municipal and central-government expenditure and income from 2004 to 2021, disaggregated into current spending and investment across sectors including agriculture, education, social protection, and health. Government income encompasses canon transfers, taxes, and donations. All fiscal variables enter the models as three-year pre-period per-capita averages expressed in thousands of 2017 soles. Municipal expenditure is aggregated across sectors because sector-specific municipal spending exhibits near-zero cross-district variance, while central-government spending is kept sector-specific given its meaningful geographic variation. Budget execution—the share of each district's approved budget actually spent—serves as an institutional proxy for local administrative capacity. Physical capital enters through firm-level assets and revenues from INEI's National Economic Survey, allocated to districts proportionally using the 2008 National Economic Census. Spatial lags of key variables are computed from a queen-contiguity weight matrix, which identifies neighboring districts that share a border or vertex.[2]

[2]Geographic fixed effects enter the regressions at two levels: 24 department dummies in the linear β-convergence regression of Section 4.1, and 3 macro-region dummies (Costa, Sierra, Selva) in the multinomial logits of Sections 5–6, where the per-class parameter cost is four times larger and finer geography would induce quasi-separation in Selva.

Central-government expenditure on social protection doubled between 2007 and 2017, and education expenditure nearly doubled, from 507 to 936 soles per capita. Municipal income from the canon (natural-resource revenue sharing) rose 28 percent over the same period, while municipal tax collection remained a modest source of local revenue relative to central transfers. The analysis below tests whether central versus municipal spending predicts the pace of district-level poverty reduction.

Table 2 reports baseline mean characteristics across the five re-ranking groups. Lagging and Improving districts had different baseline poverty rates (87 versus 69 percent) and sharply different education compositions: 6.5 percent of the working-age population in Lagging districts had completed tertiary education versus 13.4 percent in Improving; 22.4 percent had no formal education in Lagging versus 14.4 percent in Improving. The two groups also differ in working-age share and informality, but the education gap is the largest in relative terms. These contrasts motivate the multivariate analysis that follows.

Table 2: Mean characteristics by re-ranking group, 2007 baseline

	Low-poverty	**Lagging**	**Improving**	**Worsening**	**Stable**
Initial conditions					
Poverty rate (%)	21.1	**86.5**	68.5	53.5	57.2
Log consumption	6.10	**4.99**	5.37	5.38	5.44
Education composition, % of working-age population					
No education	4.6	**22.4**	14.4	12.6	13.9
Secondary	41.4	28.9	35.3	31.9	33.8
Tertiary	**34.0**	**6.5**	13.4	13.1	15.7
Labor market					
Working-age population (%)	73.6	60.4	67.7	65.3	66.0
Informal employment (%)	62.5	**86.4**	81.4	80.7	78.6
Agricultural employment (%)	25.0	**73.5**	62.2	59.7	55.6
Central-government expenditure, thous. soles per capita					
Social protection	0.639	0.231	0.404	0.267	0.375
Education	0.446	0.493	0.462	0.469	0.450
Health	0.192	0.122	0.144	0.120	0.130
Municipal expenditure, thous. soles per capita					
Total current expenditure	0.125	0.118	0.154	0.094	0.129
Total investment	0.165	0.116	0.210	0.152	0.179
Fiscal capacity and economic structure					
Budget execution (%)	52.4	43.8	40.1	48.2	47.8
GDP per capita	21,288	6,715	12,169	10,646	11,551
Districts	250	248	375	500	361

Notes: Bold values highlight substantively important contrasts discussed in the text. Low-poverty districts are in Q1 in both years. Lagging districts are in Q5 in 2007 and Q4–Q5 in 2017 (so Q5→Q4 is classified as Lagging, not Improving). Improving districts move to a strictly lower poverty quintile and did not start in Q5; Worsening districts move to a higher poverty quintile; Stable districts remain in the same quintile.
Fiscal variables are three-year pre-period averages per capita in thousands of 2017 soles. Sources: Population Census 2007, ENAHO 2007/2017, and MEF.

4 Absolute progress with relative persistence

4.1 Distributional patterns across districts

Aggregate progress masks sharply divergent local trajectories. At the district level, the average poverty rate fell from 58 percent in 2007 to 28 percent in 2017—a 30-percentage-point decline. Improving-group districts (those that moved to a strictly lower poverty quintile, except Q5→Q4 transitions, which are classified as Lagging) cut poverty by 51 percentage points on average, from 69 to 18 percent. Lagging-group districts (those that began in the poorest quintile in 2007 and remained in the bottom two quintiles in 2017) cut poverty by 42 percentage points, from 87 to 45 percent—a large absolute gain, but insufficient to close the gap with faster-improving peers. Districts that worsened in relative rank cut poverty by only about 13 percentage points in absolute terms, from 53.5 to 40.4 percent. Per-capita consumption rose 57 percent on average, but the gains concentrated in coastal areas and parts of the central Amazon and Andes. The southern Sierra—where 26.7 percent of districts classify as lagging in our analysis—saw the least relative progress.

Figure 2 makes this divergence concrete. At the group-mean level, all five groups reduced poverty in absolute terms between 2007 and 2017, but the pace varied sharply. At the individual-district level, 82 of 1,734 districts (4.7 percent) experienced small absolute increases; these are distributed across groups and do not change the group-mean pattern.

Panel (a) of Figure 2 plots the distribution of district-level poverty rates in 2007 and 2017. The 2007 distribution is bimodal—the "twin peaks" phenomenon of Quah (1996)—one mode near 35 percent poverty and a second mode near 80 percent. By 2017, the entire distribution has shifted leftward and compressed toward unimodality, confirming absolute improvement across the board. A persistent right tail remains: the districts that started poorest improved substantially but could not close the gap. Panel (c) decomposes the 2007 distribution by macro-region and reveals that the national bimodality is a between-region phenomenon. The Costa concentrates at low poverty rates, the Sierra spans the full range, and the Selva concentrates at moderate-to-high rates. The national twin peaks arise because these three regions sit at different points in the poverty distribution—not because any single region is internally polarized.

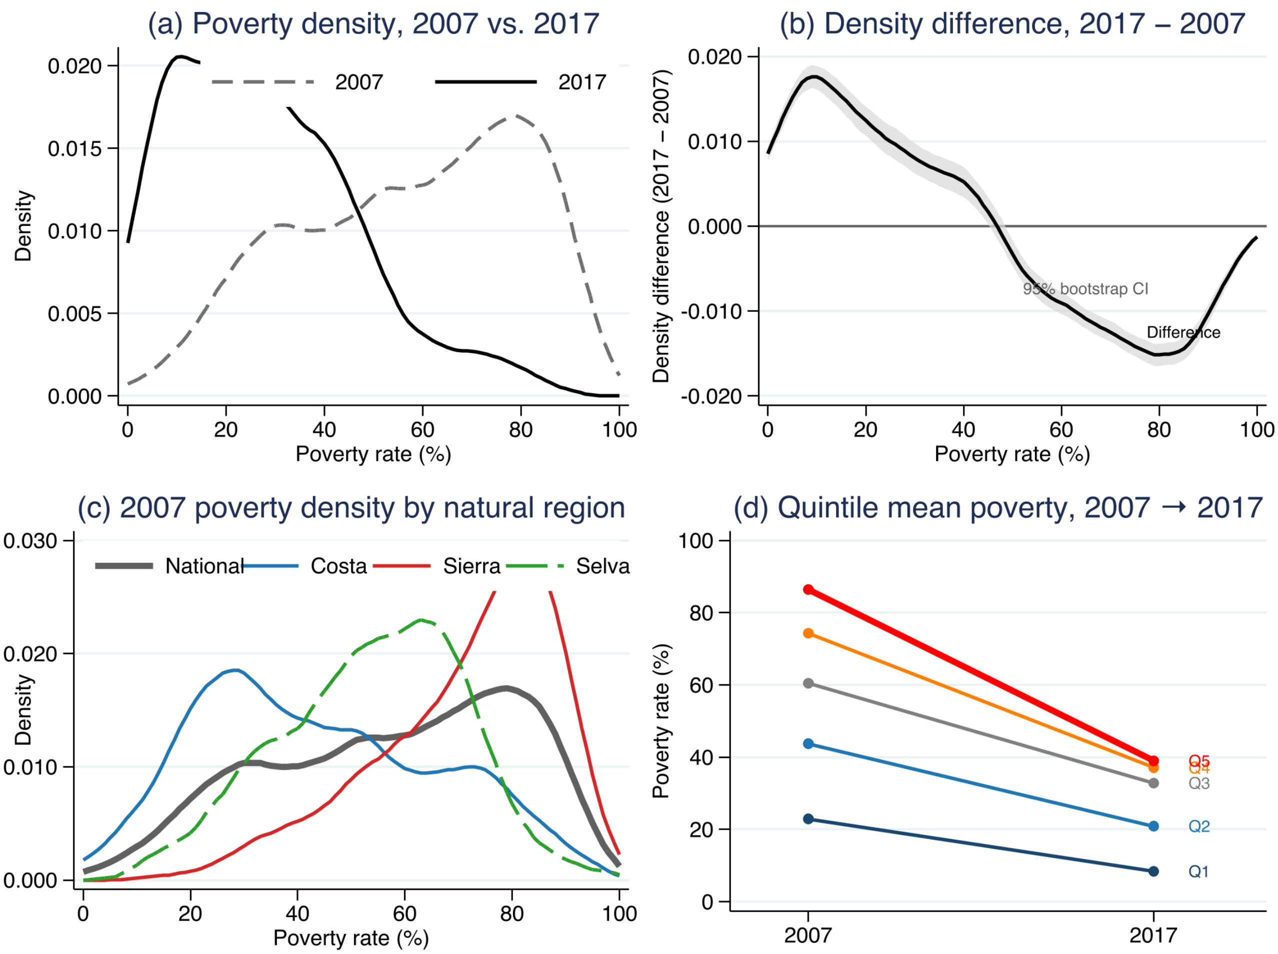

Figure 2: Distributional poverty reduction, regional structure, and re-ranking-group trajectories, 2007–2017. Panel (a) plots kernel densities of district poverty rates in 2007 and 2017. Panel (b) plots the density difference, 2017 minus 2007, with a 95 percent bootstrap confidence band. Positive values indicate poverty-rate ranges where the 2017 distribution has more mass than the 2007 distribution; negative values indicate ranges where it has less mass. Panel (c) decomposes the 2007 district poverty distribution by natural region. Panel (d) plots mean poverty rates in 2007 and 2017 by re-ranking group. Natural region explains 31.2 percent of cross-district variation in 2007 poverty rates based on a one-way ANOVA, $p < 0.001$.

Distinguishing districts that caught up from those that improved only in absolute terms requires moving from descriptive geography to multivariate analysis. Existing work suggests four candidate channels: initial poverty, human-capital composition, fiscal capacity, and spatial context. López-Calva et al. (2022) examine fiscal capacity in Mexican municipalities; Miranti (2021) examines spatial polarization in Indonesian districts; Najam and Gibson (2022) examine spatial spillovers in Pakistan; and Kim et al. (2016) document the multilevel structure of within-country poverty geography in India. None evaluates the four channels jointly at the local administrative margin where fiscal decentralization actually operates, despite the prevalence of multi-tier subnational governance in middle-income economies (Bardhan, 2002; Faguet, 2014). Peru, with 1,734 districts and a two-decade decentralization process, offers a useful setting for examining how these channels distinguish catching-up from lagging localities.

4.2 β-convergence at the district level

The classification used throughout the paper is a relative-rank concept based on quintile transitions: it documents whether districts moved up or down in the national poverty distribution but does not, by itself, speak to the speed at which initially poorer districts close the gap. The cross-country growth and poverty-convergence literatures, since Barro and Sala-i Martin (1992), summarize that speed in a single parametric coefficient—the rate at which initially worse-off units move toward a common path, conditional on observables. Adapting that benchmark to district-level poverty, in the spirit of Ravallion (2012) and Ouyang et al. (2019), we ask whether Peru's aggregate poverty reduction looked like proportional catch-up across districts or like uniform improvement that left relative ranks intact—and, conditional on baseline structure, whether convergence is detectable at all. We therefore estimate

$$g_i = \alpha + \beta \ln H_{i,2007} + X_i'\gamma + \varepsilon_i, \tag{3}$$

where

$$g_i = \frac{\ln H_{i,2017} - \ln H_{i,2007}}{10}$$

is the annualized log rate of poverty change. Since negative values of g_i indicate poverty reduction, a negative estimate of β indicates proportional β-convergence: districts with higher initial poverty experienced faster proportional poverty reduction. Because the dependent variable is a bounded poverty rate, β-convergence here should be read as proportional catch-up in poverty reduction conditional on baseline structure, not as literal convergence to a common steady-state poverty rate. The coefficient answers whether initially poorer districts closed the gap faster once observable district fundamentals are held fixed.

Table 3 reports the estimates by sample restriction. In the full district sample, the unconditional coefficient is positive and precisely estimated, implying proportional anti-convergence: initially poorer districts experienced less negative log-rate poverty growth than less-poor districts. Conditioning on baseline log consumption, education composition, agricultural employment, and department fixed effects reverses the sign and yields a negative, precisely estimated coefficient consistent with conditional β-convergence at a half-life of roughly 16 years. The Yohai MM estimator[3] preserves the unconditional OLS sign, suggesting that the full-sample divergence result is not driven by outliers.

The conditional convergence pattern is concentrated at moderate initial-poverty levels. Among districts with poverty above 60 percent in 2007, both the unconditional and conditional coefficients are negative and statistically significant. Above the 70 and 80 percent cutoffs, the log-rate coefficients are statistically indistinguishable from zero. The standard β-convergence benchmark therefore provides little evidence of robust proportional catch-up among the very poorest districts, the population emphasized by the lagging classification.

[3]The MM estimator combines a high-breakdown-point S-estimator on the same scale with an efficient redescending M-step (Yohai, 1987); it is implemented in Stata via `robreg mm` at 85 percent efficiency.

Table 3: Log-rate poverty β-convergence across Peruvian districts, 2007–2017

Panel A. Full sample and 60 percent initial-poverty cutoff

	No cutoff, full sample			$H_{2007} > 60\%$		
	Uncond. OLS	Cond. OLS	MM	Uncond. OLS	Cond. OLS	MM
Initial log poverty, $\ln H_{2007}$	0.0357***	-0.0446***	0.0338***	-0.0427***	-0.0471***	-0.0793***
	(0.0035)	(0.0078)	(0.0033)	(0.0142)	(0.0140)	(0.0145)
Controls	No	Yes	No	No	Yes	No
Observations	1,728	1,728	1,728	876	876	876
R^2	0.100	0.557	–	0.010	0.631	–

Panel B. Stricter initial-poverty cutoffs

	$H_{2007} > 70\%$			$H_{2007} > 80\%$		
	Uncond. OLS	Cond. OLS	MM	Uncond. OLS	Cond. OLS	MM
Initial log poverty, $\ln H_{2007}$	0.0046	-0.0271	-0.0383	0.0244	-0.0617	0.0161
	(0.0233)	(0.0205)	(0.0235)	(0.0442)	(0.0400)	(0.0418)
Controls	No	Yes	No	No	Yes	No
Observations	648	648	648	357	357	357
R^2	0.000	0.628	–	0.001	0.574	–

Notes: The dependent variable is annualized log poverty growth, $g_i = (\ln H_{i,2017} - \ln H_{i,2007})/10$. Negative coefficients indicate proportional β-convergence. Robust standard errors are in parentheses. Conditional specifications include baseline log consumption, education composition, agricultural employment, and department fixed effects. MM denotes Yohai MM robust regression. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3 implies incomplete convergence: districts above roughly 60 percent initial poverty appear to converge toward the middle once baseline structure is held constant, but the very poorest districts show no robust proportional catch-up. The core result is therefore conditional convergence for moderately poor districts and non-convergence at the bottom tail. Which baseline characteristics account for that conditional convergence is the question Section 5 examines directly.

The absence of catch-up at the bottom is mirrored in the cross-sectional dispersion of the poorest districts. Among the Lagging group, the standard deviation of district poverty rates *rises* from 0.041 in 2007 to 0.124 in 2017 (Table 15), a clear σ-divergence that stands in contrast to the σ-convergence observed for every other group. Far from closing, the gap separating the persistently poorest districts from the rest widens over the decade—reinforcing the bottom-tail non-convergence found in the β-convergence regressions.

The baseline β-convergence specification imposes a common poverty-reduction slope across districts, but Table 4 suggests that this is too restrictive: most conditioning covariates converge across districts, whereas tertiary attainment diverges. Table 5 therefore allows the convergence slope to vary with baseline education by interacting initial log poverty with the no-education, secondary, and tertiary shares. In the full specification, the three interactions are jointly significant, indicating that convergence rates depend on education composition rather than on initial poverty alone. The secondary margin strengthens catch-up: the interaction with initial poverty is negative, and in column (3) the implied slope steepens from -0.026 at one standard deviation below the mean secondary share to -0.058 at one standard deviation above. Tertiary attainment moves in the opposite direction. In the full interaction model, the $\ln H_{2007} \times$ Tertiary coefficient is positive (+0.066), implying a flatter convergence slope in districts with larger tertiary shares; the no-education interaction, by contrast, is not robust once all three education margins are included together. The pattern is therefore asymmetric: secondary attainment is the education margin most closely associated with faster proportional poverty reduction among initially poor districts, whereas tertiary attainment is more spatially concentrated and less closely associated with broad-based catch-up. Section 5 then turns to the re-ranking-group results as a descriptive complement, asking how this same asymmetry appears in relative-rank outcomes.

Table 4: Convergence and divergence in the conditioning covariates, 2007–2017

	σ-convergence				β-convergence	
	SD_{07}	SD_{17}	CV_{07}	CV_{17}	Full	H07>60%
Education (share of working-age population)						
No education	0.0883	0.0770***	0.6531	0.6889	−0.0213***	−0.0265***
					(0.0017)	(0.0025)
Primary	0.1330	0.1294	0.3628	0.3748	−0.0102***	−0.0108***
					(0.0008)	(0.0021)
Secondary	0.0913	0.0836***	0.2690	0.1997	−0.0235***	−0.0267***
					(0.0013)	(0.0018)
Tertiary	0.1248	0.1422***	0.7916	1.1694	+0.0036***	+0.0055
					(0.0013)	(0.0042)
Economic structure						
Working-age share	0.0646	0.0615**	0.0972	0.0853	−0.0224***	−0.0109***
					(0.0015)	(0.0025)
Agricultural employment	0.1539	0.1570	0.5373	0.5184	−0.0205***	−0.0468***
					(0.0021)	(0.0039)
Poverty headcount	0.2298	0.1847***	0.3983	0.6660	see Table 3	

Notes: The table tests whether the conditioning covariates themselves converge or diverge across districts between 2007 and 2017. The σ-convergence block reports the cross-district standard deviation (SD) and coefficient of variation (CV) in each year; stars on SD_{17} mark rejection of the equal-dispersion null in a two-sided variance-ratio F-test, so a starred fall in the SD indicates σ-convergence and a starred rise indicates σ-divergence. The β-convergence block regresses the annualized change in the covariate, $(x_{2017} - x_{2007})/10$, on its own 2007 level; a negative coefficient indicates β-convergence, with robust standard errors in parentheses. "Full" uses all districts ($N = 1{,}728$); "H07>60%" restricts to districts with a 2007 headcount above 60% ($N = 876$). All covariates are shares: the education variables are shares of the working-age population; the working-age share is working-age population over total population; agricultural employment is agricultural workers over the working-age population. The poverty-headcount β-estimate is the headline convergence result in Table 3.

Only baseline characteristics observed consistently at both dates admit this test. Log household consumption underlies the poverty headcount itself and its 2007 and 2017 series are not on a common scale in the estimation sample; district GDP per capita has no comparable two-date per-capita level and is apportioned from province-level accounts; the central- and municipal-spending and budget-execution variables are policy flows entered as pre-period averages; and informal employment has no district-level 2017 measure. The reported covariates are the predetermined human-capital and sectoral-structure characteristics measured consistently in both years.

$^{*}\, p < 0.10$, $^{**}\, p < 0.05$, $^{***}\, p < 0.01$.

Table 5: Education × initial-poverty interactions in β-convergence

	(1) Baseline	(2) ×NoEd	(3) ×Sec	(4) ×Tert	(5) ×All
$\ln H_{2007}$	-0.0446^{***}	-0.0555^{***}	-0.0421^{***}	-0.0640^{***}	-0.0584^{***}
	(0.0078)	(0.0080)	(0.0071)	(0.0071)	(0.0077)
No education		0.0371			-0.0109
		(0.0328)			(0.0376)
Secondary			0.0465^{**}		0.0010
			(0.0203)		(0.0240)
Tertiary				0.0447^{**}	0.0515^{**}
				(0.0220)	(0.0224)
$\ln H_{07}\times$NoEd		-0.1232^{**}			-0.0214
		(0.0586)			(0.0896)
$\ln H_{07}\times$Sec			-0.1735^{***}		-0.1167^{**}
			(0.0418)		(0.0456)
$\ln H_{07}\times$Tert				0.0934^{***}	0.0657^{**}
				(0.0206)	(0.0309)
N	1,728	1,728	1,728	1,728	1,728
R^2	0.557	0.560	0.570	0.574	0.579

Notes: Dependent variable is annualized log poverty growth g (negative $\ln H_{2007}$ coefficient $= \beta$-convergence). Education shares are mean-centered, so $\ln H_{2007}$ is the convergence slope at average education. All columns include log consumption, agricultural employment, and department fixed effects; robust SEs. In column (5) the three interactions are jointly significant, $F(3, 1694) = 9.55$, $p < 0.0001$. From column (3), the implied convergence slope steepens from -0.026 at one SD below mean secondary share to -0.058 at one SD above. $^{**}p < 0.05$, $^{***}p < 0.01$.

5 Human capital and re-ranking groups

5.1 Education composition and re-ranking-group membership

Section 4 shows that poverty catch-up is conditional and that the convergence slope varies with baseline education. We now use the relative-poverty re-ranking groups as a descriptive complement to the β-convergence results. The multinomial logit asks how the same baseline characteristics map into relative-rank outcomes; it is not itself a β-convergence test. Risk sets differ mechanically across origin quintiles—Q1 districts cannot enter the Lagging group and Q5 districts cannot enter the Low-poverty group—so the estimated coefficients should be read as predictors of joint position-and-re-ranking outcomes rather than as predictors of re-ranking conditional on a common risk set. Section 7 reports an ordered re-ranking specification that is closer to a pure re-ranking outcome. A second concern is that the Q5→Q4 = Lagging rule could mechanically drive the Lagging-margin findings. Appendix Table 18 re-estimates the same multinomial logit with Q5→Q4 reclassified as Improving and shows that the Improving-margin coefficients retain their sign and significance, whereas on the Lagging margin only initial poverty is robust; the education and working-age coefficients there are weak and unstable under either classification rule, consistent with our demotion of the group apparatus.

The Barro and Sala-i Martin (1992) framework offers a natural lens for the conditioning set. Convergence is conditional on the steady state, and the steady state depends on initial endowments—most prominently human capital, which shapes both absorptive capacity and the returns to public investment in lagging areas (Barro, 2001). Districts with similar starting poverty but different human-capital stocks should therefore follow systematically different trajectories. We test this by treating the re-ranking group membership rather than the growth rate as the outcome, since growth regressions collapse heterogeneous trajectories into a single slope and cannot separate poor districts that improved (Improving) from poor districts that did not (Lagging).

Let y_i denote the re-ranking group of district i and X_i a vector of baseline characteristics measured in 2007. We model membership with a multinomial logit,

$$P(y_i = j \mid X_i) \;=\; \frac{\exp(X_i'\beta_j)}{\sum_{k=1}^{5}\exp(X_i'\beta_k)}, \tag{4}$$

with $\beta_5 = 0$ so that *Stable* is the omitted base.

Education enters as three shares of the working-age population—no education, secondary, and tertiary—with primary as the omitted base. The education decomposition is substantively important: a single index of years of schooling would mask the nonlinear returns

Table 6: Education-specification comparison: multinomial logit

Model	Education shares included	AIC	ΔAIC	BIC	Pseudo R^2	Log-lik.
A	Secondary only	3,033	18	3,273	0.460	−1,472.6
B	Tertiary only	3,066	51	3,306	0.454	−1,488.8
C	Secondary + tertiary	3,024	9	3,286	0.463	−1,464.1
D	No education only	3,065	50	3,305	0.455	−1,488.7
E	**No education + secondary + tertiary**	**3,015**	**0**	**3,298**	**0.467**	**−1,455.3**

Notes: $N = 1{,}734$ districts; *Stable* is the base outcome. All columns include log consumption, poverty headcount, working-age share, informality and agricultural shares, density, surface area, and macro-region fixed effects. Primary education is the omitted education category in the full education-decomposition specifications. ΔAIC is measured relative to the lowest-AIC model. Boldface identifies the AIC-preferred specification.

flagged by Psacharopoulos and Patrinos (2018) and is mechanically unable to separate low-skill deficits from high-skill advantages. Table 6 compares the three-share specification against four leaner alternatives. The full decomposition delivers the lowest AIC, with the secondary-plus-tertiary pair a close second; the gap to any single-share specification is large in multinomial-logit terms. Substantively, the three margins carry non-redundant information about re-ranking-group membership. We keep the three-share decomposition as the human-capital block in everything that follows.

Table 7 builds the multinomial-logit model in three blocks: initial conditions (M1); human capital and labour-market composition (M2); and geography with macro-region fixed effects (M3). The largest single improvement comes between M1 and M2: adding human capital and labour-market composition cuts the AIC by 225 points and raises the pseudo-R^2 by five percentage points, more than any other block. Adding exogenous geography and region fixed effects (M3) yields a further, smaller gain and minimises the AIC among these specifications; the BIC, which penalises parameters more sharply, is minimised at M2. We take M3 as our preferred descriptive specification.

Layering municipal and central-government fiscal variables and a productive-capacity block onto M3 produces only marginal further changes in fit; we report these blocks descriptively in Appendix Table 17 and do not interpret them. Central-government social-protection spending in particular is allocated through the same poverty maps that define our outcome and therefore cannot be read as a causal driver. For this reason, any improvement in fit from adding the central-government spending block is not used as evidence for a fiscal mechanism; the substantive interpretation is restricted to the human-capital, labor-market, and geography blocks.

Table 7: Multinomial-logit model of re-ranking-group membership

Model	Block added	AIC	ΔAIC	BIC	ΔBIC	Pseudo R^2
M1	Initial conditions	3,273	258	3,338	94	0.405
M2	+ Human capital & labor	3,048	33	3,244	0	0.455
M3	**+ Geography + region FE**	**3,015**	**0**	**3,298**	**54**	**0.467**

Notes: $N = 1{,}734$ districts; *Stable* is the base outcome. ΔAIC is measured relative to M3; ΔBIC relative to the lowest-BIC model (M2). Boldface identifies our preferred descriptive specification, M3, which adds exogenous geography and macro-region fixed effects to the human-capital and labour core. Fiscal and productive-capacity covariates add little to fit and are reported descriptively, with no causal interpretation, in Appendix Table 17. Huber–White heteroskedasticity-robust standard errors are used throughout.

5.2 Nonlinear education gradients

To characterize nonlinear education gradients, we use two complementary approaches. First, we estimate the nonparametric binscatter of Cattaneo et al. (2024), which partitions the covariate support into IMSE-optimal bins, fits a cubic B-spline through the binned means, and reports 95 percent uniform confidence bands. Departures from linearity are assessed by comparing the linear fit with these bands. Second, we estimate a piecewise-linear specification with a grid search over candidate knots,

$$Y_i = \alpha + \beta_1 X_i + \beta_2 (X_i - \tau) \cdot \mathbf{1}(X_i > \tau) + \varepsilon_i, \tag{5}$$

where τ^* is the knot selected to maximize R^2 over the candidate grid. The estimation sample is restricted to districts in the fifth, poorest poverty quintile in 2007. This restriction is necessary because the Lagging outcome is defined only for Q5 starters: a district must begin in Q5 to be classified as Lagging. Including Q4 districts would therefore mechanically assign 347 observations an outcome of zero, even though those districts are not at risk of entering the Lagging group under our definition. The Q5-only sample ($N = 346$) removes these mechanical zeros and provides the consistent estimation sample throughout section 5.2.

Within the Q5-only subsample ($N = 346$), tertiary education displays the clearest nonlinear association with persistent lagging. The estimated piecewise specification identifies a descriptive kink near 12 percent tertiary attainment ($R^2 = 0.078$; slope-change $p < 0.001$, conditional on the SSR-minimizing τ^*). Below this point, Pr(Lagging) falls by roughly 4.2 percentage points for each additional percentage point of tertiary attainment; above it, the fitted slope is approximately zero. The nonparametric binscatter in the bottom-right panel (Panel D) of Figure 3 is consistent with this pattern: the 95 percent uniform confidence band excludes a linear fit over the central range of the data.

The other education margins show weaker or less stable nonlinearities. In the same Q5-only subsample, secondary attainment exhibits a kink near 42 percent: Pr(Lagging) declines with secondary attainment up to roughly that point and then turns upward beyond it (slope-change $p < 0.001$; R^2 rises from 0.061 to 0.079). The no-education share is broadly positively associated with Lagging-group membership, with a marginal kink near 24 percent ($p = 0.07$): districts with larger shares of residents without formal education are more likely to remain in the Lagging group, although the fitted relationship flattens in the upper tail.

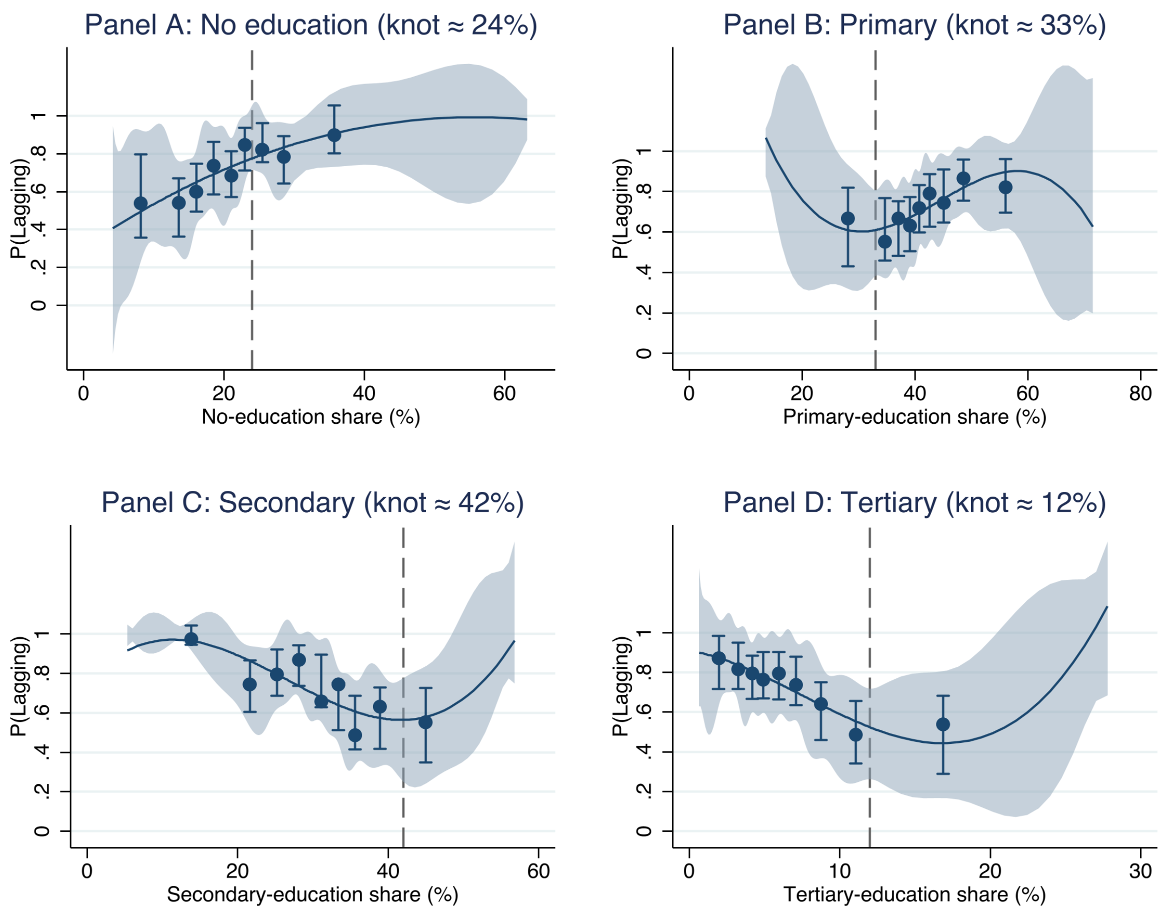

Figure 3: Education composition and persistent lagging among initially high-poverty districts. The sample is restricted to districts in the fifth (poorest) poverty quintile in 2007 ($N = 346$), the consistent sample for which the Lagging outcome is defined. The outcome in all panels is an indicator for remaining in the Lagging group by 2017. Points are binned district-level means; solid curves are cubic B-spline fits; shaded areas are 95 percent confidence bands. Dashed vertical lines mark descriptive knots from piecewise fits. The tertiary-education relationship displays the clearest nonlinear pattern: the probability of remaining Lagging falls sharply below roughly 12 percent tertiary attainment and then flattens. The figure is descriptive and should not be interpreted as evidence of causal or structural thresholds.

6 Spatial Analysis

6.1 Spatial Structure

Own-district characteristics do not fully absorb the geographic structure of relative re-ranking. Appendix Figure 5 maps re-ranking groups alongside the within-sample quintile of neighbors' secondary-education share and shows that Lagging districts overlap with low quintiles of neighbors' secondary education.

Table 8: Spatial co-clustering of education and poverty (queen-contiguity neighborhoods)

Diagnostic	Statistic	Interpretation
Pairwise correlations		
ρ(Own secondary, $W \cdot$ Secondary)	0.319	Spatial clustering of education
ρ(Own poverty headcount, $W \cdot$ Secondary)	-0.165	Poorer districts have less-educated neighbors
Univariate slopes with robust SEs		
Own secondary on $W \cdot$ Secondary	0.204 (0.014) $t = 14.76$	Spatial dependence in education
Own poverty on $W \cdot$ Secondary	-0.266 (0.038) $t = -6.96$	Spatial dependence in poverty

Notes: The spatial weight matrix W is queen-contiguity, defined by shared border or vertex. For the descriptive diagnostics reported here we leave W *unnormalised* so that $W \cdot$Secondary retains the raw spatial-sum scale of Anselin (1995)-style co-clustering; the multinomial-logit regressions in Table 9 instead use a row-normalised W so the lag is the average over contiguous neighbours and is on the same scale as the underlying variable. The qualitative co-clustering patterns reported here are unchanged under row-normalisation. The first block reports Pearson correlations between an own-district variable and the corresponding neighbors' variable. The second block reports the slope coefficient, heteroskedasticity-robust standard error in parentheses, and t-statistic from a univariate regression of an own-district variable on its spatial lag. Positive slopes indicate spatial co-clustering. The diagnostic is on the same scale as the spatial-multinomial-logit coefficients in Table 9, but is not Moran's I because W is unnormalized.

We test spatial associations by augmenting the multinomial and binary logits with neighbors' characteristics: log consumption, secondary education, informality, agricultural employment, and GDP per capita. Neighbors are defined by a queen-contiguity weight matrix (shared border or vertex). Progressive spatial models (A–D) add each neighbor variable sequentially and are compared via AIC and BIC.

Among the four neighbor variables considered, only neighbors' secondary-education share predicts re-ranking-group membership. The spatial lag on $W \cdot$ Secondary educ. share enters positively and significantly in the Low-poverty and Improving margins relative to Stable, while the analogous coefficients for Lagging and Worsening are statistically indistinguishable from zero (Table 9, Panels A and C). Neighbors' consumption, informality, and agricultural employment are uniformly insignificant. Adding neighbors' secondary education alone (Column C of Table 9) yields a modest improvement in fit over the non-spatial baseline, with the spatial-lag term significant at conventional levels.

This specificity—the spatial association operates through human capital, not through income, employment, or informality—is consistent with the education externalities emphasized by Desmet et al. (2026) and Jalan and Ravallion (2002). Conditional on a district's own characteristics, those surrounded by neighbors with lower secondary-education shares are predicted to have worse Low-poverty and Improving outcomes; the spatial association is concentrated in those two group margins (Panels A and C of Table 9) and is statistically zero at the Lagging and Worsening margins.

Table 9: Spatial multinomial logit (row-normalised queen-contiguity W): key spatial-lag coefficients

Outcome vs. Stable	Spatial lag shown	(B)	(C)	(D)
Low-poverty	$W\cdot$ Log consumption	0.067	−0.535	−0.485
		(0.276)	(0.398)	(0.529)
	$W\cdot$ Secondary education	—	8.591**	8.423**
			(3.744)	(3.969)
Lagging	$W\cdot$ Log consumption	−0.020	−0.056	−0.461
		(0.158)	(0.279)	(0.556)
	$W\cdot$ Secondary education	—	0.414	−0.401
			(3.489)	(3.859)
Improving	$W\cdot$ Log consumption	0.132	−0.288	−0.684**
		(0.141)	(0.218)	(0.345)
	$W\cdot$ Secondary education	—	6.093**	6.544***
			(2.422)	(2.478)
Worsening	$W\cdot$ Log consumption	−0.334***	−0.279	−0.350
		(0.115)	(0.191)	(0.309)
	$W\cdot$ Secondary education	—	−0.810	−1.193
			(2.212)	(2.304)
Full non-spatial controls		Yes	Yes	Yes
Additional spatial lags		No	No	Yes
Observations		1,734	1,734	1,734
Pseudo R^2		0.475	0.477	0.481

Notes: The omitted outcome is *Stable*. Models (B)–(D) sequentially add neighbours' log consumption, secondary-education share, and then neighbours' informality, agricultural employment, and GDP per capita. All models include own-district controls and macro-region fixed effects.
W is a row-normalised queen-contiguity matrix: neighbouring districts share a border or vertex, and spatial lags are averages over contiguous neighbours. W was constructed in Stata using `spmatrix create contiguity W, normalize(row)` after `spset`.
Robust standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

6.2 Regional Heterogeneity

Table 10 cross-tabulates each district's own-versus-neighbor secondary-education category against its re-ranking group, showing that "high own, high neighbors" districts concentrate in the Low-poverty group, while "low own, low neighbors" districts concentrate in the Worsening and Lagging groups.

Table 10: Own-neighbor secondary-education categories by re-ranking group

Own-neighbor education category	Low-poverty	Lagging	Improving	Worsening	Stable	Total
High own, high neighbors	76 (35.2%)	13 (6.0%)	40 (18.5%)	46 (21.3%)	41 (19.0%)	216
Low own, low neighbors	4 (1.5%)	62 (23.1%)	37 (13.8%)	106 (39.6%)	59 (22.0%)	268
High own, low neighbors	28 (21.2%)	13 (9.8%)	37 (28.0%)	25 (18.9%)	29 (22.0%)	132
Low own, high neighbors	6 (10.3%)	17 (29.3%)	11 (19.0%)	13 (22.4%)	11 (19.0%)	58
Middle range	136 (12.8%)	143 (13.5%)	250 (23.6%)	310 (29.2%)	221 (20.8%)	1060

Notes: Districts are grouped by combining the standardized own secondary-education share, z_i, with the standardized neighboring secondary-education share, z_{Wi}. "High own, high neighbors" indicates $z_i > 0.5$ and $z_{Wi} > 0.5$; "Low own, low neighbors" indicates $z_i < -0.5$ and $z_{Wi} < -0.5$; "High own, low neighbors" indicates $z_i > 0.5$ and $z_{Wi} < -0.5$; and "Low own, high neighbors" indicates $z_i < -0.5$ and $z_{Wi} > 0.5$. Districts within ± 0.5 on either dimension are classified as "Middle range." Cells report counts and row percentages, where percentages represent the share of districts in each own-neighbor education category that fall in each re-ranking group. This classification is descriptive and is not based on local Moran statistics or formal LISA significance tests.

Table 11 re-estimates the spatial multinomial logit separately by macro-region and shows that the upward (Low-poverty + Improving) association with neighbors' secondary education is largest in the Sierra and attenuates in the Selva.

Table 11: Heterogeneous spatial human-capital associations by macro-region

Outcome (vs. Stable)	**Full**	**Costa**	**Sierra**	**Selva**
Upward (Low-poverty + Improving)				
W·Secondary	7.899***	8.464***	12.088***	2.161
	(1.718)	(2.602)	(3.153)	(9.233)
Downward (Lagging + Worsening)				
W·Secondary	−0.920	−0.592	−1.175	−2.406
	(1.498)	(2.569)	(2.784)	(5.025)
Districts (N)	1,728	720	784	224
Cells (Up / Down / Stable)	625 / 748 / 361	358 / 185 / 177	232 / 414 / 141	35 / 149 / 43

Notes: Coefficients on $W \cdot$ Secondary from a three-category multinomial logit estimated separately for each region. The dependent variable collapses the five re-ranking groups into three re-ranking categories: *Upward* (Low-poverty + Improving), *Downward* (Lagging + Worsening), and *Stable* (base). The collapse keeps every cell large enough for stable identification; Selva's smallest cell becomes Upward ($N = 35$) rather than Lagging ($N = 7$ in the original five-group specification). Each regression includes the full Model C control set: own log consumption, own poverty headcount, working-age population share, the three-tier education decomposition (no education, secondary, tertiary; primary is the omitted base), informality and agricultural employment shares, population density, surface area, municipal current expenditure, budget execution rate, district GDP per capita, neighbors' log consumption ($W \cdot \log c$), and macro-region fixed effects (omitted in the regional regressions). Standard errors in parentheses are heteroskedasticity-robust. $^{*}\, p < 0.10$, $^{**}\, p < 0.05$, $^{***}\, p < 0.01$.

The spatial human-capital association is concentrated in the Sierra and Costa and is imprecise in the Selva.

6.3 Migration and Sorting

The spatial results in Section 6 leave one interpretation open. Because we read the coefficient on neighbors' education as a spatial human-capital association rather than as pure spillovers, we need to know whether that gradient reflects local complementarities or simply the spatial footprint of who moves where. Internal migration offers a direct test, and it is well suited to the convergence framing of the paper: the flows we use are measured over 2002–2007 from the 2007 Census, and are therefore predetermined relative to the 2007–2017 window over which we measure β-convergence in Section 4. If the persistently poor districts are losing population, and in particular losing their more educated residents, then the human-capital gradient documented in Section 5 is partly sustained by sorting rather than by local accumulation alone—and the failure of the bottom tail to catch up acquires a demographic, not only a structural, explanation. Tables 12 and 13 and Figure 4 examine the two margins of this process: the net movement of people and the net movement of skill.

Table 12: Net flows of people versus net flows of skill, by quintile

	Net counts		Per 100 pop ≥5	
	People	Secondary+ adults	People	Secondary+ adults
Low-poverty	+302,208	+25,837	+2.4	+0.2
Lagging	-94,244	-6,527	-6.4	-0.4
Improving	-67,674	-8,686	-3.7	-0.5
Worsening	-78,990	-3,080	-1.6	-0.1
Stable	-43,528	-3,620	-1.3	-0.1
All districts	+17,772	+3,924	+0.1	+0.0

Notes: Internal (within-Peru) migrants, 2002–2007, from the 2007 Census. Net people is in-migrants minus out-migrants; net secondary+ adults is the same net flow counting only adults aged 25+ with secondary education or more. Positive values are net gains. A group can gain people on net while losing skill, or the reverse: the neoclassical (body) and sorting (skill) margins need not move together.

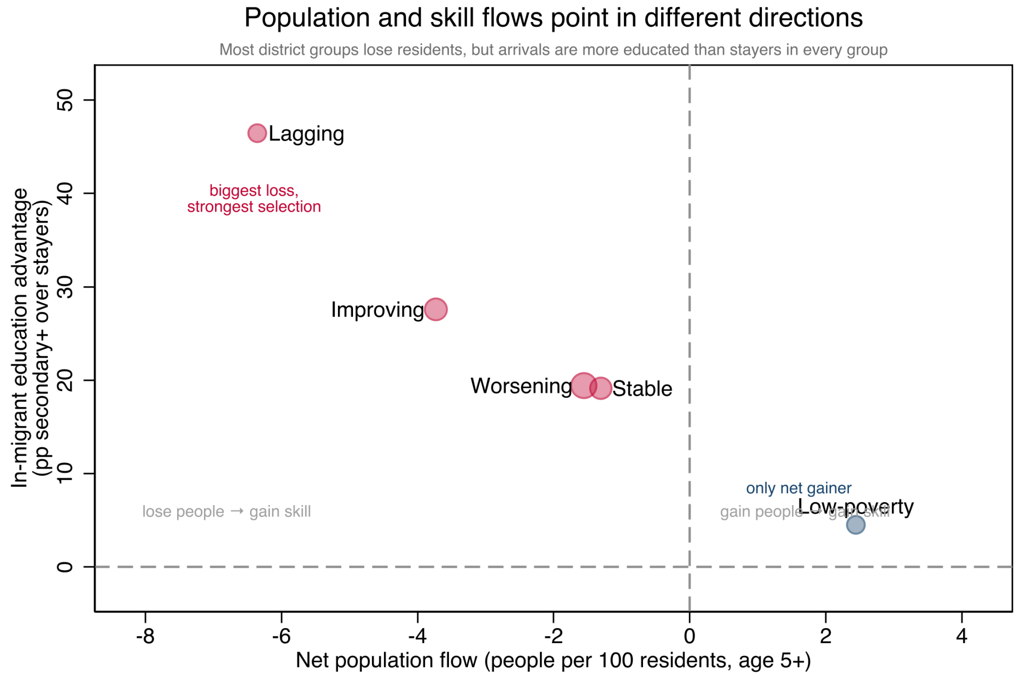

Internal migration 2002-2007 (CPV2007), predetermined relative to the 2007-2017 convergence window.
One point per district mobility group; marker area proportional to the number of districts. Population-weighted within group.

Figure 4: Population and skill flows point in different directions, 2002–2007. Each point is one re-ranking group the horizontal axis is the net internal migration rate (in- minus out-migrants per 100 residents aged 5+), the vertical axis is the percentage-point gap in the secondary-or-more share between in-migrants and stayers, and marker area is proportional to the number of districts. Internal migration is measured from the 2007 Census and is predetermined relative to the 2007–2017 convergence window. Every group lies above zero on the vertical axis (arrivals are positively selected everywhere), while most lie left of zero on the horizontal axis (net loss of residents).

Table 13: How much of each group's adult skill stock is imported by in-migration

	In-migrant share of adults (%)	Stayers-only secondary+ (%)	Resident secondary+ (%)	Skill imported by sorting (p.p.)
Low-poverty	15.5	79.4	80.1	+0.7
Lagging	3.9	20.9	22.7	+1.8
Improving	8.8	41.8	44.2	+2.4
Worsening	10.9	45.6	47.8	+2.1
Stable	10.7	46.8	48.9	+2.1
All districts	12.9	62.0	64.0	+2.0

Notes: Adults aged 25+ in the 2007 Census, population-weighted within group. The stayers-only column is the secondary-or-more share computed from non-migrant adults alone (the counterfactual skill stock absent in-migration); the resident column adds in-migrant adults. Skill imported by sorting is the resident minus stayers-only gap: the percentage-point lift in measured adult skill that in-migration mechanically delivers to the group. The lift is largest where in-migrant inflows are largest, so skill-selective sorting reinforces – rather than offsets – the cross-district skill gradient.

Table 12 sets the net flow of people against the net flow of skill for each re-ranking group. The contrast is stark and runs along the convergence ranking. The Low-poverty group is the only net recipient on either margin, gaining roughly 302,000 residents and 26,000 secondary-educated adults over 2002–2007 (about +2.4 and +0.2 per 100 residents aged 5+). Every other group loses on both counts. The Lagging districts—the persistently poorest tail—record the largest proportional outflow, shedding about 6.4 residents and 0.4 secondary-educated adults per 100, while the Improving, Worsening, and Stable groups lose at more moderate rates. Because national net internal migration is close to zero (+0.1 per 100 overall), these are redistributive flows: people and skill are not entering or leaving Peru so much as draining from the poorest districts toward the already better-off ones. Figure 4 summarizes the two margins jointly. Every group sits above zero on the vertical axis, so arrivals are positively selected on education everywhere, including in districts that are losing population overall; most groups sit to the left of zero on the horizontal axis, confirming net outflow. The poorest groups thus combine the worst of both: they lose people, and the relatively educated arrivals they do attract are not numerous enough to offset the departures.

Table 13 asks how much of each group's measured adult skill stock is imported through in-migration rather than produced locally. In-migrants make up very different shares of the adult population across groups—15.5% in Low-poverty districts but only 3.9% in Lagging ones—and they are far more educated where they are most numerous: the in-migrant secondary-or-more share is 79.4% in Low-poverty districts against 20.9% in Lagging districts. Comparing the resident secondary-or-more share with the counterfactual that strips in-migrants out (the stayers-only column) isolates the skill that sorting mechanically delivers. That lift is positive for every group and, crucially, is largest precisely where inflows are largest: in-migration raises the measured secondary-or-more share by about 2.4 percentage points in the Improving group and 2.0 overall, but by far less at the bottom. Skill-selective sorting therefore reinforces, rather than offsets, the cross-district skill gradient documented in Section 5. Read together with the spatial estimates of Section 6, this is why we do not interpret the neighboring-education coefficient as a pure spillover: the same gradient is partly built by who moves where. Reassuringly, the β-convergence result of Section 4 is not an artifact of this churn—it survives in the low-re-ranking subsamples reported in Appendix Table 20—so sorting shapes the level of the human-capital gradient without overturning the conditional catch-up that anchors the paper.

7 Robustness

7.1 Ordinal re-ranking (quintile transitions)

The headline re-ranking results hinge on three choices: the multinomial encoding of the outcome, the cross-sectional unit of analysis, and the Q5-to-Q4 classification rule. We probe each in turn—an ordered logit on the nine-category quintile-transition index, a within-group examination of dispersion, and a Lagging-group decomposition that isolates the correlates of absolute progress in a relatively persistent group.

Treating quintile transitions as ordered rather than categorical is the most natural alternative encoding, since it imposes monotonicity in the underlying re-ranking direction and discards the within-direction information the multinomial logit uses. The ordered logit on the nine-category quintile-transition index (-4 to $+4$), reported in Table 14, yields consistent sign patterns: initial poverty, the working-age share, and secondary education (relative to no-education as the omitted base) are the dominant predictors, while budget execution again does not enter significantly. The positive sign on initial poverty is partly mechanical: districts starting at Q5 can only stay at Q5 or move to lower quintiles, so the baseline-poverty coefficient absorbs the floor at the bottom of the rank distribution.

Table 14: Ordered-logit estimates of district poverty re-ranking, 2007–2017

	Poverty re-ranking
Initial conditions	
Poverty headcount, 2007	11.528*** (0.529)
Log consumption p.c., 2007	2.447*** (0.355)
Working-age population share	11.019*** (1.358)
Education shares	
Primary education	−4.772*** (0.975)
Secondary education	3.715*** (0.901)
Tertiary education	−0.492 (1.259)
Labor market and geography	
Informal employment share	−0.218 (0.549)
Agricultural employment share	1.300*** (0.361)
Population density	0.298 (0.186)
State capacity	
Budget execution rate	0.109 (0.159)
Fiscal controls	Yes
District economic controls	Yes
Region fixed effects	Yes
Observations	1,734
Log-likelihood	−2,255.48
Pseudo R^2	0.176

Notes: Ordered-logit coefficients are reported, with robust standard errors in parentheses. The dependent variable is $\text{mobility}_i = Q_{2007,i} - Q_{2017,i}$, where positive values indicate upward movement in the ranking. No education is the omitted education category; Costa is the omitted macro-region. Fiscal controls include municipal expenditure and investment (winsorised at the 99th percentile), central-government non-extractive-revenue categories (current and investment), municipal income, and the budget execution rate. District economic controls include surface area, district GDP per capita, financial assets per capita, and total revenue per capita. Eight cut points are estimated and suppressed for readability. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 14 summarises the ordered-logit estimates for the nine-category quintile-transition index across 1,734 districts. The dependent variable is signed so that positive values indicate upward re-ranking, so positive coefficients identify variables associated with upward movement in the poverty-quintile ranking and negative coefficients with downward movement. Education enters as three indicators with the no-education category omitted, while macro-region fixed effects, fiscal controls, and district economic controls are partialled out; the eight estimated cut points are suppressed for readability. Robust standard errors appear in parentheses, and the pseudo-R^2 of 0.176 together with the log-likelihood provide the usual goodness-of-fit benchmarks for the ordered specification.

7.2 Within-group dispersion

The re-ranking group classification also masks important within-group heterogeneity. Appendix Table 15 reports the within-group standard deviation of poverty rates in 2007 and 2017. The overall distribution exhibits σ-convergence, but the Lagging group becomes more internally dispersed over the decade, indicating that absolute improvement within the lagging group was itself heterogeneous. This reinforces an important scope condition of the group classification: the Lagging label identifies persistent relative position, not a uniform absolute trajectory among all districts in the group. To unpack this internal dispersion, Table 16 compares baseline characteristics of Lagging districts that achieved greater versus lower absolute progress and shows that within-Lagging improvers start with higher secondary-education shares and lower informality than within-Lagging stuck districts.

Table 15: σ-convergence: standard deviation of poverty rates by group

	SD (2007)	**SD (2017)**	**Change**	**σ-conv.?**	N
Low-poverty	0.080	0.029	−0.051	Yes	250
Lagging	0.041	0.124	+0.083	No	248
Improving	0.159	0.080	−0.079	Yes	375
Worsening	0.169	0.184	+0.015	No	500
Stable	0.139	0.094	−0.045	Yes	361

Note: Standard deviation of district-level poverty headcount rates within each re-ranking group. The overall distribution compressed (σ-convergence confirmed by variance-ratio F-test, $p < 0.01$). The lagging group became more dispersed internally, indicating heterogeneous absolute improvement within the group. Source: ENAHO 2007, 2017 and population census 2007, 2017.

7.3 Within-Lagging heterogeneity

Even within the Lagging-group districts differ substantially in their absolute progress. Splitting the 248 Lagging districts at the within-group quartiles of 2017 poverty isolates 62 "Lagging improvers" (largest absolute reductions, yet still in the bottom two quintiles in 2017) and 62 "Lagging stuck" districts (smallest absolute progress). Table 16 compares their 2007 baselines. Improvers entered the decade with marginally lower poverty headcounts, materially higher secondary- and tertiary-education shares, larger working-age population shares, and lower agricultural-employment shares than stuck districts. Informal-employment shares, population density, and the budget-execution rate are, by contrast, statistically indistinguishable across the two groups. Even among persistently lagging districts, stronger baseline education distinguishes districts with larger absolute poverty reductions.

Table 16: Within-Lagging heterogeneity: baseline characteristics of districts with greater versus lower absolute progress

	Lagging improvers mean	Lagging stuck mean	Difference Stuck − Improver	p-value
Initial and final poverty				
Poverty rate, 2007	0.852	0.879	0.027***	0.000
Poverty rate, 2017	0.337	0.620	0.283***	0.000
Education composition, 2007				
No education share	0.197	0.236	0.040**	0.013
Secondary education share	0.323	0.257	-0.066***	0.000
Tertiary education share	0.074	0.055	-0.019**	0.022
Labor market and demographics, 2007				
Working-age population share	0.621	0.588	-0.033***	0.001
Informal employment share	0.851	0.873	0.022	0.216
Agricultural employment share	0.718	0.763	0.045*	0.066
Population density	0.079	0.119	0.041	0.175
Fiscal capacity, 2007				
Budget execution rate	0.427	0.393	-0.033	0.604
Districts, N	62	62		

Notes: The sample is restricted to the Lagging group, defined as districts that began in the poorest poverty quintile in 2007 and remained in the bottom two quintiles in 2017. Within this group, districts are split by the within-group quartile of their 2017 poverty rate. *Lagging improvers* are districts in the bottom quartile of 2017 poverty within the Lagging group, corresponding to the largest absolute poverty reductions while still failing to escape the relative bottom. *Lagging stuck* districts are in the top quartile, corresponding to the smallest absolute progress. The third column reports the difference in means, calculated as Lagging stuck minus Lagging improvers. The fourth column reports the p-value from a two-sample t-test. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

8 Discussion and Conclusion

Peru's 2007–2017 poverty decline was broad, but it did not eliminate the territorial disparities of poverty. Many districts experienced large absolute reductions while remaining near the bottom of the national poverty distribution. The central lesson is therefore empirical rather than prescriptive: aggregate growth and generalized poverty reduction are not sufficient, by themselves, to generate upward movement in the relative poverty distribution. Districts differ in their ability to translate national progress into movements in the relative ranking, and those differences are closely associated with local and neighbouring human-capital conditions.

The evidence points to three policy-relevant patterns that should be interpreted as descriptive and predictive rather than causal. First, education composition predicts movements in the relative ranking more consistently than the aggregate municipal fiscal variables available in our data. This does not imply that public spending is unimportant; rather, it suggests that aggregate measures of spending capacity may not capture whether public resources address the constraints most relevant for poor districts. Second, secondary education is the most stable human-capital correlate of upward re-ranking across specifications. For lagging districts, secondary completion, dropout reduction, and school-to-work transitions therefore emerge as promising margins for future causal evaluation and policy experimentation, rather than as effects identified directly by this paper. Third, neighbouring secondary education is the only spatial lag that robustly predicts favourable re-ranking. This suggests that future work should examine whether human-capital policies operate at scales larger than individual districts, including local labour-market areas, school catchments, and contiguous clusters of poor districts.

The migration evidence qualifies this spatial interpretation. Predetermined migration flows show that population and skill flows do not move in the same way: most re-ranking groups lose residents on net, but in-migrants are more educated than stayers in every group. This pattern is consistent with a mixed process. Population losses from poorer or lagging areas are consistent with neoclassical migration away from disadvantaged districts, while the positive education selectivity of arrivals points to skill-selective sorting. The neighbouring-education coefficients should therefore not be interpreted as pure spillovers. A more cautious interpretation is that they capture territorial human-capital environments shaped by both local complementarities and selective migration.

The nonlinear tertiary-education pattern further qualifies the interpretation. Tertiary attainment distinguishes low-poverty districts and is associated with a lower probability of persistent lagging within the initially high-poverty subsample, but the estimated kink near 12 percent should not be read as a structural threshold or as a policy target. A more cautious interpretation is that very low tertiary attainment marks districts with limited professional, technical, and administrative capacity. The result points less to a narrow tertiary-enrolment prescription than to the need for future research on local skills ladders linking basic education, secondary completion, technical training, and pathways into higher education.

The fiscal results should be read with similar caution. Municipal spending, investment, and budget execution do not consistently predict upward re-ranking once baseline socioeconomic conditions are included. This finding is not evidence against decentralization or local public investment. Instead, it indicates that aggregate fiscal indicators may be too coarse to capture whether spending reaches the constraints that matter for poor districts. Future policy evaluation should therefore move beyond total spending and execution rates toward program-level measures of school quality, teacher allocation, transport connectivity, health access, social-protection coverage, and the complementarity between local and central spending.

Finally, the relative re-ranking classification provides a simple diagnostic device for territorial analysis. Poverty maps can be used not only to identify poor districts at a point in time, but also to distinguish districts that move upward in the poverty distribution from those that improve in absolute terms while remaining persistently behind. This distinction matters for research and targeting design, but the groups should not be interpreted as internally homogeneous re-ranking groups. The Lagging label marks persistent relative position, not a common behavioral path: even as the national distribution compresses, within-Lagging dispersion rises, and districts with stronger initial educational endowments experience larger absolute poverty reductions.

These results identify margins where future causal and policy work should focus. The post-2017 period, including the COVID-19 shock and recovery, offers one setting to examine whether baseline human-capital composition also predicts resilience under adverse aggregate conditions. Applying the same framework to other decentralizing economies in Latin America would show whether Peru's pattern is country-specific or reflects a broader regional association between local human capital, selective migration, and movements in the relative ranking.

References

Anselin, L. (1995). Local indicators of spatial association—LISA. *Geographical Analysis*, 27(2):93–115.

Azariadis, C. and Drazen, A. (1990). Threshold externalities in economic development. *Quarterly Journal of Economics*, 105(2):501–526.

Bardhan, P. (2002). Decentralization of governance and development. *Journal of Economic Perspectives*, 16(4):185–205.

Barro, R. J. (2000). Inequality and growth in a panel of countries. *Journal of Economic Growth*, 5(1):5–32.

Barro, R. J. (2001). Education and economic growth. In Helliwell, J. F., editor, *The Contribution of Human and Social Capital to Sustained Economic Growth and Well-Being*, pages 14–41. OECD, Paris.

Barro, R. J. and Sala-i Martin, X. (1992). Convergence. *Journal of Political Economy*, 100(2):223–251.

Bourguignon, F. (2003). The growth elasticity of poverty reduction: Explaining heterogeneity across countries and time periods. In Eicher, T. S. and Turnovsky, S. J., editors, *Inequality and Growth: Theory and Policy Implications*, CESifo Seminar Series, pages 3–26. MIT Press, London.

Canavire-Bacarreza, G., Jetter, M., and Robles, M. (2018). When does economic growth reduce poverty and strengthen the middle class? A state-level, sector-specific analysis of Peru. *Southern Economic Journal*.

Cattaneo, M. D., Crump, R. K., Farrell, M. H., and Feng, Y. (2024). On binscatter. *American Economic Review*, 114(5):1488–1514.

Conley, T. and Ligon, E. (2002). Economic distance and cross-country spillovers. *Journal of Economic Growth*.

Crespo Cuaresma, J., Klasen, S., and Wacker, K. M. (2022). When do we see poverty convergence? *Oxford Bulletin of Economics and Statistics*, 84(6):1283–1301.

Cuenca, R., Gil, R., Sanchez, D., and Guadalupe, R. (2022). *Descentralización: El fracaso de la única gran reforma institucional del estado en el Perú del siglo XXI*. IEP Instituto de Estudios Peruanos.

Desmet, K., Nagy, D. K., and Rossi-Hansberg, E. (2026). Human capital accumulation across space. Working Paper, SMU/CREI/University of Chicago.

Faguet, J.-P. (2014). Decentralization and governance. *World Development*, 53:2–13.

Gennaioli, N., La Porta, R., Lopez-de Silanes, F., and Shleifer, A. (2013). Human capital and regional development. *Quarterly Journal of Economics*, 128(1):105–164.

Jalan, J. and Ravallion, M. (1998). Are there dynamic gains from a poor-area development program? *Journal of Public Economics*, 67(1):65–85.

Jalan, J. and Ravallion, M. (2002). Geographic poverty traps? A micro model of consumption growth in rural China. *Journal of Applied Econometrics*, 17(4):329–346.

Kim, R., Mohanty, S. K., and Subramanian, S. V. (2016). Multilevel geographies of poverty in india. *World Development*, 87:349–359.

Kraay, A. (2006). When is growth pro-poor? evidence from a panel of countries. *Journal of Development Economics*, 80(1):198–227.

Kremer, M. (1993). The O-ring theory of economic development. *Quarterly Journal of Economics*, 108(3):551–575.

López-Calva, L. F., Ortiz-Juarez, E., and Rodríguez-Castelán, C. (2022). Within-country poverty convergence: Evidence from Mexico. *Empirical Economics*, 62(5):2547–2586.

Marrero, Á. S., Marrero, G. A., and Servén, L. (2025). Poverty convergence clubs. *Review of Income and Wealth*, 71(1):e12688.

Miranti, R. (2021). Is regional poverty converging across Indonesian districts? A distribution dynamics and spatial econometric approach. *Asia-Pacific Journal of Regional Science*, 5(3):851–883.

Najam, Z. and Gibson, J. (2022). Does intra-country poverty convergence depend on spatial spillovers and the type of poverty measure? Evidence from Pakistan. *Asia and the Pacific Policy Studies*, 9(3):516–535.

Ouyang, Y., Shimeles, A., and Thorbecke, E. (2019). Revisiting cross-country poverty convergence in the developing world with a special focus on Sub-Saharan Africa. *World Development*, 117:13–28.

Phillips, P. C. B. and Sul, D. (2007). Transition modeling and econometric convergence tests. *Econometrica*, 75(6):1771–1855.

Psacharopoulos, G. and Patrinos, H. A. (2018). Returns to investment in education: A decennial review of the global literature. *Education Economics*, 26(5):445–458.

Quah, D. T. (1993). Galton's fallacy and tests of the convergence hypothesis. *Scandinavian Journal of Economics*, 95(4):427–443.

Quah, D. T. (1996). Twin peaks: Growth and convergence in models of distribution dynamics. *Economic Journal*, 106(437):1045–1055.

Ravallion, M. (1997). Can high-inequality developing countries escape absolute poverty? *Economics Letters*, 56(1):51–57.

Ravallion, M. (2001). Growth, inequality and poverty: Looking beyond averages. *World Development*, 29(11):1803–1815.

Ravallion, M. (2012). Why don't we see poverty convergence? *American Economic Review*, 102(1):504–523.

Rey, S. and Montouri, B. (1999). U.s. regional income convergence: A spatial econometric perspective. *Regional Studies*.

Solow, R. M. (1956). A contribution to the theory of economic growth. *Quarterly Journal of Economics*, 70(1):65–94.

World Bank (2023). *Rising Strong: Peru Poverty and Equity Assessment.* The World Bank, Washington, D.C.

Yohai, V. J. (1987). High breakdown-point and high efficiency robust estimates for regression. *The Annals of Statistics*, 15(2):642–656.

Appendix

Table 17: Descriptive extension: fiscal and productive-capacity covariates

Model	Block added	AIC	ΔAIC	BIC	ΔBIC	Pseudo R^2
M1	Initial conditions	3,273	277	3,338	94	0.405
M2	+ Human capital & labor	3,048	52	3,244	0	0.455
M3	+ Geography + region FE	3,015	19	3,298	54	0.467
M4	+ Municipal spending	3,025	29	3,352	108	0.468
M5	+ Central-government spending	2,996	0	3,432	188	0.480
M6	+ Municipal income	3,006	10	3,486	242	0.482
M7	+ Productive capacity	3,012	16	3,558	314	0.485

Notes: N = 1,734 districts; *Stable* is the base outcome. This table extends the interpretive model (M3, Table 7) with fiscal and productive-capacity blocks and is reported for descriptive completeness only. **None of the M4–M7 coefficients can be read causally.** Central-government social-protection spending is included only as part of this descriptive appendix extension and is not interpreted as a mechanism. This restriction is important because Juntos and Pensión 65 are geographically targeted using the same poverty maps that define the outcome quintiles, making the social-protection regressor mechanically related to the dependent variable. The lower AIC in M5 should therefore be read only as an in-sample association, not as evidence that central-government social-protection spending predicts relative re-ranking in a behavioral or causal sense. ΔAIC is measured relative to the lowest-AIC model (M5); ΔBIC is measured relative to the lowest-BIC model (M2). Huber–White heteroskedasticity-robust standard errors are used throughout.

Table 18: Robustness: Multinomial logit under alternative re-ranking-group classification (Q5→Q4 reclassified as Improving)

	Lagging vs. Stable		Improving vs. Stable	
	Main rule (1)	Alt rule (2)	Main rule (3)	Alt rule (4)
Poverty rate, 2007	30.158***	25.911***	10.338***	11.513***
	(2.365)	(2.405)	(0.975)	(0.969)
No-education share, 2007	3.537*	0.743	1.129	2.198*
	(1.912)	(2.451)	(1.430)	(1.334)
Secondary-education share, 2007	2.163	-3.690	2.797**	3.433***
	(2.148)	(2.468)	(1.327)	(1.264)
Tertiary-education share, 2007	1.475	1.802	-1.105	0.056
	(3.565)	(4.916)	(1.553)	(1.536)
Working-age population share, 2007	0.344	2.298	15.274***	12.528***
	(2.912)	(3.719)	(1.703)	(1.609)
Group districts (N)	248	109	375	514
Total observations	1,734		1,734	
Log-likelihood	−1,487.8		−1,503.0	
AIC	3,048		3,078	

Notes: Multinomial-logit coefficients with Stable as the base outcome. The right-hand side includes initial conditions (log per-capita consumption and the 2007 poverty headcount, both held constant across columns and suppressed in the display) and the human-capital and labour-market block: the working-age population share, the no-education / secondary / tertiary education shares, and the informal-employment and agricultural-employment shares (the latter two also suppressed). Columns (1) and (3) use the classification rule of Section 3.3, in which Q5→Q4 transitions are coded as Lagging. Columns (2) and (4) reclassify Q5→Q4 transitions as Improving, restricting the Lagging group to Q5→Q5 transitions. The right-hand side is identical across columns; only the dependent-variable encoding changes.

We deliberately omit the municipal-fiscal, central-government, productive-capacity, and department-fixed-effect blocks of the full Model 7 of Table 7: under the alternative rule the Lagging cell shrinks from 248 to 109 districts and the full Model 7 does not converge. Estimating Model 2 instead (initial conditions + human capital + labour market) preserves the headline re-ranking predictors with a parameter count the smaller Lagging cell can support.

Huber–White heteroskedasticity-robust standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 19: Piecewise-linear education kinks on the Q5-only sample

Tier	Knot τ^*	Pre-slope	Slope-change	p	R^2 (lin $\to$ pw)
No education	$\approx 24\%$	0.020	-0.013	0.067	$0.066 \to 0.072$
Primary	$\approx 33\%$	-0.018	0.030	0.003	$0.019 \to 0.034$
Secondary	$\approx 42\%$	-0.016	0.042	< 0.001	$0.061 \to 0.079$
Tertiary	$\approx 12\%$	-0.042	0.062	< 0.001	$0.048 \to 0.078$

Notes: Constrained piecewise-linear LPM, $\Pr(\text{Lagging}) = \alpha + \beta_1 X + \beta_2 \max(0, X - \tau^*) + \varepsilon$, estimated on the consistent Q5-only subsample ($Q_{2007} = 5$, $N = 346$). The knot τ^* is chosen by grid search maximizing R^2 over the central support (p_{10}–p_{90} of each education share). "Pre-slope" is β_1 (change in Pr(Lagging) per percentage point below τ^*); "Slope-change" is β_2. The p-value tests $\beta_2 = 0$ (robust SEs). The final column reports the R^2 of the linear-only model versus the piecewise fit.

Table 20: Convergence and Migration

	Full	Low $\|net\|$ (bottom 50%)	$\|net\| \leq 2$ (per 100)	Low gross (bottom 50%)
Panel A. Log-rate convergence (dep. var. g; coef. on $\ln H_{07}$*)*				
β	-0.0446^{***}	-0.0430^{***}	-0.0677^{***}	-0.0503^{***}
	(0.0078)	(0.0123)	(0.0136)	(0.0085)
Panel B. Percentage-point convergence (dep. var. Δpov*; coef. on* H_{07}*)*				
β	-0.0835^{***}	-0.0837^{***}	-0.0916^{***}	-0.0805^{***}
	(0.0027)	(0.0042)	(0.0068)	(0.0039)
Districts	1,728	864	275	864

Notes: Each cell is the initial-poverty coefficient from a separate β-convergence regression with the full control set (baseline log consumption, education shares, agricultural employment, department fixed effects) and heteroskedasticity-robust standard errors. A negative coefficient indicates convergence. Columns restrict the sample by predetermined 2002–2007 internal-migration churn: net migration rate below its median, net rate within ± 2 per 100 residents, and gross turnover (in- plus out-migration rate) below its median. Migration is measured before the 2007–2017 convergence window. The convergence coefficient is stable across subsamples and, if anything, larger where migration is smallest, so the estimated convergence is not driven by population turnover.

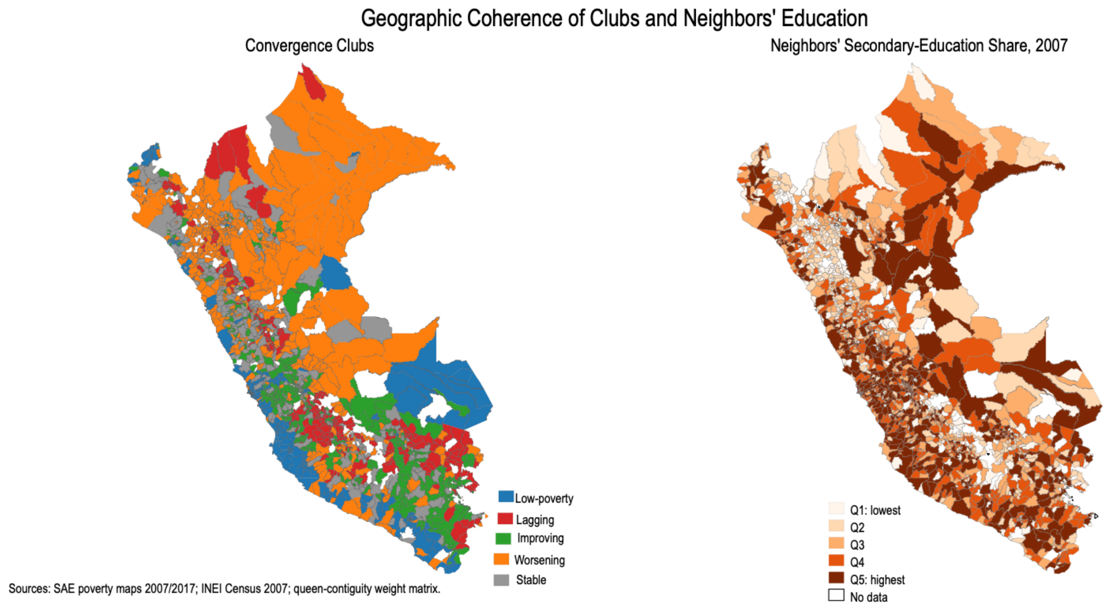

Figure 5: Geographic coherence of re-ranking groups and neighbors' secondary education, 2007. The left panel maps the five re-ranking groups (Low-poverty, Lagging, Improving, Worsening, Stable) constructed from district poverty-quintile transitions between 2007 and 2017. The right panel maps the within-sample quintile of neighbors' secondary-education share, $W{\cdot}\text{Secondary}_{2007}$. The spatial weight matrix W is the row-normalised queen-contiguity weight matrix (shared border or vertex), and the lag is the average over contiguous neighbours' values. Districts in the southern Sierra and central coast that fall in the Improving and Low-poverty groups are concentrated in the highest quintile of neighbors' secondary education; Lagging districts in the northern Andes overlap with low quintiles of neighbors' secondary education.