International Journal of Medical, Healthcare and Life Sciences

IJMHLS

Convergence Without Catch-Up: Population Weighting, Regionalisation and the Determinants of Global Diabetes Growth, 2003–2022

Jeroen Lakerveld [1], **Femke Rutters** [2], **Mirthe Muilwjk** [3*]

[1] Associate Professor, Department of Epidemiology & Data Science, Amsterdam Public Health Research Institute, Amsterdam, Netherlands
[2] Associate Professor, Department of Epidemiology & Data Science, Vrije Universiteit Amsterdam, Amsterdam, Netherlands
[3] Assistant Professor, Department of Epidemiology & Data Science, Amsterdam UMC, Amsterdam, Netherlands

*** Corresponding Author**: me.mwjk@amsterdamumc.nl

Citation: J. Lakerveld, F. Rutters, and M. Muilwjk. "Convergence Without Catch-Up: Population Weighting, Regionalisation and the Determinants of Global Diabetes Growth, 2003–2022". *International Journal of Medical, Healthcare and Life Sciences,* vol.1, No.1, pp.67–84. 2026.

Received: 20 Jan 2026

Accepted: 29 Apr 2026

Abstract

Diabetes prevalence is one of the few major non-communicable disease burdens still rising in most coun- tries. How fast it is rising, and whether the world's countries are converging or diverging, depends on analytic choices that are rarely stated. Using WHO Global Health Observatory estimates for 196 countries over 2003–2022, together with World Bank covariates, we show that the headline number is weighting-dependent in a way that reverses over time. The unweighted country mean rises from 10.01% to 13.80% (+3.79 pp); population weighting raises the 2022 level to 14.61% and raises the estimated an- nual trend by 42%, from +0.194 to +0.275 percentage points per year. The epidemic is also be- coming more regionalised — the between-region share of cross-country variance rose from 0.340 to 0.416 (+0.0052/yr), and the high-to-low regional gap widened from 7.8 to 14.6 percentage points. In a multivariate model of the 20-year change across 179 countries, the dominant correlates are the change in obesity (+0.266 pp per pp, $p < 0.001$), catch-up from a high initial level (-2.13, $p = 0.002$) and baseline health spending (-1.91, $p < 0.001$); the full specification explains 40% of the variance, against 0.1% for initial prevalence alone. The convergence evidence is specification-dependent in an instructive way: countries show significant relative convergence in the Barro–Sala-i-Martin growth form ($\lambda = 1.1$–2.3%/yr, all $p < 0.001$) while showing no absolute convergence and widening absolute dispersion (SD 5.51 to 6.44). A rate-versus-composition decomposition shows that population ageing accounts for 43% of Europe's crude rise but only 6% of Africa's, and the male–female gap reverses sign across regions in a manner that itself depends on weighting. Finally, WHO/NCD-RisC and IDF estimates, though strongly correlated ($r = 0.71$), disagree on the global level by about 3.6 percentage points. We argue that unweighted country-mean summaries of global diabetes, which remain common, systematically understate both the level and the pace of the epidemic and misplace its centre of gravity.

Keywords: Diabetes Prevalence, Global Health, Population Weighting, β-Convergence, Kitagawa Decomposition, Who Global Health Observatory, Measurement Disagreement

INTRODUCTION

Type 2 diabetes has become the clearest example of an epidemiological transition that has not run its course. Age-standardised prevalence among adults has risen on every continent for two decades, and the disease now sits alongside obesity as a defining feature of the global non-communicable disease burden [1,

Copyright © 2026 by Author/s. IJMHLS is an open access article distributed under the Creative Commons Attribution License which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

2]. Yet the quantitative literature on how that rise is distributed is thinner than the literature on its level, for a reason that is largely methodological: global prevalence is almost always summarized as a single number produced by a weighting scheme, and the sensitivity of the answer to that scheme is seldom reported.

Three questions organize in this paper.

First, who counts as the world? Global prevalence can be computed by treating each country equally (an unweighted mean over countries) or by weighting each country by its population. The two answers are usually close enough that the choice appears innocuous. We show that for diabetes over 2003–2022 it is not: the gap between the two estimates is not merely large in 2022, it has the opposite sign in 2003, because the epidemic has migrated toward the most populous countries. Any analysis that mixes the two, or that reports one while interpreting the other, will misstate both the level and the trend.

Second, are countries converging? Convergence is a well-developed concept in growth economics, with an established apparatus: the β-convergence regression of Barro and Sala-i-Martin [4], the distinction between absolute and conditional convergence, the separate notion of σ-convergence, and the recur- ring warning that β - and σ-convergence need not agree. Health economists have applied this apparatus to mortality and to life expectancy, but rarely to diabetes, and rarely with attention to the specification that has become standard in the growth literature. We apply it and find a result that looks contradictory until the specification is read carefully: relative convergence is strong and significant, while absolute convergence is absent. Low-burden countries are growing faster than high-burden countries in proportional terms; the absolute gap in percentage points keeps widening anyway.

Third, what is left of the rise once obesity is accounted for? Obesity is the best-established driver, and it is robustly correlated with diabetes in these data. It is far from sufficient: obesity change explains roughly 18% of the variance in 20-year diabetes change, leaving a substantial component associated with income, health spending, demographic structure and region. Separating the epidemiological component of the rise from the component atributable to population ageing is necessary if the two are to be addressed by different policies.

Our contribution is therefore primarily analytic rather than descriptive. We do not claim new biolog- ical findings. We claim that three measurement choices — weighting, convergence specification, and the rate/composition split — each change the answer materially, and that together they yield a coherent account of diabetes as an epidemic that is regionalising, accelerating, converging in relative terms and diverging in absolute terms. We also quantify a documented disagreement between the two major estimation programmes, WHO/NCD-RisC and the IDF Diabetes Atlas, and show that it is a disagreement about level rather than about direction

DATA

Prevalence

Prevalence estimates come from the WHO Global Health Observatory (GHO) via its OData API[1]. The primary indicator is NCD_DIABETES_PREVALENCE_AGESTD, which reports age-standardised raised fast- ing blood glucose (≥ 7.0 mmol/L) or on-medication prevalence for adults. These estimates are produced for WHO by the NCD Risk Factor Collaboration and are model-based: they pool survey data, correct for measurement diferences, and interpolate over years and countries without direct observation. We use the age-standardised series as the analytic baseline for two reasons. It removes differences in age structure, which vary enormously across the income spectrum and would otherwise be read as differences in risk. And it is the series on which WHO's own headline figures rest, which makes our results directly comparable to published aggregates.

The panel spans 2000–2022 for the 18+ age band, with a second series for 30+; both sexes and sex-specific series are available; and a crude (non-standardised) basis is published alongside. We analyse 2003–2022: 2003 is the first year for which the full covariate set is available, and 2022 is the last year in the GHO

[1] https://ghoapi.azureedge.net/api/.

series. The resulting core slice — age-standardised, both sexes, 18+, 2003–2022 — contains 3,920 country-year observations covering 196 countries. A second, independent prevalence source is the World Bank indicator SH.STA.DIAB.ZS, drawn from the IDF Diabetes Atlas and defined over ages 20–79. It is not a re-publication of the WHO series: the two use different age bands, different diagnostic thresholds and different modelling strategies. We use it as a cross-check, not as a substitute.

Covariates

Country-level covariates come from the World Bank Open Data API: total population (SP.POP.TOTL), GDP per capita in PPP terms (NY.GDP.PCAP.PP.CD), current health expenditure per capita (SH .XPD . CHEX.PC.CD), the urban share of population (SP.URB.TOTL.IN.ZS), and the share of population aged 65 and over (SP.POP.65UP.TO.ZS). Adult obesity (BMI ≥ 30) comes from WHO GHO (NCD_BMI_30C), with the same modelled-estimate caveat as prevalence. Premature non-communicable disease mortality (NCDMORT3070) is used only as a descriptive cross-reference. Every series is joined on ISO 3166-1 alpha-3 codes against the World Bank country roster. World Bank aggregate entities (regions, income groupings, "World") are excluded by requiring a non-empty region field, which only genuine countries carry. All 196 country codes in the WHO panel resolve to World Bank metadata with zero orphans.

Harmonisation

A detail maters more than it appears. WHO GHO returns prevalence in a Value field as a rounded display string of the form "12 .6 [9 .1-16 .4]" and in NumericValue, Low and High fields at full precision. The display string is not the estimate. We parse all four and work from the numeric fields; an audit of all 41,580 rows found zero genuine parse failures and confirmed that the display string is a rounding of NumericValue. Analysts who parse Value introduce quantitation into every downstream calculation, which is harmless for levels and not harmless for trends and second differences. Regions are assigned per WHO's six-region scheme and income groups per the World Bank's four-way classification, both taken from source metadata rather than reconstructed.

METHODS

Weighting

Let pit be age-standardised prevalence in country i in year t. The unweighted global series is $\bar{p}t = \frac{1}{N}\sum i\, pit$ which treats Luxembourg and India as equal observations. The population-weighted series is

$$\bar{p}_t^{w} = \sum_i \omega_{it}\, p_{it}, \quad \omega_{it} = \frac{n_{it}}{\sum_j n_{jt}} \tag{1}$$

with nit the total population. Equation (1) is the quantity that a randomly drawn person on Earth would experience. We report both throughout, and report the difference $\Delta t = \bar{p}tw - \bar{p}t$ as a substantive object rather than as a robustness footnote. Trends are estimated by OLS on the annual series with Newey–West HAC standard errors (5 lags).

The weighting scheme is validated against WHO's published regional aggregates, which are independently produced. Our recomputed weighted regional values track them closely (maximum absolute deviation 0.46 pp across six regions), which we take as evidence that the weighting is faithful and that the residual gap reflects differences in country coverage and in WHO's internal age-standardisation weights.

Regional Comparisons

Cross-regional differences at the end of the window are tested with a Kruskal – Wallis test on 2022 country-level prevalence, with $\varepsilon 2$ as the effect size, and with Levene's median-centred test for equality of variance. Post-hoc pairwise comparisons use Mann–Whitney U with Holm–Bonferroni adjustment over all 15 pairs; effect sizes are reported as rank-biserial correlations. Because the underlying quantities are modelled estimates, we read these tests as statements about the separation of the published estimates, not as inference on sampling error in real populations.

Regional differences in trend are tested with a country-level slope regression and with a linear mixed model containing country random intercepts and random slopes for year, interacting year with region.

Determinants of the 20-Year Change

We regress the 20-year change on initial conditions and concurrent changes:

$$pi, 2022 \; - \; pi, 2003 \; = \; \alpha \; + \; \beta \ln pi, 2003 \; + \; \gamma' \Delta \mathrm{x} i \; + \; \delta' \mathrm{x} i, 2003 \; + \; \varepsilon i \qquad (2)$$

with HC1 heteroskedasticity-robust standard errors. Five nested specifications are estimated: (M1) initial prevalence alone; (M2) adding the change in obesity; (M3) adding changes in urbanisation and the 65+ share; (M4) adding baseline log GDP per capita, log change in GDP per capita, and baseline log health spending per capita; (M5) adding WHO region fixed effects. Multicollinearity is assessed with variance inflation factors.

Two limitations of (2) are worth stating at the outset. It is a cross-section, so it identifies conditional associations rather than causal effects. And the regressors are themselves modeled or estimated quantities, so attenuation and common-source bias are both possible.

Convergence

We estimate β -convergence in the growth form standard in the empirical growth literature [4]:

$$\frac{1}{T} \ln \left(\frac{p_{i,2022}}{p_{i,2003}} \right) = a + b \ln p_{i,2003} + \zeta' \mathbf{x}_i + u_i, \quad T = 19 \qquad (3)$$

with the implied annual speed of convergence

$$\lambda = - \, ln(1 + bT)/T, \qquad (4)$$

defined when $bT \; > \; - 1$. Three specifications are estimated: unconditional, with covariates, and with covariates plus region fixed effects. A negative and significant b indicates catch-up.

We also report the more familiar absolute form, in which the dependent variable is the raw change $\mathrm{pi}, 2022 \; - \mathrm{pi}, 2003$, and the σ - convergence evidence from the cross-country standard deviation and coefficient of variation over the window. The two forms answer different questions and, as we show, give opposite answers here; reporting only one would be misleading in either direction.

Variance Decomposition

To quantify regionalisation we decompose the cross-country variance of prevalence in each year into a between-region and a within-region component:

$$\underbrace{\sum_i (p_{it} - \bar{p}_t)^2}_{total} = \underbrace{\sum_r n_{rt} (\bar{p}_{rt} - \bar{p}_t)^2}_{betweenregions} + \underbrace{\sum_i (p_{it} - \bar{p}_{r(i)t})^2}_{withinregions} \qquad (5)$$

and track the between-region share Bt over the window. A rising Bt means that knowing a country's region tells you more about its diabetes burden than it used to.

Rate Versus Composition

The crude prevalence rate rises both because age-specific risk rises and because populations age into higher-risk ages. Because only two overlapping age bands (18+ and 30+) are published, we cannot per- form a full Kitagawa decomposition [7] with age-specific rates. We instead decompose the crude change into the change in the age-standardised rate and the residual of set, which absorbs the age-structure term:

$$\underbrace{p_{2022}^{crude} - p_{2003}^{crude}}_{total} = \underbrace{\Delta p^{age-std}}_{epidemiologic} + \underbrace{\Delta(p^{crude} - p^{age-std})}_{agestructure} . \quad (6)$$

This is exact by construction. It identifies the compositional component only in aggregate, not by age band, and we present it as such. The decomposition is computed both as an unweighted country mean, where each country counts once, and for the world using WHO's published population-weighted aggregates.

Robustness

The regional ordering is re-estimated under five variants: the 30+ age band instead of 18+; the crude basis instead of age-standardised; female-only prevalence; restricting to countries above one million population; and trimming the top and bottom 2.5% of each region's distribution. Concordance with the baseline ordering is measured by Spearman rank correlation. The global 20-year change is also reported under each of the three weighting/basis combinations.

RESULTS

The Magnitude of the Rise Depends on Who Counts

Age-standardised prevalence among adults rose over the 20 years, but by how much depends on the weighting. Averaged across countries, it went from 10.01% in 2003 to 13.80% in 2022, a rise of 3.79 percentage points. Weighted by population, it went from 9.20% to 14.61%, a rise of 5.41 percentage points. The ratio between the two is substantial: the weighted trend is +0.275 pp/yr against +0.194 pp/yr unweighted, a factor of 1.42 (both $p < 0.001$, HAC standard errors).

Table 1. Global age-standardised diabetes prevalence, adult (18+) population, under alternative weighting and basis. WHO-published figures are that organisation's own population-weighted aggregates and are shown for reference, not as an input.

Series	2003 (%)	2022 (%)	Change (pp)	Trend (pp/yr)
Unweighted country mean	10.01	13.80	+3.79	+0.194
Population-weighted	**9.20**	**14.61**	**+5.41**	**+0.275**
Crude, unweighted	8.89	13.93	+5.04	—
WHO published, age-standardised	8.73	14.15	+5.41	+0.275
WHO published, crude	8.19	14.78	+6.59	—

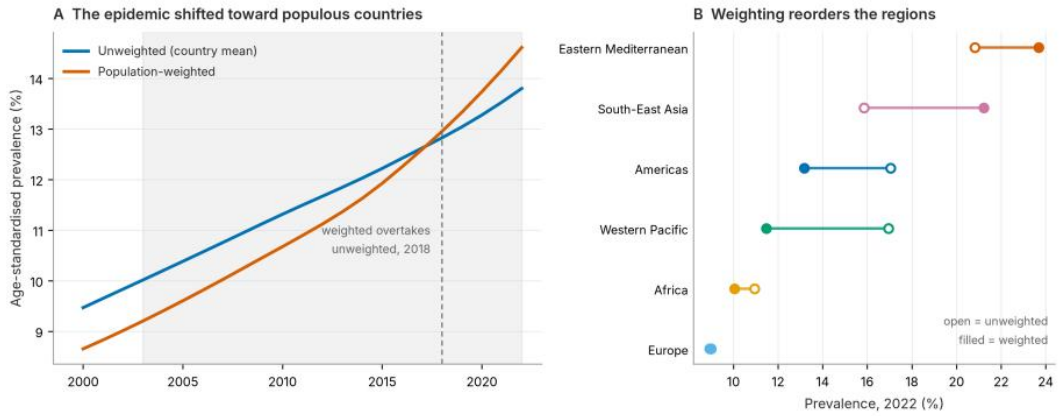


Figure 1. Weighting changes both the level and the trend. (A) Unweighted country mean versus population-weighted global prevalence, age-standardised, adults 18+. The series cross in 2018; the shaded band marks the 2003–2022 analytic window. Both series begin in 2000, the first year of the 18+ GHO series. (B) Regional prevalence in 2022, unweighted (open) versus population-weighted (filled). South-East Asia rises from fifth to second place under weighting.

Table 1 makes the structure of the discrepancy visible. The population-weighted change of +5.41 pp matches WHO's own published age-standardised change to two decimal places, which both validates the weighting implementation and confirms that WHO's headline figure is a population-weighted quantity. The unweighted series, which is what an analyst obtains by taking a simple mean across the WHO country rows, understates the change by 30%.

The reason the two diverge is not merely arithmetic. Figure 1A plots both series and shows that they cross. In 2003 the weighted series sits 0.81 pp below the unweighted one; by 2022 it sits 0.81 pp above; the crossing occurs in 2018. In words: twenty years ago, the world's most populous countries had somewhat lower diabetes prevalence than a typical country. Today they have higher. This sign reversal means the weighting bias does not merely shift the level — it changes the estimated pace of the epidemic, because the two series have different slopes, and it relocates the epidemic's centre of gravity toward South and East Asia.

Figure 1B shows the same defect on the regional ranking. Population weighting lifts South-East Asia from 15.86% to 21.23% — from fourth to second — while depressing the Western Pacific from 16.96% to 11.48%, because that region's unweighted mean is dragged upward by small Pacific island states with very high prevalence. Rank correlation between the weighted and unweighted regional orderings is 0.829: high, but not high enough to treat the two as interchangeable.

Table 2. Regional prevalence in 2022 (age-standardised, 18+), unweighted and population-weighted, with WHO's published aggregates for reference. Sorted by weighted prevalence.

WHO Region	n	Unweighted (%)	Weighted (%)	WHO Published (%)
Eastern Mediterranean	22	20.82	23.68	23.25
South-East Asia	10	15.86	21.23	21.12
Americas	37	17.04	13.17	13.06
Western Pacific	28	16.96	11.48	11.29
Africa	47	10.95	10.06	10.52
Europe	52	9.01	8.94	8.68

A formal test confirms the acceleration. Fitting a change-point model to the global series, the annual trend is +0.205 pp/yr over 2003–2012 and +0.371 pp/yr over 2013–2022; the difference of +0.167 pp/yr is significant ($p < 0.001$). The epidemic is not decelerating.

Regions have diverged, not converged

In 2022 the six WHO regions differ substantially and significantly (Kruskal–Wallis $H = 83.1$, $p < 0.001$, $\varepsilon 2 = 0.41$; Levene $W = 6.4$, $p < 0.001$). The effect size is large — region alone accounts for roughly 41% of the rank variance in country-level prevalence. Of 15 pairwise comparisons, 8 remain significant after Holm–Bonferroni adjustment. Table 2 gives the regional levels under both weightings.

The high-to-low regional gap widened from 7.8 pp in 2003 to 14.6 pp in 2022, very nearly doubling. This is the headline divergence result, and it survives weighting: the widening is visible in both the weighted and unweighted tables.

The trend evidence points the same way. Country-level slopes average +0.194 pp/yr (SD 0.208), with 85.7% of countries rising; slopes differ significantly across regions ($H = 39.4$, $p < 0.001$). In the mixed model, relative to the African Region, the Americas rises 0.177 pp/yr faster ($p < 0.001$), South-East Asia 0.194 faster ($p = 0.004$), and the Eastern Mediterranean 0.134 faster ($p = 0.007$); Europe and the Western Pacific are statistically indistinguishable from Africa. The African Region is not merely at a lower level — it is on a flatter trajectory. Whether that reflects a genuinely slower epidemic or a slower epidemiological transition yet to come is a question the data cannot settle [6].

The Epidemic Became More Regionalised

Figure 2B and Table 3 report the variance decomposition. The between-region share of total cross-country variance rose from 0.340 in 2003 to 0.416 in 2022, a trend of +0.0052 per year. Over the same period the within-region sum of squares rose from 3,905 to 4,727, so this is not a case of within-region variance

collapsing — both components grew, but the between-region component grew faster. The substantive reading is that geography is becoming a better predictor of diabetes burden, not a worse one. A world in which the epidemic were diffusing uniformly would show a falling between-region share. The opposite is happening: the disease is concentrating in identifiable regional blocks, which maters for how surveillance and treatment capacity should be allocated.

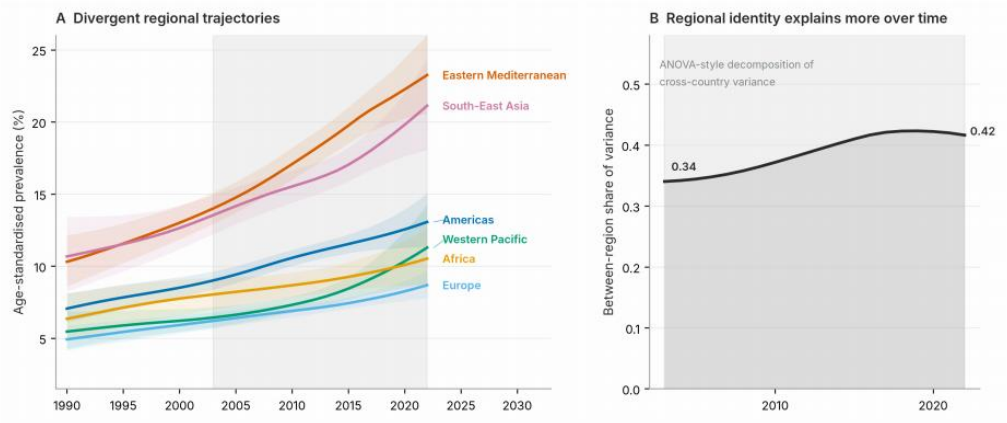


Figure 2. Regional trajectories and regionalisation. (A) WHO regional aggregates, age-standardised adults 18+, with published uncertainty intervals. (B) The between-region share of cross-country variance rose from 0.34 to 0.42 over the window, at +0.0052 per year.

Table 3. Decomposition of cross-country variance in age-standardised prevalence (196 countries)

Year	Between-region share	Within-region SS	SD (pp)	Mean (%)
2003	0.340	3,905	5.51	10.01
2010	0.372	4,453	6.03	11.31
2015	0.410	4,419	6.20	12.22
2019	0.423	4,455	6.29	13.04
2022	0.416	4,727	6.44	13.80

Obesity is Central but Far From Sufficient

The nested regressions in Table 4 tell a clean story about explanatory power. Initial prevalence alone explains essentially nothing about subsequent growth ($R^2 = 0.001$). Adding the change in obesity takes the model to $R^2 = 0.187$ — most of the explanatory work is done by a single variable. Urbanisation and ageing barely move it (0.196). Income and health spending lift it to 0.346, and region fixed effects to 0.404. The full model is reported below.

Figure 3 plots the M5 coefficients; four of them warrant comment.

Catch-up. The coefficient on log initial prevalence is __2.13 and significant ($p = 0.002$). Two countries differing by a factor of e in initial prevalence (2.7x) difer by 2.1 pp in their subsequent 20-year change, with the higher-burden country rising less in absolute terms — once con-founders are held constant. This is the conditional-convergence signal, and it is invisible in the unconditional model.

Obesity change. Each additional percentage point of adult obesity gained over the period is associated with +0.27 pp of diabetes prevalence ($p < 0.001$). This is the single most robust association in the model and the one with the clearest policy interpretation. It is also modest in magnitude: the observed inter quartile range of obesity change (6.3 pp) translates to about 1.7 pp of diabetes change, against a mean observed rise of 3.8 pp.

Income and spending. Higher baseline GDP per capita predicts faster diabetes growth, while higher baseline health spending per capita predicts slower growth. These two variables are correlated (VIF 6.2 and

6.5 respectively — the highest in the model, though below the conventional threshold of 10), so their separate coefficients should be read with caution. The pattern is nonetheless coherent with the broader literature: it is middle-income countries, not the poorest, that are converting income growth into diabetes burden, and health-system investment appears to attenuate that conversion.

Null results. Urbanisation change is not significant at any specification (p = 0.93 in M5), and the change in GDP per capita is not significant (p = 0.55). Urbanisation is so nearly collinear with development in the cross-section that its independent contribution cannot be identified here; we do not read the null as evidence that urbanisation is unimportant.

Table 4. Nested regressions of the 20-year change in age-standardised prevalence on initial conditions and concurrent changes, 179 countries with complete data. HC1 robust standard errors in parentheses.

	M1	M2	M3	M4	M5
log initial prevalence	0.280	-0.889	-0.844	-1.249*	-2.128**
	(0.749)	(0.678)	(0.716)	(0.615)	(0.678)
Δ obesity (pp)		0.401***	0.416***	0.355***	0.266***
		(0.066)	(0.067)	(0.058)	(0.063)
Δ urbanisation (pp)			0.058	0.006	0.004
			(0.039)	(0.040)	(0.041)
Δ population 65+ (pp)			-0.003	0.379**	0.265*
				(0.142)	(0.129)
log GDP pc, 2003			(0.106)	1.883***	1.585**
				(0.548)	(0.562)
Δ log GDP pc				-0.662	-0.600
				(0.801)	(1.013)
log health spending pc, 2003				-2.081***	-1.911***
				(0.361)	(0.348)
Region fixed effects	No	No	No	No	Yes
n	179	179	179	179	179
R2	0.001	0.187	0.196	0.346	0.404
Adjusted R^2	-0.004	0.178	0.177	0.319	0.361

Table 5. β-convergence under alternative specifications. b is the coefficient on log initial prevalence; λ is the implied annual speed of convergence, defined only for bT > -1 (T = 19). A negative b indicates catch-up.

Specification	b	SE	p	λ (%/yr)	R^2
Relative form: dependent variable ln(p2022 /p2003)					
Unconditional	−0.0097	0.0026	< 0.001	1.08	0.084
With covariates	−0.0146	0.0024	< 0.001	1.71	0.414
With region fixed effects	−0.0185	0.0028	< 0.001	2.28	0.452
Absolute form: dependent variable $P_{2022} - P_{2003}$					
Unconditional	+0.2800	0.7494	0.709	—	0.001
With region fixed effects	−2.1275	0.6779	0.002	—	0.404

Convergence in Relative Terms Only

This is the result most sensitive to specification, and the most instructive. Table 5 reports the convergence evidence under three specifications. In the relative (growth) form, all three specifications show significant convergence, with speeds rising from 1.08%/yr unconditional to 2.28%/yr once covariates and region fixed effects are included. These are of the same order as the ≈2%/yr modal speed reported in the growth literature [4], with the important caveat that a 19-year window over 196 units is a small sample and the estimates should be read as orders of magnitude.

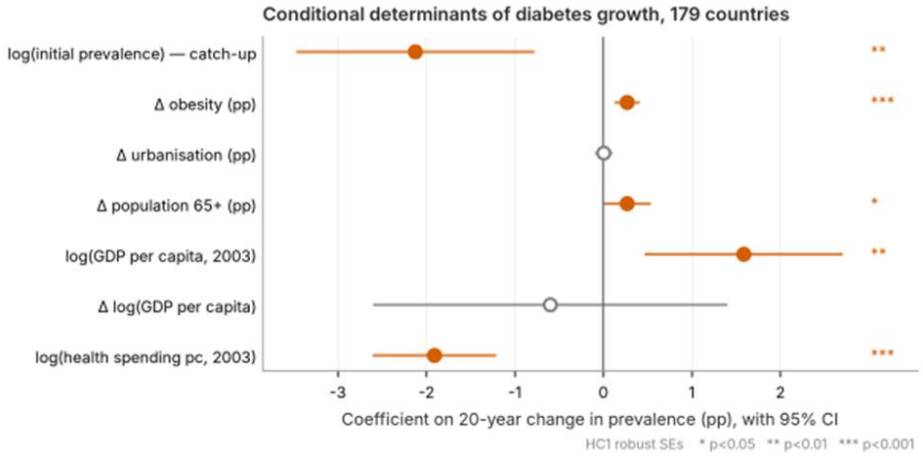


Figure 3. Conditional determinants of 20-year diabetes growth. Coefficients from specification M5 (Table 4) with 95% confidence intervals. Filled markers denote $p < 0.05$; region fixed defects are included but not shown.

In the absolute form, the unconditional specification shows no convergence: b = +0.280, p = 0.709, R^2 = 0.001. This is not a contradiction of the relative result; the two forms ask different questions. The relative form asks whether low-burden countries are growing proportionally faster. The absolute form asks whether the gap in percentage points is closing. Figure 5A resolves the apparent paradox: the cross-country standard deviation rose 17% over the window (5.51 to 6.44 pp, +17%) while the coefficient of variation fell 15% (0.550 to 0.467). Countries are moving toward each other in proportional terms and away from each other in absolute terms simultaneously. This is precisely the β-versus-σ-convergence distinction that the growth literature has long emphasised, and it means that a paper reporting either result alone would be reporting a half-truth.

For policy the distinction has teeth. Relative convergence is good news about the rate of change in the countries that were spared the worst of the epidemic, and it is the kind of convergence that health-system learning and tobacco and diet transition would be expected to produce. Absolute divergence is bad news about the stock of disease, and it is what determines how many people need care in 2030. Both are true. The second is the more actionable.

What Ageing Explains, and Where

Table 6 and Figure 6 split each region's crude rise into an epidemiological component (the age-standardised rate) and a compositional one (the crude-minus-standardised ofset). Globally, counting each country once, 75% of the crude rise is epidemiological and 25% is compositional. For the world population-weighted, WHO's published aggregates give 82% epidemiological and 18% compositional.

The regional spread is the finding. In the African Region, ageing accounts for only 6% of the rise — HIV- era mortality, conflict and high fertility keep the age distribution young, so almost all of the increase is genuine risk. In Europe, 43% of the crude rise is attributable to population ageing rather than to rising age-specific risk. The same crude number means very different things in the two stings, and a crude-rate league table that puts Europe and Africa close together conceals this entirely.

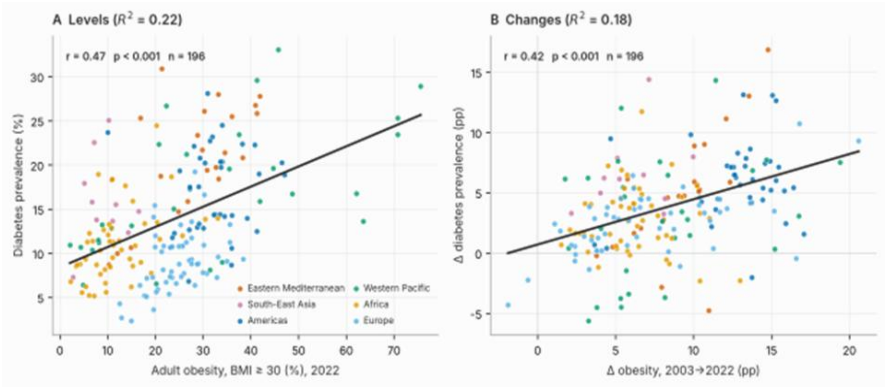


Figure 4. Obesity and diabetes, in levels and in changes. (A) Cross-sectional association in 2022. (B) Association between the 20-year change in obesity and the 20-year change in diabetes. The change-on-change relationship (r = 0.42) is nearly as strong as the level relationship (r = 0.47), which is notable given that the change measure carries two periods of measurement error. Figure 4 shows both.

Table 6. Decomposition of the crude 20-year change by region, unweighted country mean. The epidemiological share is the fraction of the crude rise attributable to the age-standardised rate rather than to population age structure.

Scope	Δ Crude (pp)	Δ Age-std (pp)	Δ Age Structure (pp)	Epidemiological Share
Global	5.04	3.79	1.25	75%
Africa	2.95	2.77	0.19	94%
Eastern Mediterranean	5.78	5.03	0.74	87%
South-East Asia	7.50	6.52	0.97	87%
Americas	7.80	6.03	1.78	77%
Western Pacific	4.68	2.97	1.71	63%
Europe	4.37	2.50	1.87	57%
World (pop-weighted, WHO)	6.59	5.41	1.18	82%

The Sex Gap Reverses — But Where Depends on Weighting

Unweighted, women have higher age-standardised prevalence than men in five of the six regions, and the global female trend is marginally steeper (+0.199 vs +0.187 pp/yr). Regionally, the unweighted gap runs from −2.22 pp in South-East Asia to +2.08 pp in Europe.

Population weighting changes the picture. Weighted, the gap is −2.38 pp in South-East Asia and −1.28 pp in Africa — both remain strongly female-predominant — but the Americas move from −1.97 to +0.09 and the Western Pacific from −0.67 to +1.98. In both cases the unweighted regional mean is driven by small states whose sex gradient difers from that of the region's large populations.

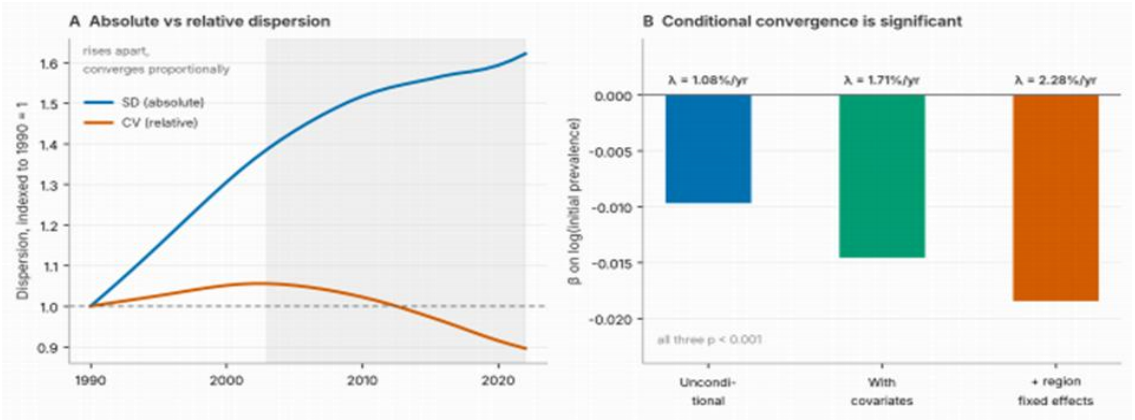


Figure 5. Relative versus absolute convergence. (A) Cross-country dispersion of age-standardised prevalence, indexed to 1990 = 1. The standard deviation rises while the coefficient of variation falls. (B) Implied speed of convergence λ under the three specifications of

Table 7. Robustness of the regional ordering and of the global change. Spearman rank correlation is between the variant's regional medians and the baseline's.

Variant	n	Rank Correlation vs Baseline
Age band 30+ (vs 18+)	196	0.943
Crude basis (vs age-standardised)	196	0.943
Female only (vs both sexes)	196	0.943
Trimmed 2.5% tails	180	1.000
Population ≥ 1 million only	158	0.829

Figure 7 plots both versions. We report this as a caution rather than a finding. The claim "women bear a higher diabetes burden than men in most world regions" is well supported unweighted and becomes region-dependent weighted. Any paper making the claim should say which population it is talking about. What is robust is the direction of the global trend: both sexes are rising at close to 0.2 pp/yr, and the sex differential is small relative to the regional differential.

Two Estimation Programmes Disagree about the Level

Whoever uses these data will eventually compare the WHO series with the IDF series carried by the World Bank. Figure 8 shows the comparison for the countries where both are available. The two are strongly correlated across countries (r = 0.71, n = 195) and agree on direction: both show a global rise. They disagree about the level by about 3.6 percentage points on average (IDF lower), with a mean absolute difference of 4.68 pp and country-level differences ranging beyond ±10 pp.

The disagreement is not an error in either series. The definitions fifer: different age bands, different diagnostic criteria for the underlying surveys, different modelling assumptions, different treatment of undiagnosed cases and different publication vintages. But the practical consequence is that a paper reporting "global diabetes prevalence" without naming its source is reporting a number that varies by roughly a quarter of its own value depending on an unstated choice. Our analysis uses WHO throughout; a replication on the IDF series would be a worthwhile check of whether the convergence and regionalisation results survive.

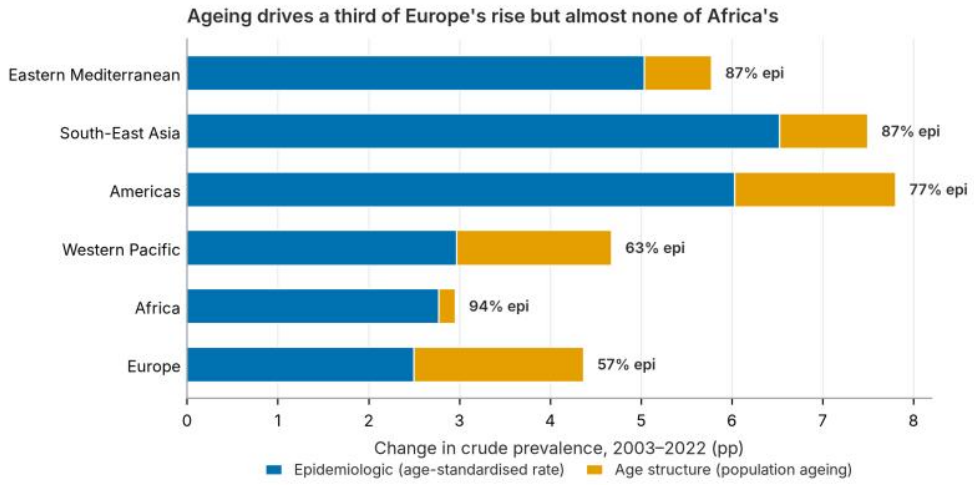


Figure 6. Rate versus composition by region. Each bar is the region's crude 20-year change, split into the age-standardised rate component (blue) and the age-structure component (orange). The epidemiologic share ranges from 57% in Europe to 94% in Africa.

Robustness

Table 7 reports the robustness grid. The regional ordering is stable under four of the five variants (p ≥ 0.94), and identical under trimming. It degrades to 0.829 when small countries are excluded, which is the

same phenomenon that produced the weighting reversal in Section 4.1: the small-state component of the unweighted mean is what makes the unweighted and weighted orderings differ, and removing it moves the unweighted ordering partway toward the weighted one.

The global change is likewise bounded: +3.79 pp age-standardised unweighted, +5.04 pp crude un weighted, +5.41 pp age-standardised population-weighted. Every specification gives a substantial rise; none gives a decline. As an external cross-reference we checked premature NCD mortality (NCDMORT3070) against diabetes change. The association is null (r = -0.064, p = 0.384, n = 185). We read this as uninformative rather than as contradicting the prevalence results: the indicator aggregates all four major NCDs, falls almost everywhere for reasons unrelated to diabetes, and would not be expected to track a single risk factor at country level over 20 years.

DISCUSSION

Why the Weighting Result Matters

The population-weighted and unweighted global series cross in 2018. Before that crossing, the world's most populous countries had below-average diabetes prevalence; after it, above. The mechanism is not mysterious — India, China, Pakistan, Indonesia and Bangladesh have all seen rapid rises, and together with a handful of other large states they now dominate the world total. But the consequence is that the analytic choice which looks least consequential, whether to use a simple mean or a weighted one, is in fact the one that determines whether the estimated trend is +0.19 or +0.28 pp/yr, and whether the epidemic looks like it is centred on small island states or on South and East Asia.

This is not merely a presentational point. The unweighted mean is dominated by the many small countries in the African and European regions, and it therefore gives a systematically different picture of where the burden is. Interventions are planned at the level of populations, not countries; so is the case for financing. The weighted series is the one that corresponds to people.

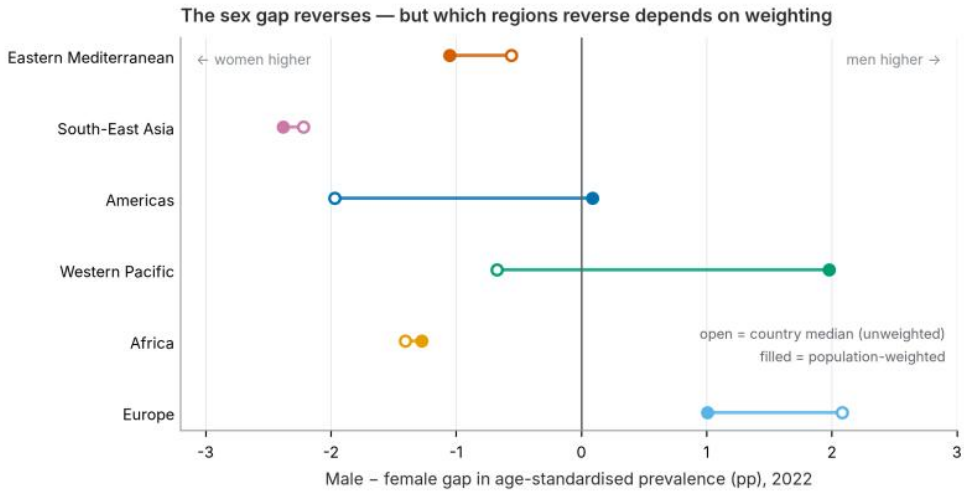


Figure 7. Male minus female gap in age-standardised prevalence, 2022. Open markers are unweighted country medians; filled markers are population-weighted regional means. The Americas and the Western Pacific change sign between the two.

Regionalisation as the Organising Fact

The between-region share of variance rose from 0.34 to 0.42 over twenty years. This is the single most robust structural finding in our analysis, and it is not obvious in advance. Globalisation of diets and the spread of sedentary work might be expected to homogenise risk; instead, the burden is concentrating into identifiable regional blocks. The Eastern Mediterranean and South-East Asia are pulling away from the pack,

while Europe and Africa remain relatively low (for diferent reasons — ageing on one side, a young population on the other).

Convergence and regionalisation coexist without tension. Countries are converging relative to one another within the global distribution, while that distribution is becoming more strongly structured by re- gion. The λ estimates of 1.1 – 2.3%/yr describe catch-up in proportional terms; the rising between-region share describes the geometry of the distribution those countries are converging within.

Obesity, and What the Residual Contains

Obesity change is the strongest single correlate of diabetes growth in these data, and the change-on-change association (r = 0.42) is nearly as strong as the level association (r = 0.47) despite carrying two periods of measurement error in both variables. Yet the full model explains 40% of the variance, which leaves 60% unexplained; adding obesity change lifts it to 19%.

The residual is not noise. It contains dietary composition independently of calorie balance, physical activity, the distribution of obesity within the population rather than its mean, the duration of expo- sure, genetic and epigenetic variation, diagnostic capacity and screening intensity, and the diabetes- specific consequences of treatments for other conditions. The atributable-risk decompositions of [9] work through several of these channels quantitatively. Our covariates capture none of them, and a model of diabetes growth that treats obesity as the sole explanatory variable is not wrong so much as incomplete — with an incompleteness that is systematic rather than random.

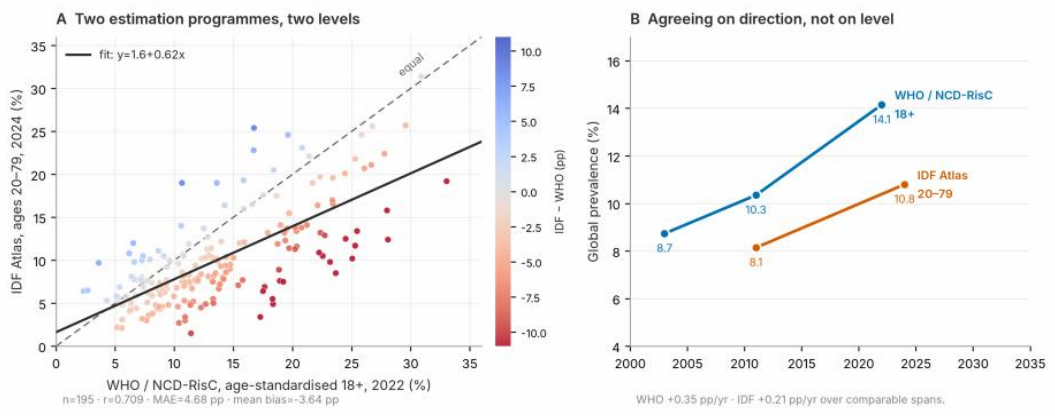


Figure 8. WHO/NCD-RisC and IDF prevalence estimates compared. (A) Country-level comparison, WHO 2022 (age-standardised 18+) against IDF 2024 (ages 20–79); the diagonal marks equality. (B) Global series from each programme over their respective spans.

A Measurement Caution about WHO and IDF

The 3.6 pp average gap between the WHO/NCD-RisC and IDF series is large enough to change policy conclusions. It is also, in our reading, healthy: two independent estimation programmes using different data and different modelling strategies, converging on the same direction and roughly the same country ordering, is stronger evidence than either alone. The problem arises only when a single number is re- ported without attribution. The seven-year vintage difference between the two series (WHO 2022, IDF 2024) compounds this, and we have not attempted to correct for it.

Implications

Three follow for practice.

First, report the weighting. Global prevalence should be stated as a population-weighted quantity or explicitly labeled as a country mean. Where both are available, both should be shown, because the difference is itself informative.

Second, separate rate from composition before comparing regions. Europe and Africa have crude rates that are closer than their epidemiological risk profiles are, in opposite directions. Age-standardisation is not a technicality in this sting; it is the difference between a statement about health and a statement about demography.

Third, treat convergence claims as specification-dependent. Our data support relative convergence and reject absolute convergence. A paper reporting only one of these would be defensible on its own terms and jointly misleading. The growth literature settled this question decades ago; the health literature has been slower to adopt its conventions.

LIMITATIONS

These estimates are modeled, not measured. WHO GHO country-year values are outputs of NCD-RisC's Bayesian hierarchical model, which borrows strength across countries and years. Countries with sparse survey data have estimates that are substantially informed by their neighbour. The consequence for our analysis is direct and important: the cross-country variance we decompose in Section 4.3 is partly model variance, and the smoothness of the annual series — which is what allows R^2 values near 0.999 in trend fits — reflects the model's temporal interpolation rather than 19 independent annual observations. Significance tests are tests on the published estimates.

Relatedly, WHO's published 95% uncertainty intervals are model uncertainty and are not propagated into any of the tests reported here. The regional separations we document are large relative to those intervals in most cases, but not in all.

The 20-year window is short for convergence analysis, and $n = 196$ is small by the standards of the growth literature, where convergence studies typically use panels of hundreds of regions over four or more decades. Our λ estimates should be read as indicative. We also estimate a single cross-section rather than a panel, so within-country dynamics and unobserved heterogeneity are not separable.

The determinant regressions are associations. Obesity, income and diabetes are jointly determined; reverse causality from diabetes to obesity is implausible but reverse causality from diabetes burden to health spending is not, and would bias the health spending coefficient toward the negative values we observe. Country fixed defects cannot be applied in a two-period cross-section, and our region fixed defects absorb variation that may itself be causal (regional food environments, for instance).

The rate/composition decomposition in Section 4.6 is exact only in the aggregate. A full Kitagawa decomposition requires age-specific rates, which GHO does not publish at the resolution needed; with two overlapping age bands we can identify the total compositional term but not its age distribution. The 43%-versus-6% Europe–Africa contrast is robust, but finer regional attribution should not be read from it.

Finally, the IDF cross-check uses a 2024 vintage against a 2022 vintage without adjustment, so the 3.6 pp mean gap conflates definition differences with a two-year difference in period. We also do not include a formal test of whether the convergence and regionalisation results replicate on the IDF series; that is the natural extension of this work

CONCLUSION

Over 2003 – 2022, global age-standardised diabetes prevalence among adults rose by 5.4 percentage points on a population-weighted basis and by 3.8 points on an unweighted one. The epidemic accelerated rather than plateaued, with the annual trend roughly doubling after 2013. It became more regionalised: the between-region share of cross-country variance rose from 0.34 to 0.42, and the gap between the highest- and lowest-burden WHO regions widened from 7.8 to 14.6 percentage points.

Two results cut against simple narratives. Convergence is real but relative: low-burden countries are

catching up in proportional terms at 1.1–2.3% per year, while the absolute gap in percentage points continues to widen and cross-country dispersion increases. And the demographic contribution varies by a factor of seven across regions — 6% of Africa's crude rise against 43% of Europe's — so that the crude-rate comparison that appears most natural is also the most misleading.

The analytic lesson is that three choices rarely reported together determine the answer: whether countries or people are counted, which convergence specification is estimated, and whether rates are separated from age composition. None is exotic. Each changes the substantive conclusion. We suggest that global diabetes reporting state all three.

REFERENCES

[1] NCD Risk Factor Collaboration, “Worldwide trends in diabetes prevalence and treatment from 1990 to 2022: A pooled analysis of 1108 population-representative studies with 141 million participants,” *The Lancet*, vol. 403, no. 10438, pp. 1507–1520, 2024.

[2] World Health Organization, “Diabetes: Fact sheet,” Geneva, Switzerland: World Health Organization, 2024. [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/diabetes

[3] International Diabetes Federation, *IDF Diabetes Atlas*, 11th ed. Brussels, Belgium: International Diabetes Federation, 2025.

[4] X. Sala-i-Martin, “The classical approach to convergence analysis,” *The Economic Journal*, vol. 106, no. 437, pp. 1019–1036, 1996.

[5] R. J. Barro and X. Sala-i-Martin, “Convergence,” *Journal of Political Economy*, vol. 100, no. 2, pp. 223–251, 1992.

[6] A. R. Omran, “The epidemiologic transition: A theory of the epidemiology of population change,” *The Milbank Memorial Fund Quarterly*, vol. 49, no. 4, pp. 509–538, 1971.

[7] E. M. Kitagawa, “Components of a difference between two rates,” *Journal of the American Statistical Association*, vol. 50, no. 272, pp. 1168–1194, 1955.

[8] L. M. Jaacks, S. Vandevijvere, K. D. Pan, et al., “The obesity transition: Stages of the global epidemic,” *The Lancet Diabetes & Endocrinology*, vol. 7, no. 3, pp. 231–240, 2019.

[9] B. Zhou, R. A. Benton, J. E. Bennett, et al., “Global, regional, and national burden of type 2 diabetes and its attributable risk factors, 1990–2021,” *The Lancet Diabetes & Endocrinology*, 2023.

APPENDIX

A.1 Variable definitions

Table 8 lists every variable used and the indicator code it comes from.

Table 8. Variable definitions and source indicators.

Variable	Source	Definition
Prevalence, age-standardised	WHO GHO	NCD_DIABETES_PREVALENCE_AGESTD; raised fasting glucose ≥7.0 mmol/L or on medication, adults 18+, age-standardised
Prevalence, crude	WHO GHO	NCD_DIABETES_PREVALENCE_CRUDE; same definition, not age-standardised
Prevalence, IDF	World Bank	SH.STA.DIAB.ZS; IDF Atlas estimates, ages 20–79
Obesity	WHO GHO	NCD_BMI_30C; BMI ≥30, adults 18+, age- standardised
Population	World Bank	SP.POP.TOTL
GDP per capita	World Bank	NY.GDP.PCAP.PP.CD, PPP, constant 2021 in- ternational $
Health spending	World Bank	SH.XPD.CHEX.PC.CD, current US$ per capita
Urbanisation	World Bank	SP.URB.TOTL.IN.ZS, % of total population
Population 65+	World Bank	SP.POP.65UP.TO.ZS, % of total population
NCD mortality	WHO GHO	NCDMORT3070; probability of dying aged 30– 70 from one of four major NCDs

A.2 Full Specification M5

Table 9 reports all coefficients of the preferred specification, with variance inflation factors alongside.

Table 9. Full specification with region fixed efects (M5), n = 179, R2 = 0.404. Africa is the reference region. VIF is the variance inflation factor of the corresponding regressor; values below 10 are conventionally treated as acceptable.

Term	Coefficient	SE	p	VIF
Intercept	0.339	4.625	0.942	—
log initial prevalence	−2.128	0.678	0.002	1.43
Δ obesity (pp)	0.266	0.063	< 0.001	1.23
Δ urbanisation (pp)	0.004	0.041	0.932	1.18
Δ population 65+ (pp)	0.265	0.129	0.040	1.99
log GDP pc, 2003	1.585	0.562	0.005	6.23
Δ log GDP pc	−0.600	1.013	0.554	1.62
log health spending pc, 2003	−1.911	0.348	< 0.001	6.50
Eastern Mediterranean (vs Africa)	2.836	1.374	0.039	—
Europe (vs Africa)	0.648	1.064	0.542	—
Americas (vs Africa)	3.050	1.033	0.003	—
South-East Asia (vs Africa)	2.806	1.393	0.044	—
Western Pacific (vs Africa)	1.281	1.045	0.221	—
n		179		
R2		0.404		