Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Convergence in Pace, Not in Outcomes: A Multilevel Growth Curve Analysis of Child Mortality Across 179 Countries

Ayşe Günsel

ayse.gunsel@kocaeli.edu.tr

Kocaeli University

Erkut Altındağ

Doğuş University

Gönül Konakay

Kocaeli University

Research Article

Keywords: Under-five mortality, Child mortality, Convergence, Health inequalities, Multilevel models, Growth curve modelling, Global health, Sustainable Development Goals

Posted Date: November 10th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11116273/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

Background

The global decline in under-five mortality since 1990 is well documented in aggregate, but aggregate trends are silent about the distribution of national experience. Whether countries with the highest mortality caught up, and whether the gap between countries narrowed, are different questions, and the child-survival agenda from the fourth Millennium Development Goal to Sustainable Development Goal target 3.2 has been framed around aggregate targets that cannot distinguish them.

Methods

We modelled annual under-five mortality for 179 countries over 1990 to 2019 (5,370 country-years; World Bank series disseminated from UN IGME estimates) with a multilevel growth curve on the logarithmic scale: country-specific random intercepts and slopes, with standardized 1990 income as a country-level predictor. Convergence in pace, the association between initial level and subsequent proportional decline, was estimated from the random-effects covariance (with cluster-bootstrap confidence intervals) and from regressions of observed change on observed baselines, including a strictly non-overlapping baseline design; convergence in outcomes was assessed from the year-by-year cross-country standard deviation of log mortality, with a within- and between-region decomposition. Robustness analyses covered piecewise and random-quadratic trends, autoregressive residuals, population-averaged estimators, and sensitivity to the time origin, baseline definition, and added measurement error.

Results

Mortality fell by 30.7% per decade on average (95% CI 28.9 to 32.4). South Asia declined fastest and sub-Saharan Africa combined the highest burden with a near-average pace. Countries with higher 1990 mortality declined proportionally faster: the intercept-slope correlation was − 0.24 (bootstrap 95% CI − 0.37 to − 0.11), essentially unchanged without income adjustment, and the non-overlapping baseline regression gave − 0.034 per decade (95% CI − 0.054 to − 0.013). Yet the cross-country dispersion of log mortality showed no net reduction, rising from 1.05 to a 1999 peak of 1.14 before receding to 1.09 in 2019; dispersion within regions fell while dispersion between regions rose. The direction and approximate magnitude of the average decline were robust across error structures, while the random-effects decomposition was sensitive to the residual specification and inference on the income-by-time interaction varied across estimators. The weak income gradient in rates of decline is consistent with, but does not directly demonstrate, the diffusion of broadly accessible health interventions.

Conclusions

Progress in child survival over 1990 to 2019 was convergent in pace but not, on net, in outcomes, and the missing narrowing lies between world regions rather than within them. Monitoring frameworks that track only aggregate averages will not register this distributional shortfall.

Background

The decline of child mortality is among the signal achievements of the past three decades, with the global rate of death before the fifth birthday falling by more than half since 1990, a transformation that has saved tens of millions of lives and reshaped the demography of the developing world [1, 2, 3]. Yet the global average conceals a multitude of national trajectories, some countries driving mortality down rapidly while others lagged or suffered reversals, and the average rate of decline tells us nothing about whether countries grew more alike or more unequal as the decline unfolded [4, 5, 6]. Understanding how child survival evolved therefore requires modeling not a single trend but a population of trajectories, characterizing both the typical path and the variation around it, and asking whether the countries that began with the highest mortality reduced it fastest or fell further behind [7, 8, 9].

This study models the trajectory of under-five mortality for 179 countries from 1990 to 2019 using a multilevel growth curve, a standard method for repeated measurements nested within units [7, 10, 11]. The model treats each country's sequence of annual mortality values as a trajectory with its own starting level and rate of change, represents these as random effects drawn from a population distribution, and estimates simultaneously the average trajectory, the variation in starting levels and in rates of decline, and the association between where a country began and how fast it moved [12, 13, 14]. That association is the heart of the convergence question: if the countries that began with the highest mortality declined proportionally fastest, it is negative [5, 6, 8].

The contributions are substantive; the paper claims no methodological novelty. First, the study quantifies the average pace of decline, roughly thirty-one percent per decade, within a single coherent model, and shows that this estimate is stable across every specification examined. Second, it documents wide heterogeneity in national trajectories and maps it geographically, distinguishing variation between the seven World Bank regions from variation among neighbors within them. Third, and centrally, it returns a divided convergence verdict: countries with the highest 1990 mortality declined proportionally faster, yet the cross-country dispersion of mortality shows no net reduction over three decades, and a decomposition locates the missing narrowing between regions rather than within them. This pattern, catch-up in pace without net compression of outcomes, qualifies the narrative of a grand convergence in global health [3, 6, 9]. Throughout, the analysis distinguishes descriptive from inferential statements, treats the modeled character of the underlying mortality estimates as a central feature of the evidence, and bounds each claim by the analysis that supports it.

The stakes of distinguishing how mortality fell, and not merely that it fell, are both scientific and practical. A declining average is compatible with uniform improvement, with laggards catching up, and with polarization, and these patterns carry different implications for the inequality of survival across

countries and for policy [5, 6, 15]. The international community framed its child-survival agenda, from the fourth Millennium Development Goal to Sustainable Development Goal target 3.2, around aggregate mortality targets, and aggregate targets are silent on distribution: a target met on average can coexist with a gap that fails to close [1, 3, 5]. Recovering the distribution of progress, and the relation between starting points and subsequent change, is what a model of trajectories provides and a model of the average cannot [7, 11].

The global decline in child mortality

The sustained decline of child mortality across the developing world has been attributed to a confluence of forces, including the diffusion of low-cost health technologies such as immunization, oral rehydration, and insecticide-treated nets, rising incomes and education, improvements in water and sanitation, and the expansion of basic health services, often supported by international assistance [16, 17, 18]. The relation between mortality and income, long described by the Preston curve, implies that economic growth lowers mortality, but the curve has shifted downward over time as health knowledge has spread, so that countries achieve lower mortality at any given income than their predecessors did [16, 19, 20]. This partial independence from income is central to the convergence question: if survival gains flow from diffusible knowledge rather than from income alone, poorer and higher-mortality countries may capture them rapidly [3, 17, 20].

The forces behind the decline differ in their dependence on national wealth. Some determinants, such as clean water infrastructure, nutrition, and the density of trained health workers, scale with income, while others, such as vaccines, oral rehydration salts, antibiotics, and bed nets, are inexpensive and portable, deployable through national programs and international support even where incomes remain low [16, 17, 18]. The downward drift of the Preston curve over the twentieth century testifies to this second channel [19, 20]. To the extent that the post-1990 decline drew on the diffusible channel it should have reached high-mortality countries quickly, whereas the income-dependent channel should have favored the already-advantaged. The role of baseline income in the growth model bears on the relative strength of the two channels, but only indirectly, since neither channel is measured here; the model can characterize the income gradient in the pace of decline, not the mechanism behind it [3, 17].

Two notions of convergence

Whether countries grow more alike over time is a question with a long pedigree in the study of economic growth, which distinguishes the catching up of laggards, beta-convergence, from the narrowing of the cross-sectional distribution, sigma-convergence, two notions that need not coincide [6, 8, 9, 21]. Beta-convergence is necessary but not sufficient for sigma-convergence, because shocks and heterogeneous trajectories can keep a distribution wide even as laggards tend, on average, to catch up [6, 21]. For child mortality the distinction is not merely technical: the period under study contained a shock capable of severing pace from outcome, the human immunodeficiency virus epidemic, which raised child mortality sharply in parts of southern and eastern Africa during the 1990s before treatment and prevention

reversed it [3, 5, 22]. The grand convergence thesis, which holds that the gap in child survival between the poorest countries and the rest could close within a generation, is a claim about outcomes, and evaluating it requires keeping the two notions apart [3].

Terms are used precisely in what follows. Convergence in pace means that a country's initial mortality level is negatively associated with its subsequent proportional rate of decline; it does not mean that countries' rates of decline became more similar to one another over time. Convergence in outcomes means that the cross-country dispersion of log mortality, an equal-country-weighted measure of relative dispersion among national rates, declined; it is not a measure of absolute mortality differences, nor of inequality among the world's children, and an unchanged dispersion does not by itself imply that country rankings were frozen [5, 6, 9].

Modelling trajectories and research questions

The multilevel growth curve represents change as a property of the unit rather than of the sample, estimating a trajectory for each country while pooling information across countries to stabilize the estimates [7, 10, 11, 23]. Time enters at the first level, nested within countries at the second; the intercept, the level at a chosen origin, and the slope, the rate of change, vary randomly across countries, with a covariance that captures whether high starting points go with fast or slow change [12, 13, 24]. Because the global decline itself is thoroughly established [1, 2], the study is framed around research questions rather than hypotheses about its existence:

RQ1. How fast did under-five mortality decline on average over 1990 to 2019, and with what uncertainty?

RQ2. How much did national trajectories vary in starting level and in pace of decline, and how is that variation distributed across and within regions?

RQ3. Did countries with higher initial mortality decline proportionally faster, and did the cross-country distribution of mortality narrow?

Methods

Data and provenance

The outcome is the under-five mortality rate, the probability that a child born in a given year dies before age five, per one thousand live births, analyzed on the natural logarithmic scale, on which the near-exponential decline is approximately linear and slopes are proportional rates of change [1, 2]. The panel comprises the 179 countries with a complete annual series from 1990 to 2019 and an observed 1990 income level, 5,370 country-year observations in all; an audit of the analysis file found no duplicate country-years, no gaps, and full internal consistency between the mortality and log-mortality columns. The country-level predictor is the logarithm of 1990 income per capita at purchasing power parity, standardized to mean zero and unit variance. Countries are assigned to the seven World Bank analytical

regions, as published at the time of this revision (September 2026); the assignment is included in the supplementary data. The source workbook names its inputs as the World Development Indicators series “Mortality rate, under-5 (per 1,000 live births)” and “GDP per capita, PPP” for 1990. Three elements of provenance are not recoverable from the supplied files: the exact indicator code and price basis of the PPP income series (current versus constant international dollars), the extraction date or vintage of both series, and the raw income values themselves, which enter the file only after log transformation and standardization. Recovering that metadata is an outstanding author action; until it is supplied, the income-adjusted analyses are reproducible from the standardized variable in the workbook but not re-derivable from the original public sources. The gap bears on interpretation asymmetrically: the core mortality findings, the average decline, the dispersion path, and the unconditional convergence quantities involve no income variable and are separable from it, whereas every income-adjusted quantity, and every statement about the income gradient, inherits this limitation until the variable’s exact identity is documented. All analyses are conditional on the supplied panel, which is provided in full as supplementary material.

A second feature of provenance is analytically consequential. The under-five mortality series disseminated through the World Bank derive from the United Nations Inter-agency Group for Child Mortality Estimation, which produces them by fitting a Bayesian penalized B-spline model to underlying survey, census, and vital registration sources of varying quality [1, 25]. The annual values are therefore estimates, smoothed across years, and thirty annual observations per country do not represent thirty independent measurements. Two consequences follow. First, within-country residuals from any smooth trend model will be strongly serially dependent, partly because of source smoothing, though genuine demographic persistence and residual misspecification can also contribute; the serial-correlation results examine the error structure directly. Second, single-year comparisons, above all those anchored at the 1990 baseline, inherit the estimation error of those years, which motivates the sensitivity analyses of the convergence parameters. Compatible joint uncertainty draws for this data vintage were not available to the authors, so formal uncertainty propagation was not attempted; the analysis is conditional on the published point estimates and this limitation is carried through the discussion.

The growth model and its estimands

The headline model is a two-level linear growth curve. At the first level, log mortality for country i in year t is a linear function of time, measured in decades since 1990, with a country-specific intercept and slope and a within-country residual, treated as conditionally independent in the headline specification. At the second level, the intercept and slope are population means plus random country deviations following a bivariate normal distribution, and standardized baseline income predicts both the intercept and, through a cross-level interaction with time, the slope [7, 10, 11]. Estimation is by restricted maximum likelihood [26]. The logarithmic scale matters for interpretation: a slope on the log scale is a proportional rate, so a country falling from one hundred to fifty has achieved the same proportional progress as one falling

from ten to five, and convergence in proportional terms is the demanding standard, absolute convergence being nearly automatic for high-burden countries [7, 9, 19].

Because several related quantities are easily conflated, the analysis names its estimands and reports them separately. Q1 is the population correlation between the random intercept and the random slope in the income-adjusted model; it describes the association between country deviations in level and in pace after the income terms, at the 1990 origin. Q2 is the same correlation from a model without income terms, the unconditional association between total starting levels and total slopes. Q3 is the descriptive Pearson correlation among the fitted country estimates (empirical Bayes predictions, which are shrunken toward the mean and are not the population parameter). Q4 is the coefficient from regressing each country's observed change in log mortality per decade on its observed baseline level, the classical beta-convergence regression, which uses no random-effects machinery at all. Sigma-convergence is assessed separately from the year-by-year cross-country standard deviation of log mortality, with a within-region and between-region decomposition. Throughout, sample standard deviations (denominator n − 1) are used, and the between/within decomposition uses the same denominator so that its parts sum to the total.

Inference

Uncertainty for fixed effects is reported as model-based 95% confidence intervals. For the central intercept-slope correlation, whose sampling distribution is not well approximated by a Wald interval, a cluster bootstrap resampling countries with replacement (500 replicates, seed 20260921) provides percentile intervals; draws enter the summary under a fixed acceptance rule, optimizer convergence and a finite correlation, and acceptance counts are reported with the results. Variance components sit on the boundary of the parameter space under the null, so their tests use likelihood-ratio statistics with mixture reference distributions [27]. In the revised analysis, heteroskedasticity-robust HC1 standard errors with a t(177) reference distribution are used for the primary country-level regression inference, and conventional standard errors are reported as a sensitivity check; where the two verdicts differ at conventional thresholds, both are shown and the difference is noted rather than resolved by the smaller p-value.

Robustness and sensitivity strategy

The headline model is deliberately parsimonious, and the analysis surrounds it with checks in four families. Functional form: the linear specification is compared by maximum likelihood with a quadratic-time model and a piecewise-linear model with knots at 2000 and 2010, all retaining country random intercepts and linear slopes; in addition, a model with country-specific random quadratic terms relaxes the assumption of linear country deviations, and its convergence diagnostics are reported. Error structure: within-country residual autocorrelation is profiled; a mixed model retaining random intercepts and slopes while adding a first-order autoregressive residual process is estimated by direct maximum likelihood, with a profile over the autoregressive parameter to assess whether the data can separate a

persistent residual process from the random effects; generalized estimating equations with an AR(1) working correlation and ordinary least squares with country-clustered standard errors are retained as complementary estimators of the fixed effects only. Geography: estimated country slopes are summarized by region; a three-level model apportions heterogeneity between regions and countries within regions, with the small number of regions (seven) treated as a limitation; a Mundlak-type decomposition separates income into regional means and within-region deviations [28], and a region-fixed-effects specification with region-by-time terms serves as a further sensitivity. Convergence parameters: the time origin is recentered, with the algebraic transformation of the covariance verified explicitly; the beta regression is re-run with alternative baselines and windows whose construction is stated exactly; and seeded perturbation simulations add noise to the mortality series under three labeled scenarios, including temporally correlated noise.

Results

The average trajectory (RQ1)

Table 1 reports descriptive statistics and Table 2 the fixed effects of the growth model. The average trajectory begins at a 1990 level corresponding to about forty-five deaths per thousand live births (the exponential of the model intercept) and declines by 30.7 percent per decade (95% CI, 28.9 to 32.4), a large and precisely estimated slope. Baseline income strongly predicts the starting level, richer countries beginning far lower. Its interaction with time is small and imprecisely estimated, + 0.022 log units per decade per standard deviation of income (95% CI, − 0.003 to + 0.047; p = .085): the point estimate implies that a country one standard deviation richer declined about 2 percentage points per decade more slowly, an effect an order of magnitude smaller than the average decline itself. The serial-correlation results and the regional decomposition below examine this interaction further; on the evidence of this model, a country’s economic starting point shaped where its trajectory began far more than how fast it fell.

Table 1
Descriptive statistics for under-five mortality (per 1,000 live births)

Statistic	1990	2019	Change
Geometric mean across countries	43.5	16.1	−63%
Median across countries	47.3	15.9	−66%
Lowest-income tercile, geometric mean	113.1	41.4	−63%
Highest-income tercile, geometric mean	15.8	5.9	−62%
Cross-country SD of log mortality	1.05	1.09	+ 4%

Note. Levels are geometric means (the exponential of the mean of log mortality), the summary consistent with the logarithmic scale of the analysis; the median is shown for comparison. Terciles split the 179 countries into thirds by standardized 1990 income. All values are computed from the supplementary data file; standard deviations are sample SDs (denominator n − 1). Mortality fell by close to two thirds in every group while the cross-country SD of log mortality did not fall.

Table 2
Fixed effects of the multilevel growth model (log under-5 mortality)

Parameter	Estimate	SE	95% CI	p
Intercept (1990 level)	3.796	0.048	[3.701, 3.890]	< .001
Time (per decade)	−0.366	0.013	[− 0.391, − 0.341]	< .001
Baseline income (z)	−0.918	0.049	[− 1.013, − 0.822]	< .001
Income × time	+ 0.022	0.013	[− 0.003, + 0.047]	.085

Note. REML estimation; 5,370 observations, 179 countries; converged, no boundary warnings. The time slope corresponds to a decline of 30.7% per decade (exp(− 0.366) − 1). SEs are model-based; standard deviations describing between-country heterogeneity appear in Table 3 *and are distinct from these standard errors.*

Figure 1 shows the 179 trajectories with the cross-country mean of log mortality overlaid; because countries are weighted equally, this mean is a summary of national trajectories, not a birth-weighted global mortality rate. The income gradient orders the levels almost completely, while the steepness of decline is far less ordered by income, the graphical counterpart of the weak interaction in Table 2. This combination, a strong income gradient in level and a weak one in pace, is consistent with accounts in which survival technologies diffused broadly across income levels [16, 17, 19], but the model measures no immunization coverage, water and sanitation, health expenditure, external assistance, schooling, or health-system capacity, so the pattern is equally compatible with other forces that operate largely independently of national income. Nor does a p-value of .085 establish that income played no role; the interval is consistent with effects in either direction of modest size.

Heterogeneity and its geography (RQ2)

Table 3 reports the variance components. The standard deviation of the random intercepts, 0.647 on the log scale, reflects the enormous range of 1990 mortality; the standard deviation of the random slopes, 0.171 per decade, means that a country one standard deviation above the mean pace declined about forty-six percent per decade while one a standard deviation below declined about seventeen percent. Both components are decisively nonzero by boundary-corrected likelihood-ratio tests (Table 3), and the heterogeneity is the substantive point: the average trajectory describes no country exactly.

Table 3
Random effects of the growth model (income-adjusted)

Component	Estimate	Inference
SD of random intercepts (1990 level)	0.647	LR = 13,157, p < .001 (mixture)
SD of random slopes (per decade)	0.171	LR = 6,215, p < .001 (mixture)
Intercept-slope correlation	−0.240	bootstrap 95% CI [− 0.365, − 0.108]
Residual SD (within country)	0.089	—

Note. Random-effect SDs describe between-country heterogeneity, not estimation uncertainty. Likelihood-ratio tests compare nested ML fits with mixture reference distributions appropriate to boundary nulls [27]. *The correlation interval is a percentile interval from a cluster bootstrap over countries (B = 500, seed 20260921; acceptance rule: optimizer convergence and a finite correlation; 500 of 500 draws accepted, 0 accepted draws carried a boundary flag; share of accepted draws positive: 0.0%). The correlation refers to deviations after the income terms, at the 1990 origin; the convergence results below report the unconditional counterpart.*

[Figure 1 near here]

Table 4 summarizes the estimated country trajectories by region. South Asia, from a very high burden, set the fastest average pace, about forty percent per decade, led by the Maldives, the fastest estimated decliner in the sample; Europe and Central Asia follows from the lowest starting levels outside North America. Sub-Saharan Africa, with much the heaviest burden, declined at about twenty-eight percent per decade, close to the sample-wide average pace but from levels several times higher. Latin America and the Caribbean shows the widest internal spread, driven by Caribbean micro-states whose small birth cohorts make the published rates volatile; the sample's only rising estimated trajectory, Dominica, is among them. The slowest estimated decliners are a mixture of such micro-states, conflict-affected countries such as the Central African Republic, and countries severely affected by the HIV epidemic.

Table 4
Estimated country trajectories by World Bank region

Region	n	Est. 1990 level, geometric mean	Est. 1990 level, arithmetic mean	Mean decline (%/decade)	SD (%/decade)
South Asia	7	103	117	−39.8	11.4
Europe & Central Asia	48	18	28	−35.1	8.6
Middle East & North Africa	16	36	45	−32.6	6.8
Sub-Saharan Africa	45	152	175	−28.1	10.9
East Asia & Pacific	29	39	59	−27.0	12.8
Latin America & Caribbean	32	36	43	−23.2	14.8
North America	2	9	9	−14.9	2.1

Note. Country-level values are empirical Bayes estimates from the income-adjusted model, per 1,000 live births; both the geometric mean (exponential of the mean estimated log level) and the arithmetic mean of the estimated levels are shown and labeled. Regions are the seven World Bank analytical regions as assigned at revision (September 2026). Percentage declines are country-specific transformed slopes averaged within region.

A three-level model, years within countries within regions, apportions this heterogeneity formally: conditional on baseline income, about one third of the remaining between-country variance in 1990 levels and about one quarter of the variance in rates of decline lies between the seven regions (full specification and estimates in the Statistical Supplement). With only seven region-level units these shares are imprecise and are offered as description, not as precisely estimated parameters. Most of the variation in pace is variation among neighbors within regions.

The regional structure also clarifies the income-pace interaction. Decomposing income into regional means and within-region deviations [28] yields a within-region coefficient on income-by-time of + 0.064 (SE 0.017, p < .001) and a between-region coefficient of − 0.029 (SE 0.019, p = .12); a specification with region and region-by-time fixed effects gives the same within-region estimate. The pooled interaction of Table 2, + 0.022, is thus a composite of a positive within-region component, poorer neighbors catching up faster in proportional terms, and a negative, imprecisely estimated between-region component. The within-region estimate is larger and better determined than the pooled one, and the offsetting structure is a further reason not to read any single mechanism off the pooled income gradient.

Convergence in pace versus convergence in outcomes (RQ3)

Table 5 assembles the convergence quantities defined in the Methods, and Figs. 2 through 4 display them. All four tell the same story about pace. The income-adjusted random-effects correlation is − 0.240 (bootstrap 95% CI, − 0.365 to − 0.108); removing the income terms leaves it essentially unchanged at − 0.243, so the association is not an artifact of conditioning; the descriptive correlation among the fitted country estimates is − 0.24; and the regression of observed change per decade on the observed 1990 level yields − 0.021 (95% CI, − 0.040 to − 0.003; p = .026 with HC1 standard errors and a t reference; p = .063 with classical standard errors), which over the full twenty-nine years is the − 0.062 total-period coefficient. Countries that entered 1990 with higher mortality tended to decline proportionally faster over the following three decades. The association is modest, and Fig. 2 shows how far it is from a rule: many high-burden countries declined slowly and many low-burden countries declined quickly.

Table 5
Convergence in pace and in outcomes: the estimands reported separately

Quantity	**Estimate**	**Uncertainty / basis**
A. Pace		
Random-effects corr., income-adjusted (Q1)	−0.240	bootstrap 95% CI [− 0.365, − 0.108]
Random-effects corr., unconditional (Q2)	−0.243	same model without income terms
Descriptive corr., EB estimates (Q3)	−0.242	Pearson, level vs. log slope (p = .001)
Observed beta regression (Q4), per decade	−0.021	95% CI [− 0.040, − 0.003]; p = .026 (HC1-t)
B. Outcomes		
SD of log mortality, 1990	1.05	sample SD, 179 countries
SD of log mortality, 1999 (peak)	1.14	
SD of log mortality, 2019	1.09	no net reduction relative to 1990
Within-region variance, 1990 → 2019	0.50 → 0.43	declined
Between-region variance, 1990 → 2019	0.60 → 0.75	increased

Note. The four pace quantities are distinct estimands (see the Estimands paragraph of the Methods) and agree here; they are reported separately rather than interchangeably. The variance decomposition uses a common n − 1 denominator so the within- and between-region parts sum to the total; it is a descriptive accounting, not a causal attribution. The beta regression is shown with HC1 standard errors and a t(177) reference, the primary inference in the revised analysis; with conventional standard errors the p-value is .063 (see Inference, Methods).

[Figure 2 near here]

The outcomes side tells a different story, and Fig. 4 gives it precisely. The cross-country standard deviation of log mortality rose from 1.05 in 1990 to a peak of 1.14 in 1999, the years in which the HIV

epidemic was raising child mortality in parts of southern and eastern Africa, an association of timing documented in the surveillance literature rather than tested here [5, 22]; it then declined for most of two decades, but at 1.09 in 2019 remained above its 1990 value. The precise statement is that dispersion first increased and then declined, with no net reduction over 1990 to 2019; there was sigma-convergence after 1999, but not enough to undo the earlier widening. Figure 3 shows the distribution shifting left with little change in width. The decomposition in the lower panel of Table 5 adds the geography: dispersion within regions fell, from 0.50 to 0.43 in variance terms, while dispersion between regional means rose from 0.60 to 0.75, its share of the total rising from 54 to 64 percent. As descriptive accounting, the absence of net global narrowing over these three decades is a between-region phenomenon; within regions, on average, distributions compressed.

[Figure 3 near here]

[Figure 4 near here]

Functional form

Table 6 relaxes the common time trend. Both the quadratic and the piecewise specifications improve the likelihood, as almost any added flexibility will with 5,370 observations, and the piecewise model is preferred by both information criteria. The piecewise fixed-effects curve implies average declines of 28.7 percent per decade in the 1990s, 34.2 percent in the 2000s, and 25.5 percent in the 2010s, changes implied by the common fitted curve rather than averages of country-specific paces; the mid-period acceleration coincides with the years in which international support for child survival expanded most rapidly, and the 2010s easing predates the coronavirus pandemic. The convergence parameters are indistinguishable across these three specifications. That comparison, however, varies only the common curve while country deviations remain linear, so it does not by itself establish robustness to country-specific curvature. A model adding country-specific random quadratic terms addresses that directly: it converges (with a boundary caution reported in the Supplement), finds real curvature heterogeneity, and reduces the residual SD from 0.089 to 0.061. In that model the correlation between the 1990 level and the instantaneous 1990 slope is + 0.07, while the correlation between the 1990 level and the whole-period average pace, the estimand this paper's convergence claim concerns, is − 0.241, matching the linear model. Local slopes and whole-period pace are different estimands, and the convergence-in-pace finding is a statement about the latter. Figure 5 illustrates this contrast at the level of individual countries, comparing the observed series with the fitted linear trajectories for the six fastest and six slowest estimated decliners: the linear fits track the fastest decliners closely, while the curvature visible among the slowest is the pattern the random-quadratic extension accommodates.

Table 6
Functional form of the time trend (ML comparison; country random intercepts and linear slopes retained)

Specification	Params	Log-lik.	AIC	BIC
Linear time	8	4,344.2	−8,672	−8,620
Quadratic time	9	4,355.1	−8,692	−8,633
Piecewise linear (knots 2000, 2010)	10	4,431.5	−8,843	−8,777
Piecewise implied decline/decade	1990s − 28.7%	2000s − 34.2%	2010s − 25.4%	
Intercept-slope correlation	−0.240	−0.240	−0.240	
Income × time (SE ≈ 0.013)	+ 0.022	+ 0.022	+ 0.022	

Note. Same 5,370 observations and likelihood definition throughout; parameter counts include fixed effects, random-effect (co)variances, and the residual variance. The implied percentage declines are transformations of the common fitted piecewise curve. The random-quadratic extension is reported in the text and, in full, in the Statistical Supplement.

[Figure 5 near here]

Serial correlation

The error structure carries a signature of how the data are made. The median within-country lag-one autocorrelation of residuals from the headline model is 0.95, and estimators that model the dependence recover an autoregressive parameter near unity, consistent with series that are themselves spline-smoothed estimates [25], though demographic persistence and residual misspecification may contribute as well. Table 7 compares four estimators. The average pace of decline moves little, between − 0.34 and − 0.37 log units per decade, and the income gradient in levels is similarly stable. The income-by-time interaction varies more: +0.022 in the mixed model and in clustered OLS (where its p-value is .045), + 0.008 with a standard error of 0.012 under the estimators that model the serial dependence. The pattern across estimators, point estimates between + 0.008 and + 0.022 with intervals that mostly include zero, supports a bounded conclusion: the income gradient in the pace of decline is small at most, and the evidence for it is fragile, without establishing its absence.

The deeper question concerns what serial correlation does to the random-effects decomposition. In a mixed model that retains random intercepts and slopes while adding an AR(1) residual process, estimated by direct maximum likelihood and validated against the independence fit at ρ = 0, the likelihood rises steeply (from 4,344 to 7,168) and the autoregressive parameter is estimated at 0.995; three dispersed starting values reach very similar log-likelihood values (equal to two decimals) and the same ρ, while the implied intercept-slope correlation varies from − 0.17 to − 0.10 across these near-optimal fits, itself an early sign of the fragility documented below. The dense likelihood profile in the Supplement shows that ρ itself is well determined, not flat: fixing ρ at 0.95 costs about 12.5 log-

likelihood units against the optimum, a decisive difference by any likelihood standard. What is fragile is the decomposition. Across the narrow interval of ρ values that fit nearly as well as the optimum (within roughly two log-likelihood units, approximately 0.985 to 0.995), the random-intercept SD falls from about 0.51 to 0.03 and the intercept-slope correlation moves with it, because a residual process this persistent can absorb much of what the independence specification attributes to permanent country levels. The estimated random-effects decomposition is therefore sensitive to the within-country residual specification, which warrants caution in interpreting the magnitude of the intercept-slope correlation; this is a finding about this dataset and specification, established by the profile diagnostics, not a general claim about growth models on thirty-year panels. The fixed effects, by contrast, are stable across residual specifications, and the convergence-in-pace conclusion rests jointly on the independence-specification model and on the observed-data beta regressions, which involve no random-effects decomposition.

Table 7
Fixed effects under alternative treatments of serial correlation

Parameter	Mixed, iid resid. (REML)	Mixed + AR(1) resid. (ML)	GEE, AR(1) working corr.	OLS, clustered SEs
Time (per decade)	−0.366 (0.013)	−0.342 (0.012)	−0.342 (0.012)	−0.366 (0.013)
Baseline income (z)	−0.918 (0.049)	−0.859 (0.045)	−0.858 (0.051)	−0.918 (0.053)
Income × time	+ 0.022 (0.013)	+ 0.008 (0.012)	+ 0.008 (0.011)	+ 0.022 (0.011)
Serial-correlation parameter	—	ρ = 0.995	ρ = 0.996	unrestricted

Note. Standard errors in parentheses. The AR(1) mixed model retains country random intercepts and slopes (direct ML; validated against the iid ML fit at ρ = 0); its random-effects components are sensitive to the value of ρ near the optimum (see text and Supplement S6) and are therefore not tabulated as settled quantities. GEE estimates population-averaged fixed effects with robust (sandwich) standard errors and a first-order autoregressive working correlation; it does not estimate the random-effects correlation. Clustered OLS allows arbitrary within-country correlation, fixed effects only. All estimators converged.

Sensitivity of the convergence parameters

Three exercises are distinguished because they answer different questions. Recentering the time origin is a reparameterization of the same model: the intercept-slope covariance at origin c equals the 1990 covariance plus c times the slope variance, and Table 8 confirms that the re-estimated correlations match this algebra exactly, − 0.240 at 1990, − 0.111 at 1995, + 0.146 at mid-period. The sign change is a property of the parameterization, not a separately estimated change in convergence over time: the model's single covariance matrix, viewed from different origins, implies that the countries still carrying high mortality by the mid-2000s were disproportionately those whose decline had been slow. The convergence-in-pace claim is anchored at 1990, the baseline year of the international goals, and should be read as origin-specific.

Changing the baseline measure and the estimation window addresses regression to the mean. Because the 1990 values are estimates, a country whose baseline happens to be overstated will show both a high starting level and a fast apparent decline, biasing the beta coefficient negative [29]. Averaging the baseline over 1990 to 1992 and measuring change over 1993 to 2019, so that no observed year enters both the baseline and the change, leaves the association intact and somewhat stronger (− 0.034 per decade, 95% CI − 0.054 to − 0.013; p = .001); a variant that starts the change window at 1992, and therefore still shares one year with the baseline average, gives − 0.028 (p = .006) and is retained in the Supplement, labeled as overlapping. Measuring change over 1995 to 2019, with the 1990 baseline entirely outside the window, strengthens the association to − 0.045 per decade (p < .001). Separating the observed years in these ways removes the direct arithmetic overlap between baseline and change, but it does not remove serially correlated measurement error, dependence introduced by the mortality-estimation procedure, or every source of regression to the mean. The strengthening in the later window is consistent with the early-window reversals working against catch-up, and is not, by itself, proof that such artifacts are absent.

Finally, seeded perturbation simulations gauge the sensitivity of the model-based correlation to added measurement error. These are labeled scenario exercises, not propagation of the source estimates' published uncertainty, which was not available in compatible form. Independent noise added to the 1990 observations, at two scenario magnitudes, leaves the correlation at − 0.240, as expected given that the model intercept is informed by a country's whole trajectory rather than the single baseline year; temporally correlated noise added to all years (AR(1), ρ = 0.9, SD 0.05) yields a mean of − 0.244 with a middle-95% range of − 0.269 to − 0.223 across replications. Within these scenarios the correlation is insensitive; the scenarios do not exhaust the ways the source estimation process could err.

Table 8
Sensitivity of the convergence parameters

Analysis	Estimate	Detail
A. Origin of the time scale (same model, reparameterized)		
Origin 1990 (headline)	−0.240	reference
Origin 1995	−0.111	algebraic prediction −0.110
Origin mid-period (2004/05)	+ 0.146	algebraic prediction + 0.146
B. Beta regression (Δ log mortality per decade on baseline; HC1-t)		
Baseline 1990; window 1990–2019	−0.021 [− 0.040, − 0.003]	p = .026 (classical p = .063)
Baseline mean 1990–92; window 1993–2019	−0.034 [− 0.054, − 0.013]	p = .001 (non-overlapping years)
Baseline 1990; window 1995–2019	−0.045 [− 0.067, − 0.024]	p < .001 (baseline outside window)
C. Perturbation scenarios (B = 100 each, seed 20260921; accepted 300 of 300)		
iid noise on 1990 obs., SD 0.05	−0.240	across-replication SD 0.001
iid noise on 1990 obs., SD 0.10	−0.241	across-replication SD 0.003
AR(1) noise, all years (ρ = .9, SD 0.05)	−0.244	middle 95% [− 0.269, − 0.223]

Note. Panel A: the covariance at origin c (decades after 1990) equals cov(1990) + c × var(slope); estimated and predicted values coincide, so recentering is not an independent replication. Panel B: outcome and baseline constructions as stated; HC1 standard errors with a t(177) reference (primary in the revised analysis); with conventional standard errors the three p-values are .063, .008, and < .001; an additional variant sharing one year between baseline and window appears in Supplement S9. Panel C: statistic is the REML intercept-slope correlation; acceptance rule as in the bootstrap, with acceptance counts in Supplement S9; Monte Carlo SE of each reported mean is at most 0.0012; noise magnitudes are assumptions, labeled as scenarios.

Discussion

The analysis returns a divided verdict on one of the period's defining health achievements. Child mortality fell by about thirty-one percent per decade on average, and it fell proportionally faster where it had been higher: every estimand of pace convergence examined, conditional and unconditional, model-based and observed-data, points the same way, with the model-based correlation at − 0.24 (95% CI, −

0.37 to − 0.11). Yet the cross-country dispersion of log mortality shows no net reduction over 1990 to 2019, having widened through the 1990s and then narrowed without returning to its starting value. The two notions of convergence, often blurred in discussion of global health progress, genuinely came apart in this period: catch-up in pace operated, and the distribution of outcomes did not compress on net.

The geography of the two findings is instructive. Within regions, dispersion fell and poorer neighbors caught up faster, the within-region income-pace gradient being positive and well determined; between regions, dispersion rose. In descriptive accounting terms, the absence of net global narrowing over these three decades is a between-region phenomenon, dominated by the gap between sub-Saharan Africa, whose burden remained several times the global level while its pace merely matched the average, and the rest of the distribution. Whether the between-region widening reflects the HIV epidemic, conflict, health-system differences, or other forces is not tested here; the epidemic's timing coincides with the 1990s widening, but attribution would require event- or cause-specific analysis beyond this design.

What drove the decline is a question this design can bound but not settle. The weak and statistically fragile income gradient in rates of decline is consistent with accounts in which survival gains flowed through channels not requiring national wealth, the diffusion of inexpensive interventions being the leading candidate in the literature [16, 17, 20], but the model measures none of those channels, and the offsetting within- and between-region components of the gradient warn against any single-mechanism reading. Only baseline income is modeled, so the analysis is silent on the role of income growth over the period. What the estimates do license is narrower: the pace of proportional decline was at most weakly ordered by economic starting point, and whatever drove the decline operated broadly across the income distribution without closing the cross-country distribution of outcomes. Identifying which interventions or investments would close the remaining gaps is beyond this design; the results locate the gaps, between regions and in the highest-burden settings, without prescribing the remedy.

The robustness program sharpens the reading and sets its limits. The average pace and the income gradient in levels are stable across functional forms and error structures. The convergence-in-pace estimand survives country-specific curvature once it is defined, correctly, as whole-period pace rather than an instantaneous slope. The serial-correlation analysis, by contrast, marks the limit: for this dataset and specification, the estimated random-effects decomposition is sensitive to the within-country residual specification, the intercept variance and the correlation shifting substantially across residual models that fit nearly as well as the optimum, so the magnitude of the model-based correlation warrants caution, and the pace-convergence conclusion is anchored jointly in that model and in observed-data regressions that need no such decomposition.

Because the analysis is conducted in relative terms, its inequality statements need careful language. Equal proportional declines leave ratios between countries unchanged while absolute differences shrink; with mortality falling nearly everywhere, absolute gaps narrowed almost mechanically, and the persistent quantity is relative dispersion, the ratio structure of national rates. The SD of log mortality is an equal-country-weighted measure of that relative dispersion; it says nothing about inequality among the world's

children, weights Iceland and India equally, and its constancy does not imply frozen rankings, since rank mobility is a separate question not examined here. On the demanding relative standard, high-burden countries did cut mortality by larger fractions, and yet the relative structure of the distribution outlasted three decades of decline; on the absolute standard, gaps closed broadly. Both statements are true, and policy discussion is better served by keeping them apart than by a single word like convergence [5, 6, 9].

Limitations and future research

The most consequential limitation concerns the data. The under-five mortality series analyzed here are estimates produced by a Bayesian spline model pooling sources of varying quality [1, 25], and the analysis is conditional on the published point estimates: compatible joint uncertainty draws for this vintage were not available, marginal intervals alone would not identify the joint temporal uncertainty, and no uncertainty propagation is claimed. The smoothness of the series limits the effective information per country well below thirty independent observations, which is precisely what the serial-correlation analysis shows, and estimation uncertainty is generally larger where vital registration is weak, which includes many high-burden countries, though this analysis did not verify source-level uncertainty country by country. Re-estimating the growth model across draws from a compatible release of the source uncertainty distributions is the natural next step [2, 30].

Four further limitations bound the conclusions. First, although the fixed trend was relaxed and a random-quadratic extension was estimated, richer country-specific dynamics, including reversal episodes, are summarized rather than modeled; a specification with structured breaks would suit the HIV-era reversals better than curvature alone [13, 24]. Second, the design is ecological and the single country-level predictor leaves the mechanism question open; time-varying measures of intervention coverage, expenditure, assistance, education, and governance, with the estimated country slopes as a second-stage outcome, would address it directly [16, 17, 18]; moreover, the income variable itself carries the provenance gap described under Data and provenance, so all income-related interpretations remain provisional until its exact identity is documented. Third, the regional analyses rest on seven units, so region-level variance components are imprecise, and the smallest states contribute volatile series that are prominent among the slowest estimated decliners; the fixed-effects sensitivity analysis and the micro-state caveats bound but do not eliminate these concerns. Fourth, the panel ends in 2019, and the pandemic years that follow will have tested the regularities documented here; extending the model through the 2020s, with the pandemic as a common shock with country-specific loadings, is an obvious continuation.

The interpretation of the random effects also warrants care. The intercept-slope correlation is a population-level summary of how starting points and paces co-vary, not a deterministic law; the scatter in Fig. 2 shows many high-burden countries declining slowly and many low-burden countries declining quickly. The countries that declined slowly despite high burdens are where the failure of outcome convergence is concentrated, and identifying their common features is the natural next step for research aimed at accelerating progress [5, 18, 22].

Conclusions

This study modeled under-five mortality trajectories for 179 countries over 1990 to 2019 with a multilevel growth curve on the logarithmic scale, quantifying the average pace of decline, the heterogeneity of national trajectories, and two distinct notions of convergence, and subjecting each headline quantity to reproduction, alternative specifications, and sensitivity analysis.

Three findings survive that scrutiny. Mortality fell by about thirty-one percent per decade on average, an estimate stable across every specification examined. National trajectories varied widely around that average, with about a quarter to a third of the heterogeneity lying between regions and the rest among neighbors: South Asia fastest, sub-Saharan Africa carrying the heaviest burden at an average pace, Caribbean micro-states supplying most of the slowest and least reliably measured paths. And the period's progress was convergent in pace but not, on net, in outcomes: higher-mortality countries declined proportionally faster, by every estimand examined, yet the cross-country dispersion of log mortality was higher in 2019 than in 1990, its post-1999 narrowing insufficient to undo the widening of the 1990s, and the missing compression is located between regions rather than within them. Baseline income shaped where trajectories began far more than how fast they fell, a pattern consistent with, though not by itself demonstrating, the diffusion of broadly accessible health interventions.

For the aspiration of a grand convergence in child survival, the results carry a double message: the observed pattern was consistent with catch-up in proportional decline, and that catch-up was not sufficient, over three decades, to narrow the distribution of outcomes across countries. Progress on the average and progress on the distribution are different achievements, and monitoring frameworks that track only the first will not register the second.

Abbreviations

AR(1)
first-order autoregressive
CI
confidence interval
EB
empirical Bayes
GEE
generalized estimating equations
HC1
heteroskedasticity-consistent variance estimator
HIV
human immunodeficiency virus
LR
likelihood ratio

ML
maximum likelihood
OLS
ordinary least squares
PPP
purchasing power parity
REML
restricted maximum likelihood
SD
standard deviation
SE
standard error
U5MR
under-five mortality rate
UN IGME
United Nations Inter-agency Group for Child Mortality Estimation
WDI
World Development Indicators.

Declarations

Ethics approval and consent to participate

Not applicable. The study uses only publicly available, aggregate country-level data; no human participants, individual-level human data, or human tissue were involved.

Consent for publication

Not applicable.

Competing interests

The authors declare that they have no competing interests.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Author Contribution

AG conceived and supervised the study, contributed to the interpretation of the results, and critically revised the manuscript. EA designed the analysis, performed the statistical modelling, and drafted the manuscript. GK contributed to the literature review, the interpretation of the results, and the revision of the manuscript. All authors read and approved the final manuscript.

Acknowledgements

Not applicable.

Data Availability

All data analysed during this study are included in this published article and its additional files. The study uses only publicly available, aggregate country-level data: the under-five mortality rate per 1,000 live births for 179 countries (1990–2019) from the World Bank World Development Indicators (disseminated from UN IGME estimates) and 1990 GDP per capita at purchasing power parity in standardized logarithmic form. Additional file 2 contains the complete analysis-ready dataset with a codebook, and Additional file 3 contains an executable analysis package that reproduces all reported results from the supplied data. The exact indicator code, price basis, and vintage of the income series are not recorded in the source file and are disclosed in the Methods as a limitation.

References

1. You D, Hug L, Ejdemyr S, Idele P, Hogan D, Mathers C, et al. Global, regional, and national levels and trends in under-5 mortality between 1990 and 2015. Lancet. 2015;386(10010):2275–86.
2. Wang H, Liddell CA, Coates MM, Mooney MD, Levitz CE, Schumacher AE, et al. Global, regional, and national levels of neonatal, infant, and under-5 mortality during 1990–2015. Lancet. 2016;388(10053):1725–74.
3. Jamison DT, Summers LH, Alleyne G, Arrow KJ, Berkley S, Binagwaho A, et al. Global health 2035: a world converging within a generation. Lancet. 2013;382(9908):1898–955.
4. Rajaratnam JK, Marcus JR, Flaxman AD, Wang H, Levin-Rector A, Dwyer L, et al. Neonatal, postneonatal, childhood, and under-5 mortality for 187 countries, 1970–2010. Lancet. 2010;375(9730):1988–2008.
5. Moser KA, Leon DA, Gwatkin DR. How does progress towards the child mortality millennium development goal affect inequalities between the poorest and least poor? BMJ. 2005;331(7526):1180–2.
6. Quah DT. Twin peaks: growth and convergence in models of distribution dynamics. Econ J. 1996;106(437):1045–55.

7. Singer JD, Willett JB. Applied longitudinal data analysis: modeling change and event occurrence. New York: Oxford University Press; 2003.
8. Barro RJ, Sala-i-Martin X. Convergence J Polit Econ. 1992;100(2):223–51.
9. Sala-i-Martin X. The classical approach to convergence analysis. Econ J. 1996;106(437):1019–36.
10. Laird NM, Ware JH. Random-effects models for longitudinal data. Biometrics. 1982;38(4):963–74.
11. Raudenbush SW, Bryk AS. Hierarchical linear models: applications and data analysis methods. 2nd ed. Thousand Oaks: Sage; 2002.
12. Bryk AS, Raudenbush SW. Application of hierarchical linear models to assessing change. Psychol Bull. 1987;101(1):147–58.
13. Curran PJ, Obeidat K, Losardo D. Twelve frequently asked questions about growth curve modeling. J Cogn Dev. 2010;11(2):121–36.
14. Snijders TAB, Bosker RJ. Multilevel analysis: an introduction to basic and advanced multilevel modeling. 2nd ed. London: Sage; 2012.
15. Bourguignon F, Morrisson C. Inequality among world citizens: 1820–1992. Am Econ Rev. 2002;92(4):727–44.
16. Cutler D, Deaton A, Lleras-Muney A. The determinants of mortality. J Econ Perspect. 2006;20(3):97–120.
17. Soares RR. On the determinants of mortality reductions in the developing world. Popul Dev Rev. 2007;33(2):247–87.
18. McGuire JW. Basic health care provision and under-5 mortality: a cross-national study of developing countries. World Dev. 2006;34(3):405–25.
19. Preston SH. The changing relation between mortality and level of economic development. Popul Stud (Camb). 1975;29(2):231–48.
20. Deaton A. The great escape: health, wealth, and the origins of inequality. Princeton: Princeton University Press; 2013.
21. Baumol WJ. Productivity growth, convergence, and welfare: what the long-run data show. Am Econ Rev. 1986;76(5):1072–85.
22. Murray CJL, Lopez AD. Global mortality, disability, and the contribution of risk factors: Global Burden of Disease Study. Lancet. 1997;349(9063):1436–42.
23. Diggle PJ, Heagerty P, Liang KY, Zeger SL. Analysis of longitudinal data. 2nd ed. Oxford: Oxford University Press; 2002.
24. Bollen KA, Curran PJ. Latent curve models: a structural equation perspective. Hoboken: Wiley; 2006.
25. Alkema L, New JR. Global estimation of child mortality using a Bayesian B-spline bias-reduction model. Ann Appl Stat. 2014;8(4):2122–49.
26. Verbeke G, Molenberghs G. Linear mixed models for longitudinal data. New York: Springer; 2000.
27. Stram DO, Lee JW. Variance components testing in the longitudinal mixed effects model. Biometrics. 1994;50(4):1171–7.

28. Mundlak Y. On the pooling of time series and cross section data. Econometrica. 1978;46(1):69–85.
29. Barnett AG, van der Pols JC, Dobson AJ. Regression to the mean: what it is and how to deal with it. Int J Epidemiol. 2005;34(1):215–20.
30. Goldstein H. Multilevel statistical models. 4th ed. Chichester: Wiley; 2011.