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.

Global inequalities in recent health system capacity and projected growth in landscape fire smoke mortality

Yanrui Li

likevrui@gmail.com

[1]University of Toronto Schools

Research Article

Keywords: Landscape-fire smoke, Fine particulate matter, Health-system capacity, Health inequality, Climate adaptation, Mortality projections, Global health

Posted Date: November 11th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11273253/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

Landscape-fire smoke mortality is projected to rise under scenarios holding vulnerability and adaptation fixed. We examined whether recent health-system capacity is lowest where projected mortality growth is greatest.

Methods

We constructed a health-system capacity index for 191 countries from five World Bank indicators and compared it with the WHO universal health coverage service coverage index. We mapped 165 countries to 17 populated projection regions. Published mortality ratios incorporate smoke exposure, demographic change and baseline mortality; inference was region-level.

Results

Lower-capacity regions faced greater projected proportional mortality growth (Spearman $\rho = -0.85$, $n = 17$; permutation $p < 0.001$). The association persisted across index construction, regional weighting, projection-uncertainty checks and country resolution, and when the WHO benchmark replaced the index. Capacity correlated with the WHO benchmark at $\rho = 0.85$. Capacity deficits concentrated among poorer populations (concentration index − 0.30). Some 0.88 billion people, one in nine of those analysed, lived where the lowest capacity tertile coincided with the highest projected-growth tertile. Adjusted estimates depended on how development and demographic structure were represented.

Conclusions

Recent health-system capacity and projected smoke-mortality growth were inversely co-located, with misalignment concentrated among poorer populations. The analysis is descriptive: fixed-adaptation projections cannot estimate capacity effects, and the association could not be cleanly separated from development and demographic structure.

Trial registration

Not applicable.

Background

Landscape-fire smoke is an escalating public-health hazard[1–3]. Exposure to fire-derived fine particulate matter (PM2.5) is consistently associated with respiratory morbidity and mortality, with more mixed evidence for cardiovascular outcomes[2–5], and landscape-fire smoke was estimated to cause an average of roughly 339,000 premature deaths annually during 1997–2006[6]. Fire activity and burned area are being reshaped by climate change[7, 8], and the burden of fire smoke is substantial, and projected to grow, even in high-income settings such as the United States[9, 10], Australia[11, 12] and Canada[13]. Quantifying the mortality attributable to ambient

PM2.5 relies on exposure–response functions[14, 15] and on global exposure and burden estimates[16, 17], and Earth-system projections indicate that fire-related PM2.5 mortality could rise roughly six-fold by the end of the century, with the steepest increases in Africa and South Asia[18].

Those projections quantify where the smoke burden will grow and how warming reshapes its distribution[18]; they say nothing about whether the regions concerned are equipped to respond.

Yet the mortality *impact* of a given exposure is not fixed. It is shaped by the capacity of health systems to monitor air quality, warn populations, and treat respiratory and cardiovascular complications[4, 19–21], and by the capacity of vital-registration systems to *record* the resulting deaths accurately[22, 23]. Cause-of-death quality varies systematically within as well as between countries: in Brazil, the share of deaths assigned to ill-defined causes (deaths for which the death record does not identify a specific underlying cause) is strongly patterned by municipal development and health infrastructure[24–26]. If both the future burden and the ability to detect it are unevenly distributed, then the alignment between where smoke mortality is projected to grow and where health systems are currently weakest is itself a policy-relevant quantity that the projections alone cannot supply. The direction of that alignment may seem foreseeable (poorer, faster-growing regions tend to have weaker health systems), but foreseeable is not measured: adaptation finance and preparedness planning allocate against quantified gaps. Vulnerability and climate-and-health readiness frameworks already combine hazard, sensitivity and adaptive capacity[27, 28], and measures of health inequality are well established[29]. What has not been done is narrower, and is what this study contributes: linking one regionally resolved projection of landscape-fire PM2.5 mortality growth to current formal health-service supply, with inference matched to that projection's geographical resolution, and bounding how far the resulting co-location is attributable to development and demography rather than to capacity itself.

A second, conceptually distinct question is whether *present-day* capacity modifies how present-day fire activity relates to mortality and its recording. That question concerns realized fire activity and contemporaneous deaths rather than projected future burden, so it can neither confirm nor refute the gap finding. It is addressed by a separate, hypothesis-generating within-country analysis reported in full in Additional file 1 and summarised only briefly here.

It is worth separating, at the outset, what this design can deliver from what it cannot, because the two are easily conflated. What the study contributes is descriptive and quantitative: how tightly present capacity and projected smoke-burden growth are co-located, the resolution at which that co-location is statistically defensible, and how much of it can be attributed to measured development and demographic structure rather than to capacity itself. What it does not contribute is any estimate of an effect of capacity on future smoke mortality — and not merely because confounding is hard to rule out at 17 regions. The projected outcome is produced by a framework that holds vulnerability and adaptation fixed, so a capacity effect is not weakly identified here; it is absent from the outcome by construction. No result below should be read as an estimate of how much mortality a given investment in health-system capacity would avert. Where terms such as "capacity deficit" and "response capacity" appear, they describe where a region's health-system resources currently stand relative to other regions, not what those resources would do to future burden.

Two questions follow. First, are regions with lower recent health-system capacity projected to face greater relative growth in landscape-fire smoke mortality? Second, how sensitive is that relationship to national wealth and to alternative definitions of health-system capacity? We address both, reporting significant and null results alike, with inference matched to the coarse regional resolution of the underlying projections. These two questions constitute

the primary analysis of this paper. The third and separate question noted above is taken up only as an exploratory within-country analysis in Additional file 1; it bears on mechanism, and cannot confirm or refute the regional finding either way.

Methods

Health-system capacity index

We compiled seven World Bank World Development Indicators[30] for all countries, taking each country's most recent non-null value in 2015–2024 (Table 1). The 2015–2024 window is the most recent decade of reporting: taking the latest value within a fixed window maximises country coverage under differing national reporting lags, while excluding older values that no longer describe recent capacity. The Primary index uses five direct health-system indicators: hospital beds per 1,000, physicians per 1,000, nurses and midwives per 1,000, health expenditure per capita (USD), and health expenditure as a share of gross domestic product (GDP). These five were selected because they are the World Development Indicators' direct measures of the resources an acute smoke-health response draws on — clinical workforce (physicians, nurses and midwives), inpatient infrastructure (hospital beds), and health financing (spending level and spending effort) — and because they are reported for nearly all countries in recent years (Additional file 1: Table S7). The Secondary (development) index adds GDP per capita (purchasing power parity, PPP) and internet users (%) as a sensitivity analysis. Results are reported for the Primary index unless stated otherwise. Right-skewed indicators (health expenditure per capita, and GDP per capita in the Secondary index) were $\log_{10}$-transformed before standardisation. Observed values were converted to z-scores using each indicator's available-country mean and population standard deviation. Missing z-scores were filled with income-group medians and, if still missing, the overall column median after that step. We took the leading eigenvector of the post-imputation indicator correlation matrix as weights and applied them directly to the imputed z-scores, without re-centring or re-scaling the columns. The weighted sum was oriented to correlate positively with the imputed physician-density indicator and then standardised to mean zero and unit population standard deviation. This is a score derived from the first principal component (PC1) direction; because imputation can alter column variances, it differs from conventional correlation-PCA scoring. This component direction was chosen because the five indicators are strongly intercorrelated and a single data-driven axis avoids assigning arbitrary weights; an equal-weight composite gives a nearly identical gap correlation (Additional file 1: Table S1). In practical terms, the index summarises whether a country has relatively more hospital beds, health workers and health spending than other countries; it is a comparative measure, not a direct measure of landscape-fire smoke preparedness. The income-group median rule uses only group membership, a strong routinely available predictor of these indicators, and adds no country-specific information; imputation was required for 17.8% of countries for hospital beds and for ≤ 2.1% for each remaining indicator (Additional file 1: Table S7), and the gap correlation is unchanged under complete-case analysis without imputation (Additional file 1: Table S1). The index measures clinical workforce, hospital beds, and health financing; it does not directly measure smoke forecasting, air-quality monitoring, public-warning coverage, availability of clean-air shelters, access to filtration, emergency preparedness, or the quality of death certification, all of which are plausibly relevant to smoke-mortality outcomes. Capacity scores for all 191 countries, with income group, percentile, projection region and imputation count, are listed in Additional file 1: Table S17 and released as a machine-readable file in the archive, so the index can be reused or audited without rebuilding it. Of the 217 World Bank reporting economies, we retained the 191 with no more than two of the seven compiled indicators missing, so that every country's score rests on at least three of the five observed index indicators; the 26 excluded are mostly small territories and dependencies. The index covered 191

countries; the leading eigenvalue accounted for 65.3% of the total eigenvalue sum of the post-imputation indicator correlation matrix. Of these countries, 165 could be mapped to the 17 represented mortality-projection regions and form the gap-analysis sample. As an external check that the composite tracks an independently constructed measure of health-system performance, we compared it with the World Health Organization's universal health coverage (UHC) service coverage index (Sustainable Development Goal indicator 3.8.1)[31], taking each country's most recent value in 2015–2023. Table 1 gives indicator transformations, median data years and score weights; the 26 index countries without a projection-region match are listed in Additional file 1: Table S2. The index correlates strongly with the WHO UHC service coverage index across the 191 index countries (Spearman ρ = 0.85; Additional file 1: Table S1b), supporting its interpretation as a measure of formal health-service supply. Three further construct checks are reported: the UHC index substituted for the composite throughout the gap analysis (Additional file 1: Table S4c); the GDP adjustment repeated with a six-indicator latent development factor and with added demographic covariates (Additional file 1: Table S10d); and, because the index is a recent snapshot juxtaposed with end-of-century projections, per-country 1990–2024 indicator trends extrapolated to 2050 and 2100 with the regional correlation re-estimated at each horizon — naive trend continuation, not a forecast (Additional file 1: Table S14 and Figs. S7 and S8).

Table 1
Capacity-index construction (Primary health-only index, 191 countries).

Indicator	Transformation	Countries observed (imputed)	Data years, median (range)	Median (IQR) across countries	Score weight
Hospital beds per 1000	—	157 (34)	2022 (2015–2023)	2.42 (1.11–3.95)	0.40
Physicians per 1000	—	187 (4)	2022 (2015–2023)	1.66 (0.40–3.40)	0.48
Nurses and midwives per 1000	—	190 (1)	2022 (2016–2023)	3.87 (1.64–6.54)	0.50
Health expenditure, % GDP	—	191 (0)	2023 (2021–2024)	6.53 (4.35–8.83)	0.31
Health expenditure per capita (USD)	$\log_{10}$	191 (0)	2023 (2020–2024)	469 (94–1553)	0.51

Each country contributes its most recent non-null value in 2015–2024; "imputed" counts countries whose missing indicator z-scores were filled using their World Bank income-group median, with the overall column median after that step as a fallback. Medians and interquartile ranges are across the 191 index countries on the raw (untransformed) scale, and show the wide between-country spread the index summarises — hospital beds range more than threefold and health spending per capita more than sixteen-fold across the interquartile range alone. All five weights are positive and of similar size, so the weighted index represents a broad capacity gradient rather than a single dominant indicator. The leading eigenvalue of the post-imputation indicator correlation matrix accounts for 65.3% of its total eigenvalue sum (eigenvalues 3.26, 0.82, 0.49, 0.28, 0.15). The Secondary (development-inclusive) index additionally incorporates GDP per capita and internet users as a sensitivity analysis.

Projected smoke-mortality growth and the gap analysis

Region-level ratios of end-of-century to historical fire-PM2.5 mortality were read directly from the Zhao et al. Source Data[18] (panel 4a of the source article; SSP2-4.5 primary, SSP5-8.5 sensitivity). SSP2-4.5 pairs a "middle-of-the-road" socioeconomic development pathway with intermediate emissions (4.5 W m^{-2} end-of-century radiative forcing); SSP5-8.5 pairs fossil-fuel-intensive development with very high emissions (8.5 W m^{-2}) and serves here as an upper-bound scenario. The Zhao mortality ratio is defined as the projected mean annual landscape-fire PM2.5-attributable premature deaths in 2095–2099 divided by mean annual deaths in the historical period 2010–2014, incorporating projected climate changes, unmitigated baseline population growth, population ageing, and baseline mortality rates, but assuming no future changes in community vulnerability or adaptation (see panel 4a of Zhao et al.[18]). For example, a ratio of 6 means that projected annual fire-PM2.5-attributable deaths in 2095–2099 are six times the historical annual number in 2010–2014; it does not mean that smoke concentrations alone increase six-fold. We analyse this fold-change scale because proportional growth captures policy-relevant stress on existing health infrastructure and because it is the outcome Zhao et al. publish at regional resolution; a ratio-scale outcome necessarily amplifies regions with small historical baselines, a property addressed in the Limitations. All global quantitative results in this paper are defined within the Zhao et al. (2025) projection framework, and are conditional on it; no attempt is made here to compare or reconcile alternative published landscape-fire mortality projections. Two properties of this outcome govern interpretation throughout. First, because vulnerability and adaptation are held fixed, every channel through which health-system capacity could alter future smoke mortality is excluded from the projection by construction; associations with this outcome are therefore descriptive co-locations of present capacity with projected burden growth, not estimates of any effect of capacity on that growth. Second, Zhao et al.'s published sensitivity experiments attribute most of the projected global rise to demographic change rather than to climate-driven emission change (Additional file 1: Table S12); cross-region variation in the ratio therefore partly reflects where population growth and ageing concentrate, and a strong inverse correlation with a development-patterned index is expected partly on demographic grounds. The Source Data do not include a regional decomposition that would allow correlating capacity against an exposure-only component. Additional file 1: Fig. S6A summarises these relationships as a directed graph. These region-level ratios were mapped to countries via a frozen country–region crosswalk (Additional file 1: Table S9; 18 Zhao regions, of which 17 are populated by the 165 mapped countries in the baseline sample — the crosswalk does not assign any country to Equatorial Africa, a limitation examined in the selected-country sensitivity analysis below). Because these ratios take only 17 distinct (region-level) values across the sample, a naïve country-level correlation test is pseudo-replicated; we therefore report the region-level Spearman correlation (n = 17 regions) with a regional permutation test. Rank-based (Spearman) correlation is used because the regional mortality ratios are strongly right-skewed (spanning 1.03 to 12.41), so that with only 17 regions a single extreme value could dominate a product-moment correlation, and because the claim under test is ordinal — whether low capacity and high projected growth co-locate. In the permutation test, the 17 regional mortality ratios are permuted across regions and the regional Spearman statistic recomputed (100,000 draws; two-sided), so the permutation statistic is identical to the reported statistic; the permutation p-value therefore answers a simple question — how often would randomly reshuffling the 17 ratios across regions produce an alignment as strong as the one observed? All permutation p-values use the add-one correction $p = (b + 1)/(B + 1)$, where b is the number of permuted statistics at least as extreme as the observed statistic and B the number of draws; they are Monte Carlo estimates from the stated number of draws, not large-sample analytic approximations. B = 100,000 (resolution $1/(B + 1) = 10^{-5}$) for the primary, scenario, index-construction and Freedman–Lane tests, and B = 20,000 (resolution 5×10^{-5}) for the 17 leave-one-region-out re-estimates, the index-robustness variants (Additional file 1: Table S18) and the capacity-trend horizons (Additional file 1: Table S14). Sensitivity analyses reassign equatorial/central-African countries to Zhao's Equatorial Africa region (giving 18 regions), repeat under SSP5-8.5, and re-estimate the regional correlation with unweighted (rather

than population-weighted) regional capacity means. No regional uncertainty intervals are published for the ratios in panel 4a of Zhao et al.[18]. We therefore perturbed each regional ratio with independent multiplicative lognormal noise calibrated to Zhao et al.'s published global 95% confidence interval (log-scale standard deviation, SD, 0.31; harsher variant 0.37 treating the two periods' uncertainties as independent), recomputing the regional correlation over 10,000 draws (Additional file 1: Table S11). The global interval includes exposure–response and model components that are shared across regions; shared multiplicative components would tend to cancel within each region's future/historical ratio and, if common across regions, would not alter regional ranks. Treating the interval as independent regional noise therefore overstates instability, making this check conservative. Even so, region-specific uncertainty and cross-region covariance are not published, so these perturbations cannot capture all shared structural errors. Finally, because the regional ratios constrain inference to 17 observations, we added a country-resolution robustness analysis using the country-level premature-death values published in the same article's Fig. 2 Source Data (historical, SSP2-4.5 and SSP5-8.5 panels). These are map-grade values apportioned from gridded model output rather than the article's regional-ratio product (for four of the five single-country regions the implied country ratios differ from the regional ratios published in panel 4a of Zhao et al.[18] by 1–12%; China is not covered by this check), so they support a robustness check rather than a replacement of the regional design. Each country's own projected/historical ratio is related to its own capacity score, which removes the shared-regional-outcome constraint; inference uses a region-cluster bootstrap (10,000 draws, resampling whole regions with all their countries) to respect within-region dependence, and results are reported without baseline filtering and with historical baselines of at least 10 and 100 deaths per year, because ratios with tiny denominators are unstable (Additional file 1: Table S4b).

Equity and sensitivity to national wealth

To size the misalignment in people rather than in correlation, countries were cross-classified on two tertiles: their own capacity score, and the projected mortality ratio of the region they map to (so every country in a region carries that region's ratio). Both sets of cuts are taken over the 165 mapped countries, tertiles being country-count based rather than population-weighted; the double-jeopardy cell is the lowest capacity tertile crossed with the highest projected-growth tertile, and the denominator for "one in nine" is the summed World Bank population of those 165 countries (8.05 billion), using the same population series as the population-weighted regional means. Median-split and quartile definitions are reported alongside so that the dependence of the headline count on where the cuts fall is visible (Additional file 1: Table S15), and the country list is in the released index dataset. We computed a population-weighted concentration index[29] of the capacity deficit (distance below the 90th-percentile "benchmark" capacity), ranked by GDP per capita. The concentration index ranges from − 1 to + 1: zero indicates a deficit spread evenly across the income distribution, and negative values indicate concentration among poorer populations. The 90th percentile was selected as an aspirational but empirically attained benchmark (the capacity level already achieved by the top decile of countries) because anchoring the deficit to the distribution's maximum would make it sensitive to a single outlying country; alternative benchmarks are reported in Additional file 1: Table S8. The index is read as a distributional summary of where the deficit sits, not as evidence of an association independent of development. To test whether the gap relationship is merely a wealth proxy, we compared the primary health-only index to a secondary development-inclusive capacity index, both evaluated with the regional permutation test. We also computed a GDP-adjusted partial Spearman correlation at the regional level: population-weighted regional primary capacity versus the regional mortality ratio, partialling out population-weighted regional log GDP per capita. Because permuting the raw mortality ratios would also destroy the mortality–GDP relationship (testing marginal rather than conditional independence), inference for the partial correlation uses Freedman–

Lane[32] residual permutation: the ranked mortality ratios are residualized on ranked regional GDP, the residuals are permuted (100,000 draws), the outcome is reconstructed as fitted value plus permuted residual, and the partial statistic is recomputed. Because a null adjusted estimate is only informative if the design could have detected a residual association, we also characterised the detection limit of this test: holding the observed capacity and GDP values fixed, we injected residual partial associations of known size into the GDP-residualised outcome and recorded how often each exceeded the Freedman–Lane critical value (2,000 simulated datasets per size; see Additional file 1: Table S10c). We interpret this adjusted estimate as a test of separability, not of whether capacity "matters". Log GDP per capita is a single, noisy proxy for a broader latent development construct, so a null residual association cannot establish that development rather than capacity drives the gap; and with capacity and GDP correlated at $\rho = 0.94$ across regions, the partial correlation is estimated from very little residual variation. To make this non-identification explicit, we additionally computed a partial-identification bound: treating log GDP as an imperfect stand-in for latent development, with rank-scale reliability λ denoting how faithfully it tracks the underlying construct ($\lambda = 1$: perfectly), we evaluated the full range of development-adjusted partial correlations consistent with the observed rank correlations and a valid joint correlation matrix, for λ from 1.00 down to 0.70, a floor below which GDP would be an implausibly weak correlate of latent development (Additional file 1: Table S13). We obtained the 95% confidence interval as the 2.5th–97.5th percentiles of 10,000 bootstrap resamples of the 17 regional observations drawn with replacement, recomputing the partial Spearman statistic on each resample; resamples with fewer than three distinct mortality-ratio values were discarded, and the population-weighted regional aggregates were held fixed within each resample (weights are not recomputed).

Separate exploratory analysis: within-country fire activity and cause-specific mortality

A separate exploratory analysis asks a different question from the primary global test, with a different estimand: whether present-day municipal hospital-bed capacity modifies the association between present-day local fire activity and ill-defined-cause mortality. It uses a Brazilian municipality-month panel for 2015–2019 built from DATASUS/SIM death records[33], INPE Programa Queimadas active-fire detections[34] and CNES hospital-bed counts[35], with municipal populations from IBGE[36]. The primary specification is an interaction between log fire activity and standardised beds per 1,000 on the ill-defined-cause mortality rate, with municipality-by-calendar-month and year-month fixed effects and municipality-clustered standard errors. Because satellite fire detections index where burning occurs rather than what is breathed, the primary specification was re-estimated with modelled landscape-fire PM2.5 from the Hu et al. (2025) product[37, 38], area-weighted onto the IBGE municipal mesh[39], substituted for fire counts, and the same design was attempted as replications in Mexico (INEGI microdata[40] with NASA FIRMS fire[41]) and Chile (DEIS microdata[42] with CONAF regional fire[43]). The replication capacity measures used the INEGI DENUE establishment directory[44] and the DEIS REM bed-day series[45], respectively. The exposure substitution, the joint fire-plus-smoke model and the widened set of cause-specific outcomes were added after the fire-count results were seen and are reported as post hoc. This analysis is exploratory throughout and is unadjusted for multiplicity; the full specification, estimates, falsification tests and replication results are in Additional file 1 (Supplementary Methods; Tables S5–S6 and S16; Supplementary Note; Fig. S9).

Statistical analysis plan

The primary analysis is the region-level association between the health-only capacity index and the SSP2-4.5 mortality ratio, using population-weighted regional means, the region-level Spearman correlation, and regional

permutation inference. All Brazilian municipality analyses are treated as exploratory because the exposure is satellite-detected fire activity rather than a smoke-exposure field. Sensitivity analyses comprise the development-inclusive index, SSP5-8.5 projections, alternative imputation and index-construction choices, leave-one-region-out and region-reassignment checks, unweighted regional capacity means, and Monte Carlo perturbation of the projection ratios (see the gap analysis above). Analysis code, derived datasets and reproduction instructions are archived on Zenodo[46].

No p-value reported here is corrected for multiplicity. The regional capacity–mortality association is the primary test; the adjustment analyses report several closely related specifications of the same confound, and are exploratory and unadjusted for multiplicity. Only one of them clears $\alpha = 0.05$, and it should be read in that light.

The Brazilian panel is not a sensitivity analysis of that test: it uses different data to answer a different question, and no result from it can confirm or refute the regional finding. Within it, robustness checks comprise binary and tertile capacity codings, municipality-trend and state-by-year-month adjustments, conditional Poisson estimation, extension to 2015–2020, and the Mexico and Chile replication attempts. The modelled-landscape-fire PM2.5 exposure substitution, the joint fire-plus-smoke model and the widened outcome set were added after the fire-count estimates had been seen, in direct response to the exposure-proxy limitation recorded above, and are reported as post hoc.

Use of large language models

The author used a large-language-model assistant (Claude, Anthropic) for coding and debugging support while developing the analysis scripts and while preparing the manuscript text. The model was not used to generate data, select analyses, or produce results: every numerical value reported here was computed by the archived, version-controlled pipeline described above and checked by the author against that pipeline's outputs. The author takes full responsibility for the content of this publication.

Results

Main analysis: global inequality between current health-system capacity and projected mortality growth

Health-system capacity is strongly inversely co-located with projected smoke-mortality growth (region-level Spearman $\rho = -0.85$, $n = 17$; regional permutation $p = 5\times10^{-5}$, 100,000 draws; Fig. 1). In practical terms, regions with weaker present-day health systems tended to rank among those projected to experience the fastest proportional growth in smoke-attributable deaths. This rank-order alignment remained strong across the examined sensitivity analyses (alternative projection scenarios, region reassignment, regional weighting, index constructions, complete-case analysis, leave-one-region-out re-estimation, and Monte Carlo perturbation of the mortality ratios), with correlations ranging from approximately − 0.69 to − 0.92 (Fig. 2C; Additional file 1: Figs. S2, S3 and S5 and Tables S1, S3, S4 and S11). Two checks are worth reporting numerically: under the higher-emissions SSP5-8.5 scenario ρ was − 0.69 to − 0.72 (all permutation $p \le 0.004$), and under independent lognormal perturbation of the 17 ratios calibrated to the published global 95% confidence interval, the median perturbed correlation was − 0.80 with all of 10,000 draws below $\rho = -0.5$. The attenuation under SSP5-8.5 is concentrated in two regions: Canada and the USA, whose projected ratios rise from 1.27 and 1.54 under SSP2-4.5 to 11.01 and 9.56, moving both from the lowest-growth tier to mid-table. Excluding them, the two scenarios agree closely ($\rho = -0.80$ under SSP2-4.5 and − 0.82

under SSP5-8.5, n = 15), so the alignment weakens specifically among high-capacity regions in the scenario where the climate-driven emission component is largest (see Discussion, interpretation), not uniformly across the distribution. The alignment also persists at country resolution: relating each of up to 157 countries' own projected mortality ratio to its own capacity score gives ρ = −0.56 to − 0.64 across baseline filters, with region-cluster bootstrap 95% CIs excluding zero in every variant, and capacity is similarly aligned with projected annual deaths in 2095–2099 per million people in the recent population (ρ = −0.57 to − 0.59 across variants). The denominator is each country's latest available World Bank population estimate from 2020 onward, so this measure is not a future per-capita mortality rate. Together, these checks suggest that the alignment is not confined to fold-change outcomes or regional aggregation (Additional file 1: Table S4b). Nor does it lean on the mixed granularity of the source regions. Six of the 17 regions contain only one or two index countries and sit mostly at the high-capacity end (region size and capacity correlate at ρ = −0.54), yet dropping all six at once leaves the alignment intact among the 11 multi-country regions (ρ = −0.69; permutation p = 0.023), the dropped block shows the same gradient on its own (ρ = −0.77 across those six units), and weighting regions by their country count or population strengthens the correlation (ρ = −0.92 and − 0.91; Additional file 1: Table S11). Nor is it specific to this composite: substituting the WHO UHC service coverage index reproduces the alignment (regional ρ = −0.70, permutation p = 0.003; country resolution ρ = −0.43, cluster 95% CI − 0.65 to − 0.04; Additional file 1: Table S4c). Because the index contains two financing indicators, and health expenditure per capita is strongly development-correlated, we also rebuilt it from physical capacity alone — hospital beds, physicians and nurses, with both financing measures removed: the alignment persists (ρ = −0.77, permutation p = 0.0005). It does not rest on any single component either, the correlation staying between − 0.79 and − 0.83 across five leave-one-indicator-out variants (Additional file 1: Table S18). This study compares current health-service supply with a fixed-adaptation future burden scenario; that temporal mismatch is a feature of the research question rather than something extrapolation can resolve. As a supplementary stress test only, under naive continuation of each country's 1990–2024 indicator trends the inverse alignment is preserved (2100 horizon: ρ = −0.68 under SSP2-4.5 and − 0.74 under SSP5-8.5, both permutation p ≤ 0.004; Additional file 1: Table S14 and Figs. S7 and S8), because trend continuation raises capacity levels while largely preserving the regional ordering; Fig. S7 maps the resulting capacity–burden gap at both horizons. Table 2 shows all 17 projection regions with their population-weighted mean capacity and projected mortality ratio.

Table 2
Characteristics of the 17 Zhao et al. projection regions.

Region	Countries (n)	Pop.-weighted mean capacity	Mortality ratio (SSP2-4.5)
Northern Africa	29	-1.10	10.66
Southern Africa	23	-1.06	10.10
South Asia	8	-1.03	9.57
Southeast Asia	11	-0.74	3.34
Northern South America	7	-0.32	12.41
Middle East	15	-0.29	5.49
Central America	12	-0.19	6.45
Brazil	1	+ 0.40	4.10
Central Asia	6	+ 0.44	7.10
China	1	+ 0.45	3.20
Southwest South America	4	+ 0.83	7.35
Oceania	8	+ 0.87	1.68
Russia	1	+ 1.22	1.77
Canada	1	+ 1.30	1.27
Europe	35	+ 1.36	1.03
USA	1	+ 2.01	1.54
Korea and Japan	2	+ 2.15	1.59

Projected mortality ratios are from Zhao et al. (2025), Source Data Fig. 4a[18]. *Country counts and population-weighted health-system capacity scores were calculated for this study. Country-level capacity was combined within each region using population-weighted means. Country-level capacity distributions within each projection region are shown in Additional file 1: Fig. S1; inference is based on these 17 regional points. Note: Equatorial Africa (18th Zhao region) is unpopulated in the baseline crosswalk; the sensitivity analysis reassigning countries to this region is reported in Additional file 1: Table S4.*

Within Table 2, Northern South America shows the largest positive departure from the overall declining gradient (mortality ratio 12.41, the highest in the table, at capacity − 0.32, the fifth-lowest of 17); the rank-based correlation limits its leverage, and dropping the region leaves the association essentially unchanged (ρ = −0.86). The single most influential omission is instead Southeast Asia, whose removal strengthens the association to ρ = −0.89 (Additional file 1: Table S3).

The lowest-capacity quartile of countries faces a mean 9.2-fold projected increase in fire-PM2.5 mortality versus 2.0-fold for the highest quartile. A descriptive stratification by World Bank income group produced a similar gradient (low-income 10.0×, lower-middle 8.5×, upper-middle 6.0×, high-income 2.7×; Fig. 2B). Both contrasts are unweighted means across countries: the quartiles contain near-equal numbers of the 165 mapped countries

(42/41/41/41), and each country carries the (repeated) mortality ratio of its projection region, so these are descriptive summaries only and are not treated as independent country-level inference (see Methods, gap analysis). The income-group gradient is likewise not fully independent of the capacity relationship, because income, capacity, and development are strongly related.

Expressed as population rather than correlation, the misalignment has a definable size. Some 0.88 billion people (10.9% of the analysed population, in 25 countries) live where the lowest capacity tertile coincides with the highest projected-growth tertile, the deep-red combination of Fig. 1B; Nigeria, Ethiopia, Egypt and Sudan are the largest. The count is 0.70 billion under the stricter quartile definition and 3.46 billion (42.9%) under a median split, and the three sets nest (Additional file 1: Table S15). At the opposite corner, 2.74 billion people live where the highest capacity tertile meets the lowest projected growth. These counts inherit the regional resolution of the projections (every country carries its region's ratio), so they size the misalignment rather than estimating country-specific risk.

Equity and sensitivity to national wealth

The capacity deficit is concentrated among low-income populations (concentration index = − 0.30; the 20% of the world's population living in the lowest-GDP-per-capita countries bears 32% of the global deficit; n = 187 countries with capacity, GDP per capita and population; Fig. 2A). This concentration is not an artifact of the benchmark choice: across 75th-, 80th-, and 95th-percentile benchmarks the concentration index remains between − 0.38 and − 0.27 and the poorest quintile's share of the deficit remains above 30% (Additional file 1: Table S8). Because countries are ranked by GDP per capita and capacity is strongly GDP-correlated (regional Spearman ρ = 0.94; country-level 0.86), some concentration is close to definitional; we read this index as a distributional summary of where the deficit sits, not as evidence independent of the development gradient (see Methods, equity and wealth sensitivity). Given this close capacity–wealth correlation, we then asked how much of the gap relationship wealth can explain. The primary index contains no macroeconomic variables (its inputs are health-sector workforce, hospital beds and health financing only) and a development-inclusive alternative that adds GDP per capita and internet use gives a similar association (ρ = −0.83, permutation $p = 1\times10^{-4}$), so the finding was not sensitive to the index specifications examined (Additional file 1: Table S1). Adjusting for GDP per capita at the regional level removed most of the association: the regional partial Spearman correlation was ρ = −0.19 (Freedman–Lane permutation p = 0.47; region-resampling bootstrap 95% confidence interval (CI) − 0.61 to + 0.06; Fig. 3). Three properties of this single-proxy adjustment, however, mean its null carries little information. It is symmetric: the unadjusted GDP–ratio association is comparably strong (regional ρ = −0.87; permutation p < 0.001), and partialling capacity out of it likewise leaves no statistically distinguishable residual (partial ρ = −0.40; Freedman–Lane p = 0.15; bootstrap 95% CI − 0.86 to + 0.05; Additional file 1: Table S10b) — adjusting either gradient for the other nulls both. It is nearly powerless: at 17 regions and ρ = 0.94 collinearity, the test only rejects when the partial correlation reaches 0.49 in absolute value, and its power against a true residual partial correlation of 0.2, 0.3 or 0.4 is 0.09, 0.15 and 0.25 respectively, with 80% power first reached at 0.7 (Additional file 1: Table S10c). And it adjusts for a noisy proxy rather than the construct: a partial-identification analysis shows the GDP-adjusted association is only weakly identified unless GDP is treated as a perfectly reliable measure of latent development (Fig. 3D; Additional file 1: Table S13). The single-proxy nulls therefore measure the collinearity of capacity with one development proxy — they are not a verdict that development rather than capacity drives the alignment.

In an exploratory sensitivity analysis, replacing GDP alone with a six-indicator development factor (log GDP per capita, internet use, urban share, electricity access, basic sanitation and secondary enrolment) — less collinear with capacity (ρ = 0.89) — yielded a larger residual capacity association (partial ρ = −0.64; Freedman–Lane p = 0.010; bootstrap 95% CI − 0.88 to − 0.07).

The estimate remained negative in every leave-one-region-out analysis, between − 0.45 and − 0.71, although its p-value exceeded 0.05 after omission of either of two regions, Oceania and Northern South America (Additional file 1: Table S10e). The reverse adjustment is null (measured development given capacity: partial ρ = +0.10; p = 0.74; Additional file 1: Table S10d).

Because this was one of several exploratory specifications unadjusted for multiplicity, involved only 17 regions, and depends on how both capacity and development were measured (with different indicators, different reliability and continuing high collinearity) it does not establish that capacity contributes independently of development. It indicates only that the observed alignment was not eliminated by this particular representation of development. The three adjusted tests do not differ meaningfully in resolving power either: their detection limits are 0.49, 0.51 and 0.54, and none reaches 80% power against a true residual association below about 0.7 (Additional file 1: Table S10c).

The relationship to demographic structure is likewise beyond what this design can resolve. Adding demographic covariates (fertility, share aged 65+, population growth) leaves a partial correlation of − 0.41 (p = 0.14; bootstrap 95% CI − 0.87 to + 0.54), but that four-covariate model has 11 residual degrees of freedom and its own detection limit is 0.54 — larger than the estimate itself. This estimate is therefore an unresolved test rather than a demonstrated attenuation: it is neither evidence that demographic structure absorbs the alignment nor evidence that it does not. The country-resolution estimates show the same pattern — the development-factor partial − 0.40 (cluster 95% CI − 0.68 to + 0.01), attenuating to − 0.17 with demographic covariates (Additional file 1: Table S10d), while the country-resolution GDP-adjusted partials are negative in every variant (− 0.35, − 0.25 and − 0.11; Additional file 1: Table S4b).

No adjustment shows development uniquely absorbing the alignment, and none identifies a causal capacity effect. Because income-group-based imputation could in principle strengthen the income–capacity relationship mechanically, we recomputed the equity and GDP-adjusted quantities on complete-case data: the concentration index (− 0.31), the poorest-quintile share (32.2%), and the null adjusted association (partial ρ = +0.01, permutation p = 0.96) were all preserved (Additional file 1: Table S10). The complete-case samples are as follows: of the 191 index countries, 154 have no imputed capacity indicator, of which 137 map to a populated projection region and enter the complete-case gap correlation; the equity and GDP-adjusted quantities additionally require GDP per capita and population, leaving 150 countries for the equity statistics and 133, spanning the 17 regions, for the GDP-adjusted partial correlation. Taken together (see Methods, equity and wealth sensitivity): the inverse association is strong unadjusted; it was not eliminated by the six-indicator representation of development, though that does not establish an independent capacity contribution; and it cannot be separated from the demographic structure that development tracks (see Methods, gap analysis) with the current 17-region data. This qualifies attribution, not existence: the unadjusted co-location reported above (response capacity lowest where projected growth is largest) is a robust feature of these data within the Zhao et al. projection framework, whichever underlying gradient produces it, and it is that descriptive misalignment, not a causal capacity effect, that the abstract and conclusions claim.

Separate exploratory analysis: within-country fire activity and cause-specific mortality

In the Brazilian panel the interaction between fire activity and hospital-bed capacity on ill-defined-cause mortality was strongly specification-dependent: negative under the primary municipality-by-calendar-month specification (β_2 = −0.100 per SD beds; p = 0.009) but null under three of the six specifications examined, so it is not a robust feature

of these data. It did not survive the exposure substitution: on the common smoke-available sample, with both exposures rescaled to unit standard deviation, the interaction fell from $\beta_2 = -0.085$ ($p = 0.009$) under fire counts to -0.013 ($p = 0.71$) under modelled landscape-fire PM2.5. In a joint model including both exposures, the contrast between their interaction coefficients was not significant (fire minus smoke, -0.083; $p = 0.12$). Neither the Mexican nor the Chilean replication reproduced the pattern, and the Mexican external-cause placebo failed under fire counts, indicating that fire detections partly proxy rural and agricultural geography. Taken together these results give no robust evidence that capacity modifies the fire–mortality association, and they bear on the mechanism question they pose rather than on the global alignment. All estimates, falsification tests and replication details are in Additional file 1 (Tables S5–S6 and S16; Supplementary Note; Fig. S9).

Discussion

Principal findings

Two findings constitute the global inequality result. First, a strong unadjusted regional inverse association: the regions facing the largest projected increases in landscape-fire smoke mortality are those with the lowest recent health-system capacity (region-level $\rho = -0.85$, $n = 17$). Second, the capacity deficit (measured on a primary index built from health-sector indicators alone, without macroeconomic variables) is concentrated among populations living in the lowest-income countries (concentration index = − 0.30, stable between − 0.38 and − 0.27 across alternative benchmarks and preserved under complete-case analysis).

Separately, and bearing on a different estimand that can neither confirm nor refute those two findings, an exploratory Brazilian panel did not yield robust evidence that health-system capacity modifies the association between fire activity and ill-defined-cause mortality: the interaction was specification-dependent, and the replications in Mexico and Chile were null. Replacing fire-detection counts with a modelled landscape-fire PM2.5 field yielded an interaction estimate not statistically distinguishable from zero, as well as a null Mexican placebo, so what fragility remained is best read as a property of the fire-detection proxy rather than of smoke exposure. The first two findings constitute the principal global inequality result; the third is a hypothesis-generating mechanism analysis.

Interpretation

Two structural features of the underlying projections shape what this alignment can and cannot mean. First, the Zhao et al. ratios hold community vulnerability and adaptation fixed, so the channels through which health-system capacity would act on future smoke mortality (surveillance, warning, treatment) are excluded from the outcome by construction. The observed correlation is therefore a co-location of present capacity with projected burden growth; no analysis of this outcome, at any sample size, could estimate an effect of capacity on that growth. Second, the projected growth is predominantly demographic. In Zhao et al.'s published sensitivity experiments (Additional file 1: Table S12), evolving population size and age structure alone raise projected global fire-PM2.5 deaths 6.4-fold (more than the full projection's 5.9-fold), whereas evolving climate-driven emissions alone raise them by 3% under SSP2-4.5 (46% under SSP5-8.5). Cross-region variation in the mortality ratio therefore partly tracks where population growth and ageing will concentrate — patterns that are themselves development-correlated. A strong inverse correlation with any development-patterned index is thus partly expected on demographic grounds alone. The Source Data provide no regional decomposition against which capacity could be correlated with the exposure-driven component only; such a decomposition would be the sharpest refinement of this design.

The global gap thus identifies a mismatch between recent health-system capacity and projected future landscape-fire smoke risk; it does not establish that capacity causally reduces smoke mortality. Zhao et al.'s own inequality analysis already shows that the projected deaths concentrate among low-income populations[18]. What the present analysis adds is the response side of that inequality (the health-system resources that would have to absorb the burden are arrayed inversely to it) together with explicit boundaries on what this added gradient can mean. The double-vulnerability pattern itself, in which populations facing the greatest climate-health risk have the least service coverage, is consistent with existing composite risk frameworks; the contribution here is its quantification for projected landscape-fire smoke mortality specifically, with resolution-matched inference and with symmetry, detection-limit, benchmark-substitution, region-composition and capacity-horizon analyses that state exactly what the alignment can and cannot support.

That the direction of the pattern is expected does not substitute for measuring it: expected relationships can prove weak, non-monotonic, or driven by a handful of units, and the analyses above show this one is none of those. Both capacity and risk track development, and how far the two can be told apart depends on how development is measured. Against its single conventional proxy, they cannot: capacity and GDP per capita carry no demonstrable information beyond each other (both directional partials null; Additional file 1: Table S10b), and that test could not have detected a moderate residual association in any case (detection limit $|\rho| \geq 0.49$; Additional file 1: Table S10c). In the exploratory six-indicator development-factor specification, the residual capacity association was statistically distinguishable from zero (partial $\rho = -0.64$, $p = 0.010$, 95% CI − 0.88 to − 0.07), whereas the reverse partial was not (+ 0.10, $p = 0.74$; Additional file 1: Table S10d). However, because the constructs were measured differently, remained highly collinear, and were analysed across only 17 regions, this result does not establish a contribution of capacity independent of development.

Demographic structure is a different matter: with four covariates at 17 regions the demographic-adjusted test resolves nothing smaller than $|\rho| = 0.54$, and the estimate it returns (− 0.41) falls below that, so the question of whether demographic structure absorbs the alignment is left open by these data rather than answered by them. That is a limit of resolution, not a finding, and it is consistent with health-system supply being embedded in (and projected burden growth being driven by) the same demographic and structural conditions, such as poverty, housing quality, energy access and baseline population health. This residual entanglement is itself policy-relevant. It cautions against expecting narrowly targeted infrastructure investments (additional hospital beds or clinical staff in isolation) to offset projected smoke-mortality growth where broader development lags.

The policy content of these results is therefore descriptive: they identify where projected smoke-burden growth and present health-system capacity are most misaligned — information relevant to prioritising adaptation finance, surveillance and monitoring investment, and preparedness planning. Whether strengthening clinical capacity would reduce future smoke mortality is a causal question these data cannot answer, for the reason given above. The exploratory within-country analysis reported in Additional file 1 did not change this reading: the fire-count interaction was specification-dependent and did not survive substituting modelled landscape-fire PM2.5 for fire detections, so what those models tracked is better described as burning geography than as smoke exposure. Because that analysis is hypothesis-generating and concerns present-day fire activity rather than projected burden growth, it is not evidence of differential cause-of-death ascertainment and bears on the mechanism question it poses, not on the global alignment.

Estimating whether capacity independently protects against smoke mortality would require a different design built on three elements. First, realized smoke-attributable mortality rather than fixed-adaptation projections. Second,

exposure identification from modelled smoke rather than local fire counts: plume-dispersion or trajectory-based smoke products (for example, smoke-day classifications informed by the HYbrid Single-Particle Lagrangian Integrated Trajectory (HYSPLIT) model), chemical-transport model output, satellite-and-model-fused daily smoke-PM2.5 estimates, or atmospheric-composition reanalyses such as the Copernicus Atmosphere Monitoring Service (CAMS) capture long-range transport, and smoke arriving from distant upwind fires supplies variation in exposure plausibly exogenous to local conditions. Third, exogenous variation in capacity itself, such as facility openings or closures, staffing shocks, or programme rollouts with clean timing. Thus, the geographical inequality is robust, but its attribution specifically to health-system capacity is not identifiable in these data.

Strengths and limitations

The analysis uses public data and an archived analysis pipeline; the Brazilian panel is treated as exploratory, and significant and null results are reported alike. No p-value here is corrected for multiplicity: the regional capacity–mortality association is the primary test, but the adjustment analyses report several specifications, and the one that clears $\alpha = 0.05$ should be read in that light rather than as an isolated confirmatory result. Its limitations fall into three groups.

First, projection limitations. This analysis does not establish that capacity causally reduces smoke mortality; the reason is structural rather than statistical, and is given in the Background. Every claim involving projected mortality derives from a single upstream projection framework (Zhao et al.); the within-framework sensitivity checks demonstrate stability to analytic choices but cannot test whether the regional pattern is specific to that framework, and an alternative projection with comparable regional resolution was not evaluated for cross-model corroboration. The projected mortality ratios were treated as fixed regional point estimates: we did not propagate uncertainty from the underlying climate models, fire-emission projections, demographic assumptions, baseline mortality estimates, or exposure–response functions, and region-specific uncertainty with cross-region covariance is not available, so the calibrated Monte Carlo perturbation (see Methods, gap analysis; Additional file 1: Table S11) bounds but cannot fully capture shared structural errors. The effective inferential sample is 17 projection regions, not 165 independent countries — except in the country-resolution robustness analysis, whose own map-grade data limitations are stated in Methods. The outcome is proportional (fold-change) growth: because several low-income regions had small 2010–2014 baselines, the ratio scale mechanically amplifies proportional signals there, and projected relative growth does not directly equal future absolute burden, which also depends on baseline population size and disease rates; the country-resolution check, in which capacity aligns similarly with projected annual deaths in 2095–2099 per million people in the recent population (Additional file 1: Table S4b), provides a complementary scale comparison but cannot resolve the effects of future population size and demographic composition. Region-based estimates and the double-jeopardy population counts are also conditional on the frozen country–region crosswalk. In particular, the baseline assigns the two Congos, Gabon and Equatorial Guinea to Southern Africa, although Zhao et al.'s Supplementary Fig. 5 places this central-African area within Equatorial Africa[18]. The selected-country reassignment sensitivity examines one alternative and does not validate the full allocation.

Second, measurement and confounding. The GDP-adjusted null does not show that development rather than capacity is the operative gradient, since adjusting for a single development proxy cannot license that conclusion; conversely, the development-factor result shows a residual capacity association, not a capacity effect — the factor is one defensible measurement of development, and its estimate rests on the same 17 regions. Adding demographic covariates spends four degrees of freedom on 17 observations and yields a test that cannot resolve anything below $|\rho| = 0.54$, so demography and capacity remain unseparated: that configuration is reported for

completeness, not as evidence either way. The capacity index is cross-sectional, combines indicators measured in different recent years, and represents recent rather than end-of-century capacity; under naive continuation of 1990–2024 indicator trends the alignment persists to 2100 (Additional file 1: Table S14), but SSP-consistent capacity trajectories were not evaluated in this study, so end-of-century capacity forecasting is beyond these data. The index does not measure the smoke-specific interventions most proximate to exposure (smoke forecasting, air-quality monitoring, public-warning coverage, clean-air shelters and filtration access, and death-certification quality; see Methods). It tracks the WHO UHC service coverage index closely (ρ = 0.85), which supports it as a measure of general health-service supply but does not validate it as a measure of smoke-specific response capability; given its 0.94 regional correlation with GDP, the index is best read as measuring the formal-sector health-service supply that accompanies development, not validated smoke-response capability.

Third, limitations of the exploratory within-country analysis. Its design is ecological, its fire-detection exposure is an imperfect proxy for breathed smoke, and the modelled smoke field that replaces it is gridded at 0.25 degrees, so exposure misclassification would attenuate interaction estimates toward zero and a true null cannot be separated from that attenuation. It does not bear on the regional finding in either direction: its estimand is the modification of a contemporaneous fire-mortality association within one country, whereas the regional analysis concerns projected end-of-century burden growth under fixed adaptation. Brazil does contribute to its own regional capacity aggregate, and the regional correlation is essentially unchanged when it is dropped (ρ − 0.846 to − 0.856; Additional file 1: Table S3). All coefficients from that analysis are read as descriptive associations; the full set of limitations, including interference across municipal boundaries, the time-invariant 2019 capacity measure and the heterogeneous-effects concern for two-way fixed effects[47, 48], is given in Additional file 1.

Conclusions

Across 17 world regions, recent (2015–2024) health-system capacity is strongly and inversely aligned with the projected end-of-century growth in landscape-fire PM2.5 mortality under SSP2-4.5 in the Zhao et al. (2025) projection framework, pending cross-model corroboration: the regions expected to face the largest proportional increases are those with the weakest current formal-sector health-service supply (clinical workforce, hospital beds and health financing) and this capacity deficit is concentrated among populations living in the world's lowest-income countries. The alignment persists when an external WHO capacity benchmark is substituted for the index and at country resolution. It is descriptive by design and does not estimate a causal effect of capacity. Adjusted estimates were sensitive to how development and demographic structure were represented, so a capacity-specific association could not be cleanly separated.

The contribution is the quantified measurement of how strong, how geographically stable and how distributionally concentrated that misalignment is within the Zhao et al. (2025) region-resolved projection framework, together with explicit boundaries on what the measurement can support. Sized in population rather than in correlation, 0.88 billion people, one in nine of those analysed, live where the weakest capacity tertile meets the fastest projected growth. That is a count of people in a doubly disadvantaged position, and the kind of quantity against which adaptation finance and preparedness planning are allocated; it is not an estimate of lives that capacity investment would save. Establishing whether capacity independently protects against smoke mortality will require direct transported-smoke exposure estimates and plausibly exogenous variation in health-system capacity.

Abbreviations

CAMS
Copernicus Atmosphere Monitoring Service
CI
Confidence interval
CNES
Cadastro Nacional de Estabelecimentos de Saúde
CONAF
Corporación Nacional Forestal
DATASUS
Departamento de Informática do Sistema Único de Saúde
DEIS
Departamento de Estadísticas e Información de Salud
FIRMS
Fire Information for Resource Management System
GDP
Gross domestic product
HYSPLIT
HYbrid Single-Particle Lagrangian Integrated Trajectory
IBGE
Brazilian Institute of Geography and Statistics (Instituto Brasileiro de Geografia e Estatística)
INEGI
National Institute of Statistics and Geography (Instituto Nacional de Estadística y Geografía)
INPE
National Institute for Space Research (Instituto Nacional de Pesquisas Espaciais)
PC1
First principal component
PM2.5
Fine particulate matter with aerodynamic diameter ≤ 2.5 μm
PPP
Purchasing power parity
REB
Research ethics board
SD
Standard deviation
SIM
Mortality Information System (Sistema de Informação sobre Mortalidade)
SSP
Shared Socioeconomic Pathway
TCPS 2
Tri-Council Policy Statement:Ethical Conduct for Research Involving Humans (2nd edition)
UHC
Universal health coverage

Declarations

Ethical approval

The study used publicly accessible, deidentified administrative mortality records and aggregate environmental and health-system data obtained through the providers' public distribution channels. No direct identifiers were accessed, no individuals were recruited or contacted, and no attempt was made to identify individuals. Mortality records were aggregated to municipality-month or region-year units for statistical analysis, and no individual-level mortality records are reported. The author assessed the study as falling within the public-information exemption described in the Tri-Council Policy Statement: Ethical Conduct for Research Involving Humans (TCPS 2, 2022), Article 2.2[49]. This was the author's own assessment; no formal ethics committee approval, waiver or exemption determination was obtained.

Consent to participate

Not applicable. The study involved secondary analysis of publicly accessible, deidentified records and aggregate datasets, without recruitment or direct contact with individuals. No new individual consent was sought.

Consent to Publish

Not applicable. The manuscript and supplementary information contain no identifiable individual-level information.

Competing interests

The author declares that they have no competing interests.

Funding

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

Author Contribution

Y.L. was solely responsible for the conceptualization, methodology, software, data curation, formal analysis, visualization, and writing of the original draft, as well as review and editing of the manuscript. Y.L. read and approved the final manuscript.

Acknowledgements

Not applicable.

Data Availability

The datasets generated during the current study are available in the Zenodo repository, https://doi.org/10.5281/zenodo.21815304. The repository contains derived analytic datasets, analysis code, and reproduction instructions. Original source datasets are available from the providers identified in the manuscript and repository documentation. The Zhao et al. (2025) Source Data spreadsheets are not included in this repository,

and selected source-value fields are omitted from the deposited tables. Some figures display these source quantities. The original Source Data are available with the article at https://doi.org/10.1038/s41586-025-09612-9. Instructions for obtaining the required inputs, restoring the omitted values, and reproducing the analyses are provided in the repository.

References

1. Xu R, Yu P, Abramson MJ, Johnston FH, Samet JM, Bell ML, et al. Wildfires, global climate change, and human health. N Engl J Med. 2020;383:2173–81. 10.1056/NEJMsr2028985.
2. Liu JC, Pereira G, Uhl SA, Bravo MA, Bell ML. A systematic review of the physical health impacts from non-occupational exposure to wildfire smoke. Environ Res. 2015;136:120–32. 10.1016/j.envres.2014.10.015.
3. Reid CE, Brauer M, Johnston FH, Jerrett M, Balmes JR, Elliott CT. Critical review of health impacts of wildfire smoke exposure. Environ Health Perspect. 2016;124:1334–43. 10.1289/ehp.1409277.
4. Cascio WE. Wildland fire smoke and human health. Sci Total Environ. 2018;624:586–95. 10.1016/j.scitotenv.2017.12.086.
5. Black C, Tesfaigzi Y, Bassein JA, Miller LA. Wildfire smoke exposure and human health: significant gaps in research for a growing public health issue. Environ Toxicol Pharmacol. 2017;55:186–95. 10.1016/j.etap.2017.08.022.
6. Johnston FH, Henderson SB, Chen Y, Randerson JT, Marlier M, DeFries RS, et al. Estimated global mortality attributable to smoke from landscape fires. Environ Health Perspect. 2012;120:695–701. 10.1289/ehp.1104422.
7. Jones MW, Abatzoglou JT, Veraverbeke S, Andela N, Lasslop G, Forkel M, et al. Global and regional trends and drivers of fire under climate change. Rev Geophys. 2022;60:e2020RG000726. 10.1029/2020RG000726.
8. n der Werf GR, Randerson JT, Giglio L, van Leeuwen TT, Chen Y, Rogers BM, et al. Global fire emissions estimates during 1997–2016. Earth Syst Sci Data. 2017;9:697–720. 10.5194/essd-9-697-2017.
9. Burke M, Driscoll A, Heft-Neal S, Xue J, Burney J, Wara M. The changing risk and burden of wildfire in the United States. Proc Natl Acad Sci U S A. 2021;118:e2011048118. 10.1073/pnas.2011048118.
10. Ford B, Val Martin M, Zelasky SE, Fischer EV, Anenberg SC, Heald CL, et al. Future fire impacts on smoke concentrations, visibility, and health in the contiguous United States. Geohealth. 2018;2:229–47. 10.1029/2018GH000144.
11. Borchers Arriagada N, Palmer AJ, Bowman DMJS, Morgan GG, Jalaludin BB, Johnston FH. Unprecedented smoke-related health burden associated with the 2019–20 bushfires in eastern Australia. Med J Aust. 2020;213:282–3. 10.5694/mja2.50545.
12. Johnston FH, Hanigan IC, Henderson SB, Morgan GG, Bowman DMJS. Extreme air pollution events from bushfires and dust storms and their association with mortality in Sydney, Australia 1994–2007. Environ Res. 2011;111:811–6. 10.1016/j.envres.2011.05.007.
13. Matz CJ, Egyed M, Xi G, Racine J, Pavlovic R, Rittmaster R, et al. Health impact analysis of PM2.5 from wildfire smoke in Canada (2013–2015, 2017–2018). Sci Total Environ. 2020;725:138506. 10.1016/j.scitotenv.2020.138506.
14. Burnett R, Chen H, Szyszkowicz M, Fann N, Hubbell B, Pope CA 3rd, et al. Global estimates of mortality associated with long-term exposure to outdoor fine particulate matter. Proc Natl Acad Sci U S A. 2018;115:9592–7. 10.1073/pnas.1803222115.

15. Lelieveld J, Evans JS, Fnais M, Giannadaki D, Pozzer A. The contribution of outdoor air pollution sources to premature mortality on a global scale. Nature. 2015;525:367–71. 10.1038/nature15371.
16. GBD 2021 Risk Factors Collaborators. Global burden and strength of evidence for 88 risk factors in 204 countries and 811 subnational locations, 1990–2021: a systematic analysis for the Global Burden of Disease Study 2021. Lancet. 2024;403:2162–203. 10.1016/S0140-6736(24)00933-4.
17. van Donkelaar A, Hammer MS, Bindle L, Brauer M, Brook JR, Garay MJ, et al. Monthly global estimates of fine particulate matter and their uncertainty. Environ Sci Technol. 2021;55:15287–300. 10.1021/acs.est.1c05309.
18. Zhao J, Zheng B, Ciais P, Chen Y, Gasser T, Canadell JG, et al. Global warming amplifies wildfire health burden and reshapes inequality. Nature. 2025;647:928–34. 10.1038/s41586-025-09612-9.
19. Ebi KL, Vanos J, Baldwin JW, Bell JE, Hondula DM, Errett NA, et al. Extreme weather and climate change: population health and health system implications. Annu Rev Public Health. 2021;42:293–315. 10.1146/annurev-publhealth-012420-105026.
20. Haines A, Ebi K. The imperative for climate action to protect health. N Engl J Med. 2019;380:263–73. 10.1056/NEJMra1807873.
21. Kruk ME, Gage AD, Arsenault C, Jordan K, Leslie HH, Roder-DeWan S, et al. High-quality health systems in the Sustainable Development Goals era: time for a revolution. Lancet Glob Health. 2018;6:e1196–252. 10.1016/S2214-109X(18)30386-3.
22. Naghavi M, Makela S, Foreman K, O'Brien J, Pourmalek F, Lozano R. Algorithms for enhancing public health utility of national causes-of-death data. Popul Health Metr. 2010;8:9. 10.1186/1478-7954-8-9.
23. Mathers CD, Ma Fat D, Inoue M, Rao C, Lopez AD. Counting the dead and what they died from: an assessment of the global status of cause of death data. Bull World Health Organ. 2005;83:171–7.
24. França E, Teixeira R, Ishitani L, Duncan BB, Cortez-Escalante JJ, Morais Neto OL, et al. Ill-defined causes of death in Brazil: a redistribution method based on the investigation of such causes. Rev Saude Publica. 2014;48:671–81. 10.1590/S0034-8910.2014048005146.
25. Mello Jorge MHP, Laurenti R, Gotlieb SLD. [Quality analysis of Brazilian vital statistics: the experience of implementing the SIM and SINASC systems]. Cien Saude Colet. 2007;12:643–54. 10.1590/S1413-81232007000300014.
26. Kanso S, Romero DE, Leite IC, Moraes EN. [Geographic, demographic, and socioeconomic differences in quality of data on cause of death in Brazilian elders]. Cad Saude Publica. 2011;27:1323–39. 10.1590/S0102-311X2011000700008.
27. Chen C, Noble I, Hellmann J, Coffee J, Murillo M, Chawla N. University of Notre Dame Global Adaptation Index country index technical report. Notre Dame, IN: University of Notre Dame; 2015.
28. Watts N, Amann M, Arnell N, Ayeb-Karlsson S, Beagley J, Belesova K, et al. The 2020 report of The Lancet Countdown on health and climate change: responding to converging crises. Lancet. 2021;397:129–70. 10.1016/S0140-6736(20)32290-X.
29. Wagstaff A, Paci P, van Doorslaer E. On the measurement of inequalities in health. Soc Sci Med. 1991;33:545–57. 10.1016/0277-9536(91)90212-U.
30. World Bank. World Development Indicators [dataset], Washington DC, The World Bank Group. ; 2026. Capacity-index and gap-analysis indicators SH.MED.BEDS.ZS, SH.MED.PHYS.ZS, SH.MED.NUMW.P3, SH.XPD.CHEX.GD.ZS, SH.XPD.CHEX.PC.CD, NY.GDP.PCAP.PP.CD, IT.NET.USER.ZS, SP.POP.TOTL; development and demographic indicators used in the adjusted analyses SP.URB.TOTL.IN.ZS, EG.ELC.ACCS.ZS,

SH.STA.BASS.ZS, SE.SEC.ENRR, SP.DYN.TFRT.IN, SP.POP.65UP.TO.ZS, SP.POP.GROW (2015–2024 values). Retrieved via the World Bank Indicators API (portal: https://data.worldbank.org/indicator). Accessed 13 July and 4 August 2026.

31. World Health Organization. UHC service coverage index (SDG 3.8.1), indicator UHC_INDEX_REPORTED [dataset]. Geneva: World Health Organization Global Health Observatory. https://ghoapi.azureedge.net/api/UHC_INDEX_REPORTED. Accessed 4 August 2026.
32. Freedman D, Lane D. A nonstochastic interpretation of reported significance levels. J Bus Econ Stat. 1983;1:292–8. 10.1080/07350015.1983.10509354.
33. Ministério da Saúde. (Brazil), DATASUS/OpenDataSUS. Sistema de Informação sobre Mortalidade (SIM) — Mortalidade Geral, microdata files 2015–2020 [dataset]. https://dadosabertos.saude.gov.br/dataset/sim (files retrieved from the official OpenDataSUS distribution bucket. https://diaad.s3.sa-east-1.amazonaws.com/sim/). Accessed 14 July 2026.
34. Instituto Nacional de Pesquisas Espaciais (INPE). Programa Queimadas — focos de calor, satélite de referência (AQUA_M-T), annual files 2015–2021 [dataset]. https://dataserver-coids.inpe.br/queimadas/queimadas/focos/csv/anual/Brasil_sat_ref/. Accessed 14 July 2026.
35. Ministério da Saúde. (Brazil), DATASUS. Cadastro Nacional de Estabelecimentos de Saúde (CNES) — Leitos (hospital beds), 2019 monthly files [dataset]. https://cnes.datasus.gov.br/. Accessed 13 July 2026.
36. Instituto Brasileiro de Geografia e Estatística (IBGE). Estimativas da População residente para os municípios brasileiros, 2015–2020 [dataset]. https://ftp.ibge.gov.br/Estimativas_de_Populacao/. Accessed 14 July 2026.
37. Hu Y, Tian C, Yue X, Lei Y, Cao Y, Xu R, et al. Global high-resolution fire-sourced PM2.5 concentrations for 2000–2023. Earth Syst Sci Data. 2025;17:3741–56. 10.5194/essd-17-3741-2025.
38. Hu Y, Tian C, Yue X, GFED. Fire-sourced PM2.5 concentrations dataset [dataset]. Zenodo. 2025. https://doi.org/10.5281/zenodo.15493914.
39. Instituto Brasileiro de Geografia e Estatística (IBGE). Malha municipal digital do Brasil, 2019 [dataset]. https://geoftp.ibge.gov.br/organizacao_do_territorio/malhas_territoriais/malhas_municipais/municipio_2019/. Accessed 5 August 2026.
40. Instituto Nacional de Estadística y Geografía (INEGI, Mexico). Estadísticas de Defunciones Registradas (EDR) 2019, microdatos [dataset]. https://www.inegi.org.mx/programas/edr/. Accessed 13 July 2026.
41. NASA Fire Information for Resource Management System (FIRMS). MODIS Collection 6.1 standard-quality active-fire detections (MCD14ML), Mexico, 2019 [dataset]. https://doi.org/10.5067/FIRMS/MODIS/MCD14ML. Available from: https://firms.modaps.eosdis.nasa.gov/download/. Accessed 14 July 2026.
42. Departamento de Estadísticas e Información de Salud (DEIS). Ministerio de Salud (Chile). Defunciones por causa de muerte 1990–2022, cifras oficiales [dataset]. https://deis.minsal.cl/. Accessed 14 July 2026.
43. Corporación Nacional Forestal (CONAF. Chile). Estadísticas históricas de incendios forestales — superficie afectada por región, 1977–2024 [dataset]. https://www.conaf.cl/centro-documentals/estadisticas-historicas/. Accessed 14 July 2026.
44. Instituto Nacional de Estadística y Geografía (INEGI. Mexico). Directorio Estadístico Nacional de Unidades Económicas (DENUE), health-sector establishment directory, May 2025 snapshot [dataset]. https://www.inegi.org.mx/app/descarga/?ti=6. Accessed 13 July 2026.
45. Departamento de Estadísticas e Información de Salud (DEIS), Ministerio de Salud (Chile). Resumen Estadístico Mensual (REM), hospital bed-day series, 2014–2025 [dataset]. https://deis.minsal.cl/ (Datos

Abiertos, REM series). Accessed 14 July 2026.
46. Li Y. Recent health-system capacity and projected growth in landscape-fire smoke mortality: a global inequality analysis [dataset]. Zenodo; 2026. https://doi.org/10.5281/zenodo.21815304.
47. de Chaisemartin C, D'Haultfœuille X. Two-way fixed effects estimators with heterogeneous treatment effects. Am Econ Rev. 2020;110:2964–96. 10.1257/aer.20181169.
48. de Chaisemartin C, D'Haultfœuille X. Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: a survey. Econom J. 2023;26:C1–30. 10.1093/ectj/utac017.
49. Canadian Institutes of Health Research, Natural Sciences and Engineering Research Council of Canada, Social Sciences and Humanities Research Council of Canada. Tri-Council Policy Statement: Ethical Conduct for Research Involving Humans. 2022. Chapter 2, Article 2.2. https://ethics.gc.ca/eng/tcps2-eptc2_2022_chapter2-chapitre2.html. Accessed 15 October 2026.
50. Cameron AC, Gelbach JB, Miller DL. Robust inference with multiway clustering. J Bus Econ Stat. 2011;29(2):238–49. 10.1198/jbes.2010.07136.
51. Instituto Nacional de Estadísticas (INE. Chile). Censo de Población y Vivienda 2017, regional population counts [dataset]. https://censo2024.ine.gob.cl/estadisticas-2017/. Accessed 15 October 2026.

Figures

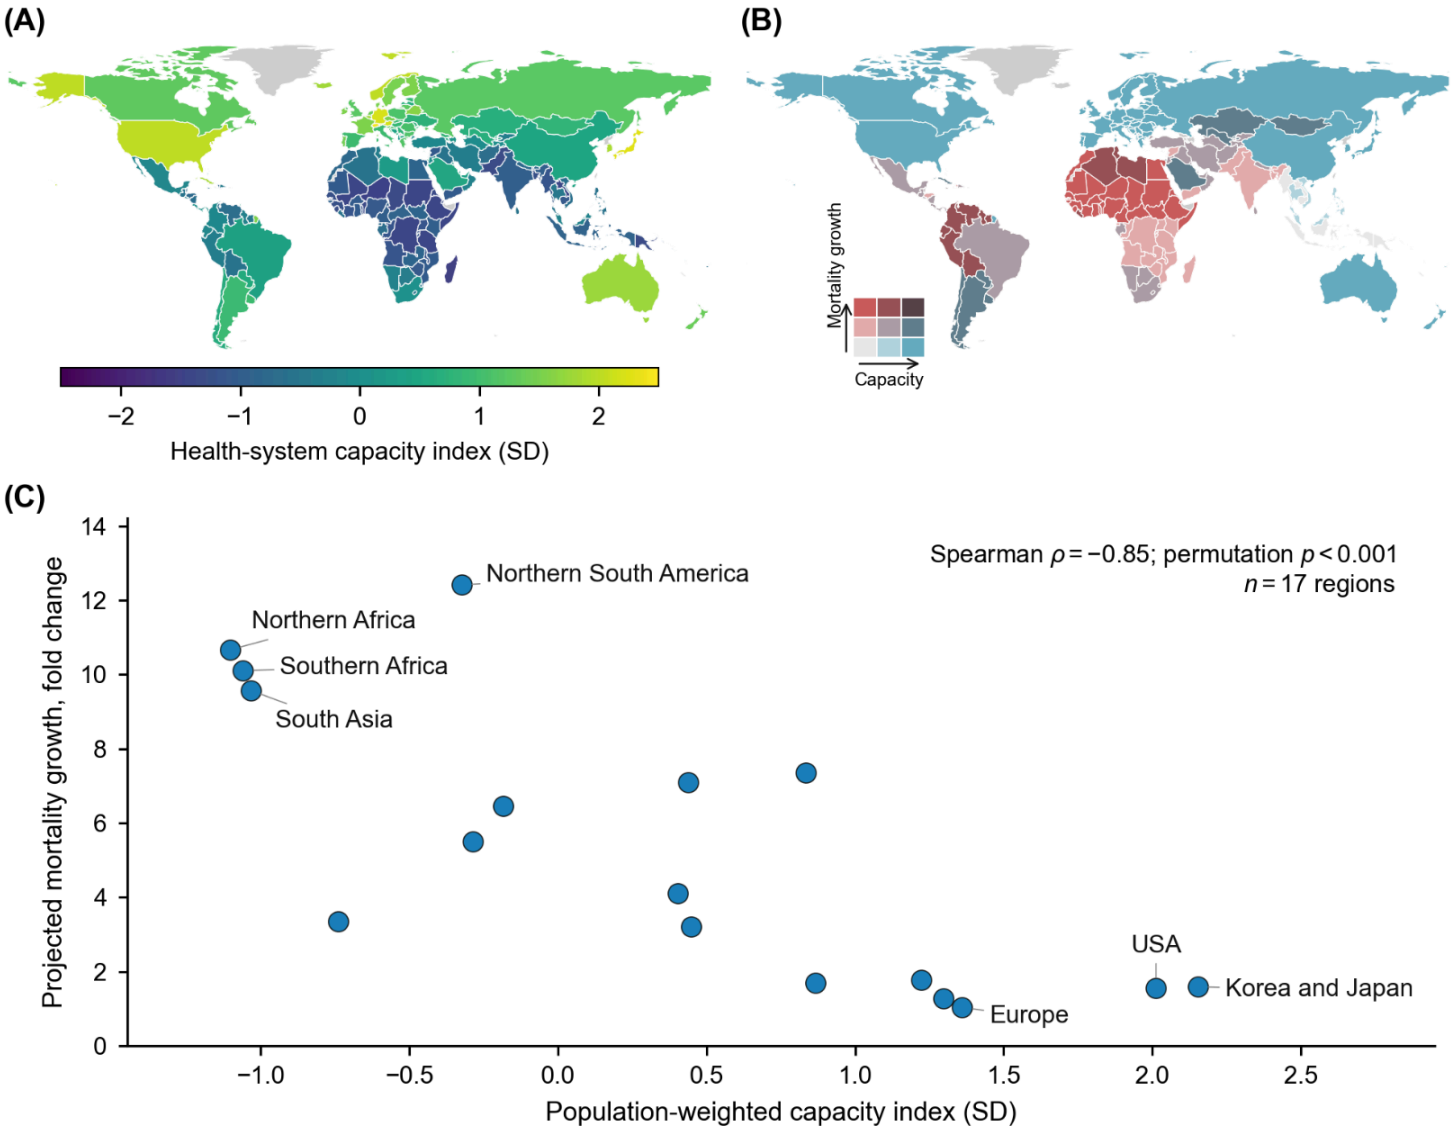


Figure 1

Co-location of recent health-system capacity and projected fire-PM2.5 mortality growth. (A) Recent health-system capacity index (primary, health-only; standardised) for 191 countries. (B) The two together: each country is shaded by the combination of its capacity tertile and its projected-mortality-growth tertile (3 × 3 key; capacity increases to the right, projected growth upward). Projected growth is the 2095–2099 versus 2010–2014 fold change under SSP2-4.5 from Zhao et al. (2025), Source Data Fig. 4a[18], carried at the value of the country's projection region rather than as an independent country-level estimate. Deep red marks countries in the lowest capacity tertile facing the highest projected growth; teal marks the converse. (C) Region-level association between population-weighted capacity and projected mortality growth (Spearman ρ = −0.85; regional permutation p < 0.001; n = 17 regions; equal marker sizes; each point is one projection region). Grey shading indicates countries not in the capacity index or not mapped to a projection region. Country boundaries: Natural Earth (public domain).

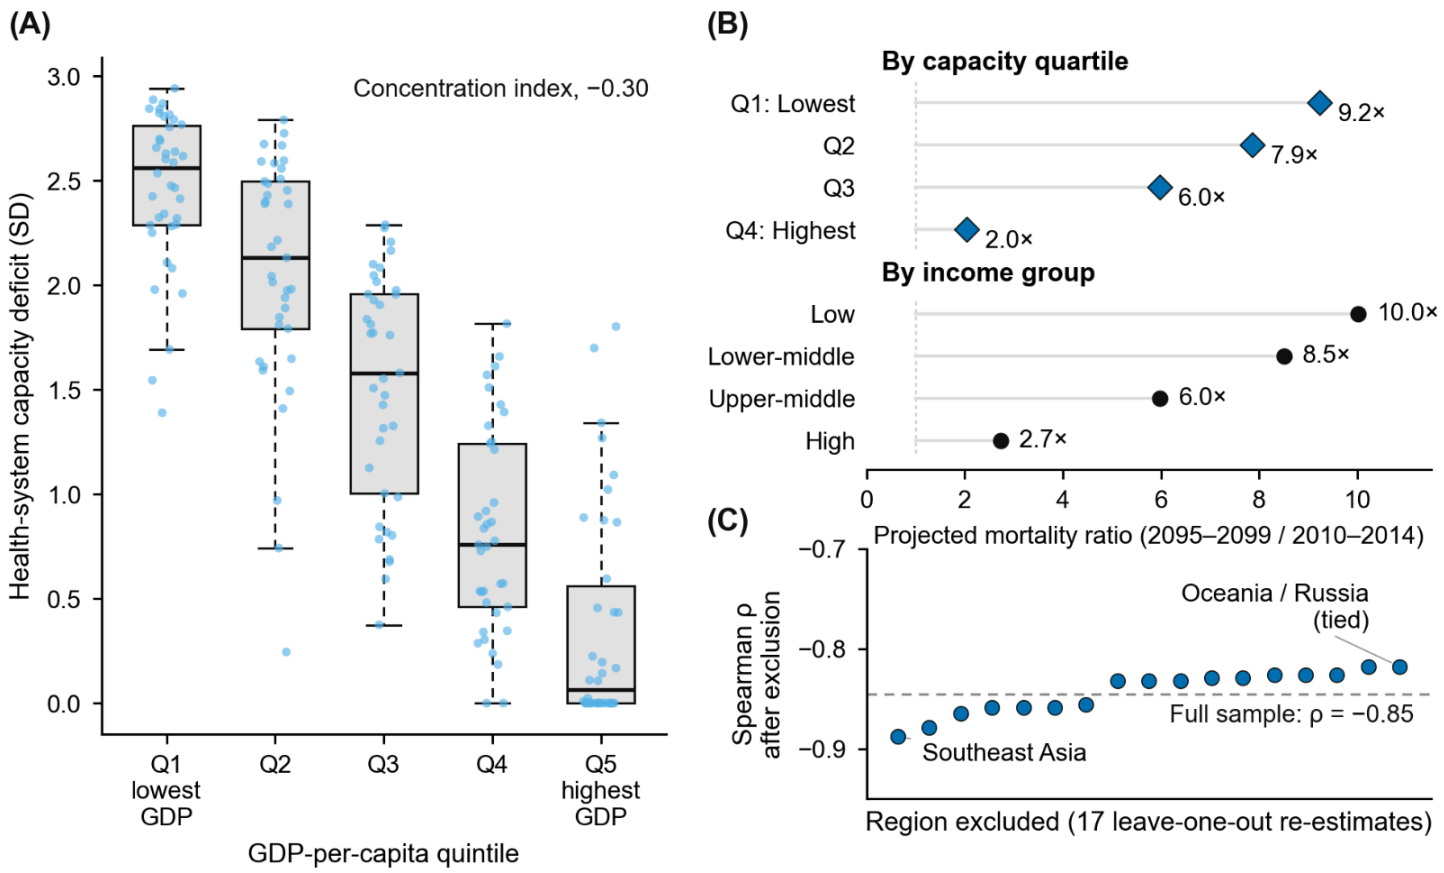


Figure 2

Capacity deficit, projected mortality growth, and leave-one-region-out robustness. (A) Distribution of the health-system capacity deficit (benchmark: 90th percentile of the capacity index) across countries grouped into GDP-per-capita quintiles; boxes show the median and interquartile range, whiskers extend to 1.5× the interquartile range, and small points are individual countries. Poorer quintiles carry larger deficits (concentration index −0.30). (B) Mean projected mortality growth (fold change; dashed line at a ratio of 1, no projected change) by capacity quartile (diamonds) and by income group (circles). (C) Regional Spearman correlation re-estimated 17 times, each omitting one projection region; the dashed line is the full-sample value (ρ = −0.85); points are ordered by the re-estimated correlation and the two extremes are labelled (strongest, Southeast Asia excluded, ρ = −0.89; weakest, Oceania and Russia excluded, tied at ρ = −0.82). All 17 values are given in Additional file 1: Table S3. Panel A boxes summarise the unweighted country distribution and panel B values are descriptive unweighted country means, whereas the concentration index annotated in panel A is population-weighted; countries within the same projection region share the same mortality ratio, so no country-level inferential comparison is intended. Mortality inputs are from Zhao et al. (2025), Source Data Fig. 4a[18]; grouped summaries and correlations were calculated for this study.

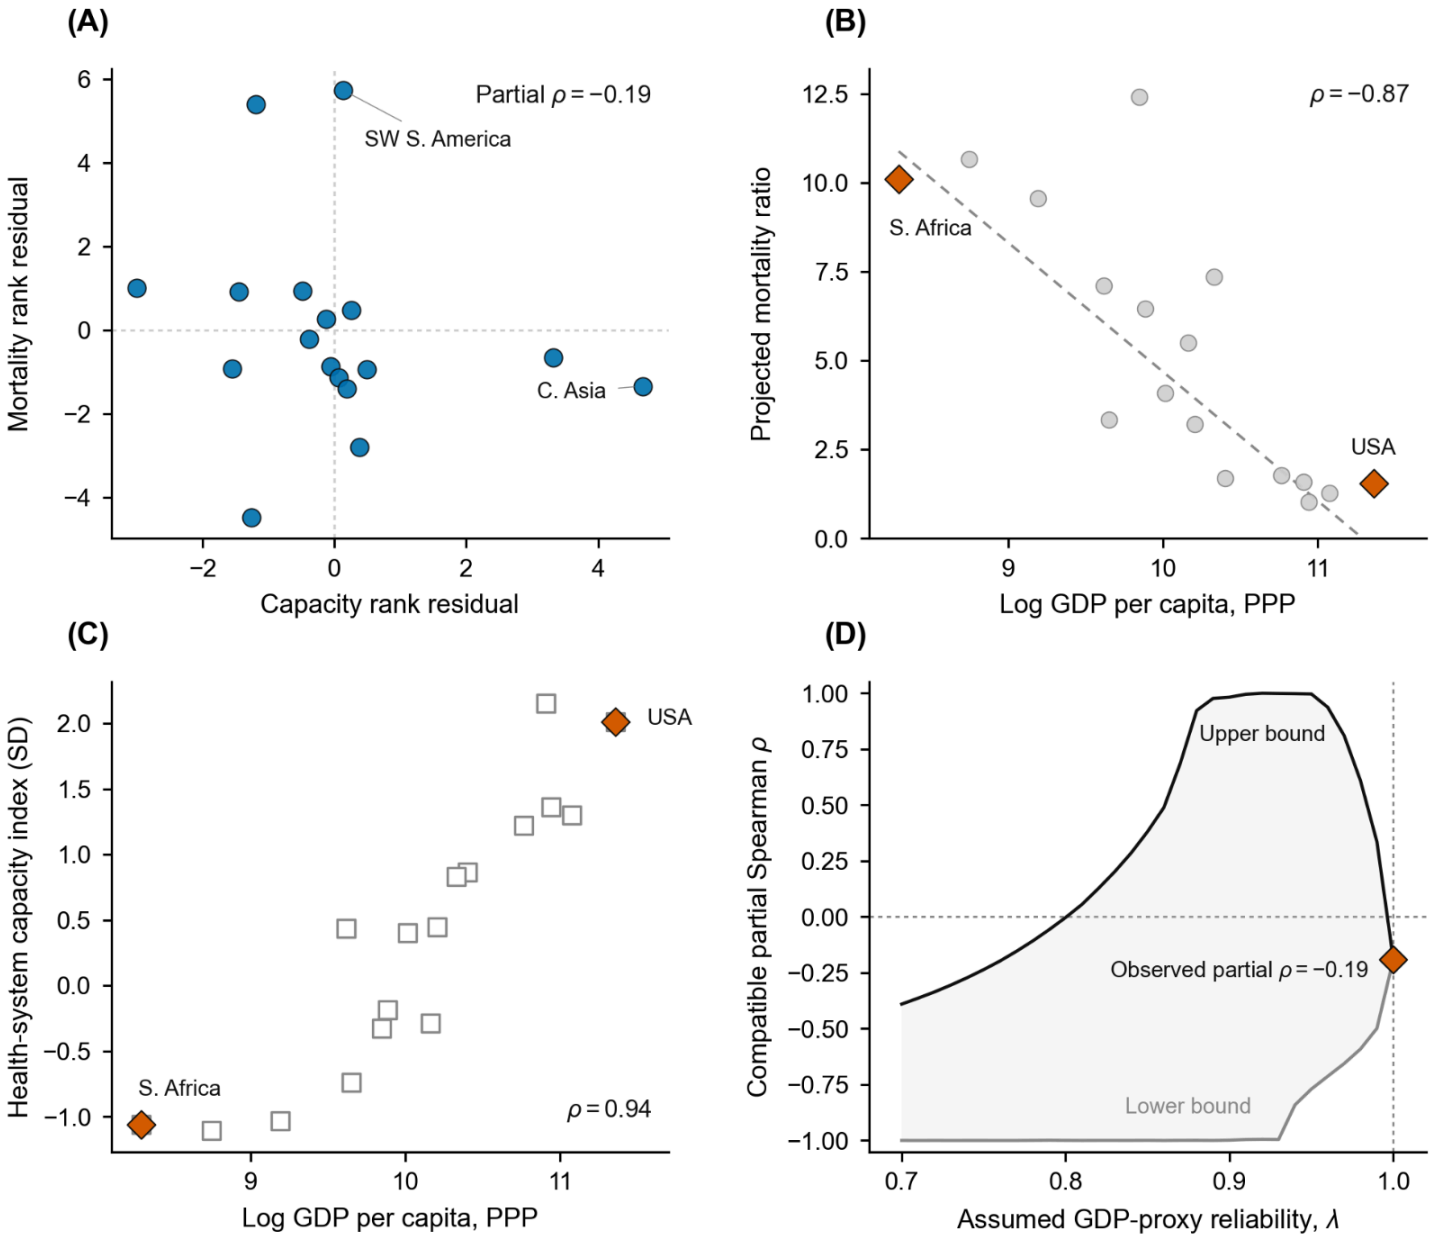


Figure 3

GDP adjustment and partial identification across 17 projection regions. (A) Capacity and mortality-growth rank residuals after adjustment for ranked regional log GDP per capita (partial ρ = −0.19); labelled regions are Southwest South America and Central Asia. (B) Log GDP per capita (PPP international dollars) versus mortality ratio (ρ = −0.87); dashed line: Theil–Sen fit. (C) Log GDP versus capacity (ρ = 0.94). Diamonds in B and C identify Southern Africa and USA. (D) Compatible development-adjusted partial correlations across assumed GDP reliability λ, evaluated in increments of 0.01. Shading shows the bounds; the diamond and vertical line mark λ = 1, and the horizontal line marks zero. The development-factor adjustment is reported in Table S10d of Additional file 1. Mortality ratios are from Zhao et al. (2025), Source Data Fig. 4a[18]; residuals, associations and bounds were calculated for this study.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- Additionalfile16.docx
- Additionalfiles.docx