Reconstructed Mortality Series Understate Forecast Uncertainty: Evidence from The Gambia and Thirty Countries

Abdoulie Balisa[1]

[1]Department of Statistics and Actuarial Science, Kwame Nkrumah University of Science and Technology, Kumasi, Ghana

Draft, September 6, 2026

Abstract

Mortality models applied to countries without death registration produce prediction intervals that appear implausibly narrow. For The Gambia, classical Lee–Carter, a Bayesian Poisson Lee–Carter and the Li–Lee coherent extension all place life expectancy in 2074 within an interval three to five years wide, against roughly twenty years for the United Nations projection. This paper asks why. Three candidate explanations are tested and the first two are rejected. Fitting the same model to three successive World Population Prospects revisions, which are independent re-estimates of the same history, gives a between-reconstruction standard deviation of 0.30 to 0.58 years against a within-model standard deviation of 1.54 to 1.62 years, so propagating it widens the interval by at most 6%. Drift estimated on rolling equal-length windows varies by no more than its own standard error (ratio 0.9). A horizon analysis shows that the absence of a long-run asymptote in random-walk extrapolation accounts for roughly 14% of the gap: the ratio of interval widths is already 4.1 at a one-year horizon and grows only 1.27-fold over the following fifty years. The residual gap is multiplicative and independent of horizon, which points to the innovation standard deviation σ. Across thirty countries, median σ is 2.95 where civil registration is complete and 1.02 where it is not (Mann–Whitney $p < 0.0001$); The Gambia sits at 1.02. A normalisation-invariant version of the same comparison gives a smaller separation (1.57-fold), so the effect is bounded rather than pinned. The evidence is consistent with reconstruction removing the short-term variation from which σ is estimated, though the country groups also differ in income and age structure and a matched design would be needed to separate these. Separately, re-basing a validated cohort-component projection on the 2024 census lowers the projected 2074 population from 5.35 to 4.66 million (95% interval 4.35 to 4.98 million).

Keywords: mortality forecasting; Lee–Carter; uncertainty quantification; data-sparse populations; population projection; The Gambia

1 Introduction

Population projections underwrite long-range public decisions, including how many teachers to train, how large a future pension liability will be, and where to site clinics. The credibility of a projection rests on its stated uncertainty as much as on its central estimate.

For countries with complete vital registration, mortality forecasting rests on a mature methodology that begins with the log-bilinear model of Lee and Carter (1992) and has been extended in many directions since (Booth and Tickle, 2008). That literature assumes an input most of the world does not have: a long, directly observed series of age-specific death rates. Where registration is incomplete or absent, the input series is instead reconstructed. It is assembled from census age distributions, sibling and orphanhood histories in surveys, sample

registration and model life table systems (Wilmoth et al., 2012), then smoothed across age, time and neighbouring countries.

Analysts fit forecasting models to these reconstructions and report prediction intervals as though the input were observed. A reconstruction is a model output carrying its own error, so the natural suspicion is that treating it as an observation discards that error and produces intervals that are too narrow. This paper tests that suspicion and does not find support for it. It then tests whether the constant-drift assumption in random-walk extrapolation is responsible, and again finds no support. A third explanation, the absence of a long-run asymptote in the random walk, is supported but accounts for only a small share of the discrepancy. What remains is a multiplicative offset that is present at the shortest horizons, and Section 4.6 locates it in the innovation variance.

Two of the four results reported here are null. This is deliberate. The interval gap documented in Table 1 admits several plausible explanations, and establishing which of them fail against evidence is what makes the surviving explanation informative.

The Gambia is a useful case for this question. It has no functioning death registration and does not appear in the Human Mortality Database. It has two long-running demographic surveillance sites, Farafenni from 1981 and Basse from 2007, which provide empirical mortality data over a limited geography (Jasseh et al., 2015; Ezeani et al., 2025). In 2024 it completed its first fully digital census, which became available after the current UN projection round was finalised and therefore serves as an out-of-sample check on the reconstruction.

The paper makes the following contributions.

1. It documents a consistent gap between the prediction intervals produced by three standard mortality models fitted to a reconstruction and those produced by the UN hierarchical Bayesian method (Raftery et al., 2012; Alkema et al., 2011; Ševčíková and Raftery, 2016) (Section 4.2).

2. It tests and rejects the hypothesis that the gap arises because reconstruction uncertainty goes unpropagated (Section 4.3).

3. It tests and rejects the hypothesis that the gap arises from an unstable drift treated as constant (Section 4.4).

4. It shows by horizon analysis that the structural explanation accounts for about 14% of the gap, the remainder being independent of horizon (Section 4.5).

5. It locates the residual offset in the innovation standard deviation and shows across thirty countries that this parameter is systematically smaller where the mortality series is reconstructed (Section 4.6).

6. It validates a cohort-component projection engine against the UN projection and produces an independent projection for The Gambia anchored on the 2024 census (Sections 4.7 and 4.8).

2 Data

All inputs are public and retrieval is version-pinned, so the analysis reproduces exactly.

UN World Population Prospects 2024. Single-age death rates for 1950–2100, single-age population for 1949–2023, life expectancy estimates and the official probabilistic projection (United Nations, Department of Economic and Social Affairs, Population Division, 2024). Data were extracted for The Gambia and for a West African reference group (Senegal, Guinea-Bissau, Guinea, Mali, Sierra Leone, Mauritania and Burkina Faso) used by the coherent model.

Earlier WPP revisions. The 2022 and 2019 revisions, used in Section 4.3 as alternative reconstructions of the same mortality history. The 2022 revision shares the single-age annual grid of the 2024 revision. The 2019 revision is published on an abridged five-year grid.

The 2024 census. The Gambia Bureau of Statistics enumerated approximately 2,422,712 persons, with 40.8% aged under 15 and 3.0% aged 65 and over (Gambia Bureau of Statistics, 2024). This supplies the projection base population.

Validation sources. The 2013 and 2019–20 Demographic and Health Surveys (Gambia Bureau of Statistics and ICF, 2021), and published life tables from the Farafenni and Basse surveillance sites (Jasseh et al., 2015; Ezeani et al., 2025).

Status of the mortality series. The WPP series for The Gambia is a reconstruction rather than a record. It is treated as such throughout, and Section 4.3 makes that distinction an object of study rather than a caveat.

3 Methods

3.1 Life tables

Period life tables are constructed from age-specific death rates from first principles following Preston et al. (2001), using the Coale–Demeny rule for infant person-years and a standard close-out for the open age interval. As an internal check, the implementation reproduces the WPP published life expectancy for 2023 (65.86 years) to within 0.001 years.

3.2 Classical Lee–Carter

The model writes log death rates as

$$\log m_{x,t} = a_x + b_x k_t + \varepsilon_{x,t}, \tag{1}$$

identified by $\sum_x b_x = 1$ and $\sum_t k_t = 0$, estimated by singular value decomposition, with k_t extrapolated as a random walk with drift (Lee and Carter, 1992). Forecast uncertainty combines the innovation variance with uncertainty in the estimated drift.

3.3 Bayesian Poisson Lee–Carter

Following Brouhns et al. (2002) and Czado et al. (2005), deaths are modelled as

$$D_{x,t} \sim \text{Poisson}(E_{x,t} \exp(a_x + b_x k_t)), \tag{2}$$

where $E_{x,t}$ denotes exposure. We fix a_x at the mean log rate following the Lee–Carter identification, constrain b_x to the simplex and model k_t as a random walk with drift. Fixing a_x resolves an identifiability problem that otherwise prevents convergence. Sampling uses PyMC (Abril-Pla et al., 2023) and convergence is assessed by the rank-normalised $\widehat{R}$ of Vehtari et al. (2021), with a maximum of 1.010 across all parameters in the reported run.

3.4 Li–Lee coherent extension

The Gambia is forecast jointly with the West African reference group (Li and Lee, 2005). A common factor shared by the group carries the long-run trend, and a country-specific deviation is modelled as a mean-reverting process so that the country cannot drift indefinitely from its neighbours.

3.5 Cohort-component projection

A single-age, two-sex Leslie projection. Survivorship comes from the forecast life tables and fertility from the total fertility rate combined with the age pattern of childbearing, together with the sex ratio at birth and net migration. Credible intervals come from 1,000 simulations propagating our mortality uncertainty together with the UN fertility uncertainty.

4 Results

4.1 The census gap

WPP places The Gambia at 2,728,905 in 2023, approximately 13% above the 2024 census count of 2,422,712. The UN series also runs above every previous Gambian census. Because the WPP 2024 round was finalised before census results became available, the discrepancy is informative rather than contradictory. It provides an out-of-sample measurement of reconstruction error in the level of the population.

4.2 Mortality fit and forecast

The first singular component explains 99.3% of variation in log death rates, and estimated life expectancy rises from 31.2 years in 1950 to 65.9 in 2023.

In an out-of-sample backtest fitting only to 1950–2010 and predicting 2011–2023, mean absolute error in life expectancy is 0.65 years and the target falls inside the 95% interval in all 13 held-out years. The model slightly under-predicts the post-2010 improvement.

Table 1: Forecast life expectancy at birth in 2074, by method.

Method	e_0 in 2074	95% interval
Classical Lee–Carter	75.6	73.0–77.8
Bayesian Lee–Carter	75.0	73.8–76.2
Li–Lee coherent	74.8	73.0–76.6
UN WPP (hierarchical Bayes)	73.2	63.5–83.4

Table 1 sets out the problem the rest of the paper addresses. Three methodologically distinct models agree on the central estimate to within 0.8 years and all three produce intervals three to five years wide. The UN interval is approximately twenty years wide, a factor of four to six.

Agreement between the three models on the central estimate is reassuring, since they are independent in their treatment of trend. Their agreement on interval width is less informative, because all three are fitted to the same input series.

4.3 Reconstruction uncertainty

Successive WPP revisions re-estimate the same mortality history from different source data and methods. Fitting an identical model to several revisions therefore isolates the component of uncertainty attributable to the choice of reconstruction. We decompose the variance of forecast e_0 by the law of total variance,

$$\mathrm{Var}(e_0) = \underbrace{\mathbb{E}_r[\mathrm{Var}(e_0 \mid r)]}_{\text{within reconstruction}} + \underbrace{\mathrm{Var}_r[\mathbb{E}(e_0 \mid r)]}_{\text{between reconstructions}}, \tag{3}$$

where r indexes the reconstruction. The first term is what a single-source forecast reports; the second is what such a forecast omits.

All revisions are fitted on a common window ending in 2015 so that no revision benefits from a longer or fresher series, which would otherwise confound data vintage with reconstruction. Arm A compares the 2024 and 2022 revisions on their identical native single-age annual grid and requires no harmonisation. Arm B adds the 2019 revision on the abridged five-year grid, with the two modern revisions aggregated by exposure weighting and the life table closed at age 85.

Table 2: Variance decomposition of forecast e_0 in 2074 by source of uncertainty.

	Arm A (2024 vs 2022)	Arm B (2024, 2022, 2019)
Within-reconstruction SD (years)	1.62	1.54
Between-reconstruction SD (years)	0.58	0.30
Total SD (years)	1.72	1.57
Interval inflation factor	1.06×	1.02×

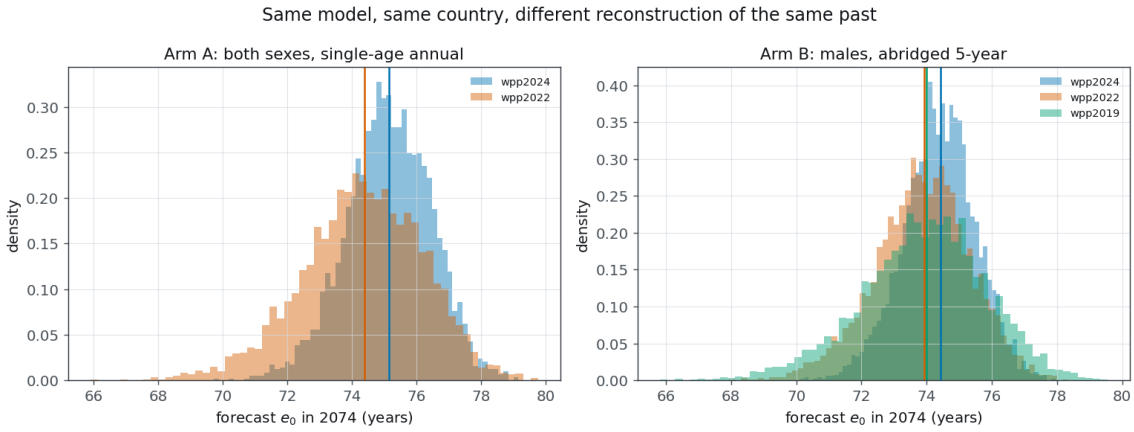

Figure 1: Forecast e_0 in 2074 under each reconstruction, for both arms.

The hypothesis is not supported. Between-reconstruction variation is between a fifth and a third of within-model variation, and propagating it widens the interval by 6% at most, against a gap of four to six times. The revisions agree closely about The Gambia's mortality history even though none of them observed it directly.

This null result requires one qualification. Revisions share source data, methodology and institutional practice, so they are not independent draws from the space of plausible reconstructions, and the measured between-revision spread should be read as a lower bound. That the revisions agree with each other is therefore weak evidence that they are individually accurate; as Section 4.1 shows, all of them sit 13% above the 2024 census. The narrower claim supported here is that whatever error the reconstructions share is not expressed as forecast-interval width, and so cannot be recovered by refitting across revisions.

4.4 Drift stability

The second candidate explanation concerns the treatment of trend. Lee–Carter forecasts k_t as a random walk with drift, estimating the drift d once and treating it thereafter as known up to a standard error $\sigma/\sqrt{T-1}$. That standard error describes uncertainty about d conditional on a single constant drift governing the entire forecast period. It cannot represent a change in the pace of mortality decline, which is a live possibility for a country whose life expectancy rose from 31 to 66 years in seven decades.

We fit the identical model to rolling 30-year windows. Equal window length makes the drift estimates directly comparable, and we compare variation across windows against the standard error reported within a window.

Table 3: Lee–Carter drift across rolling 30-year fitting windows.

Quantity	Value
Drift range across windows (per year)	-1.615 to -1.059
SD of drift across windows	0.185
Mean standard error within a window	0.214
Ratio	0.9×

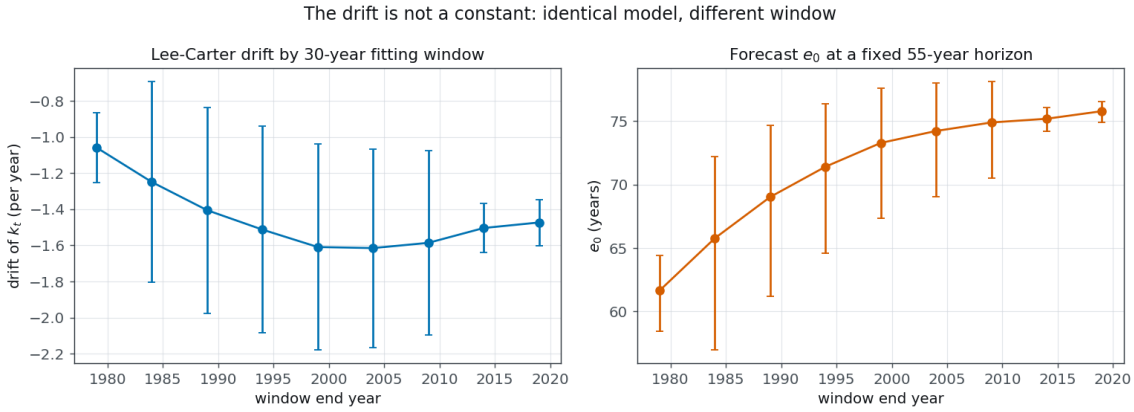

Figure 2: Drift by fitting window with 95% intervals (left), and forecast e_0 at a fixed 55-year horizon (right). The right panel is diagnostic only; see the text on confounding.

This hypothesis is also not supported. The point estimates move from -1.06 to -1.61 per year, but that movement is no larger than the sampling error of a single window's estimate, giving a ratio of 0.9. On this evidence the drift is consistent with a constant plus noise, and the constant-drift assumption does not account for the narrow intervals.

Two qualifications apply. The windows overlap, so the estimates are not independent and the ratio is descriptive rather than a formal test. In addition, comparing forecast e_0 across windows is confounded in either direction: holding the horizon fixed makes each window forecast to a different target year, while holding the target year fixed makes the horizons differ. The drift parameter is the only quantity cleanly comparable across windows, so it is reported alone.

4.5 Growth of the gap with horizon

The UN projections are produced by the Bayesian hierarchical framework of Raftery et al. (2012) and Alkema et al. (2011), implemented in `bayesPop` (Ševčíková and Raftery, 2016), in which the pace of mortality decline decelerates toward a ceiling estimated by borrowing strength across countries. A random walk with drift has no such ceiling.

This difference makes a testable prediction about shape. If the additional UN width comes from uncertainty about a long-run ceiling, that uncertainty is negligible one year out and substantial fifty years out, so the ratio of interval widths should grow with horizon. If the two methods instead differ by a fixed multiple at every horizon, the explanation lies elsewhere.

Table 4: 95% interval width for e_0 (years) by forecast horizon.

Horizon (years after 2023)	1	10	25	50
Lee–Carter (classical)	0.87	2.72	3.77	4.76
Lee–Carter (Bayesian)	0.56	1.39	1.91	2.37
Li–Lee coherent	0.83	2.34	3.31	3.54
UN WPP 2024	2.98	9.07	14.54	19.77
Ratio, UN to classical	3.42	3.33	3.86	4.15

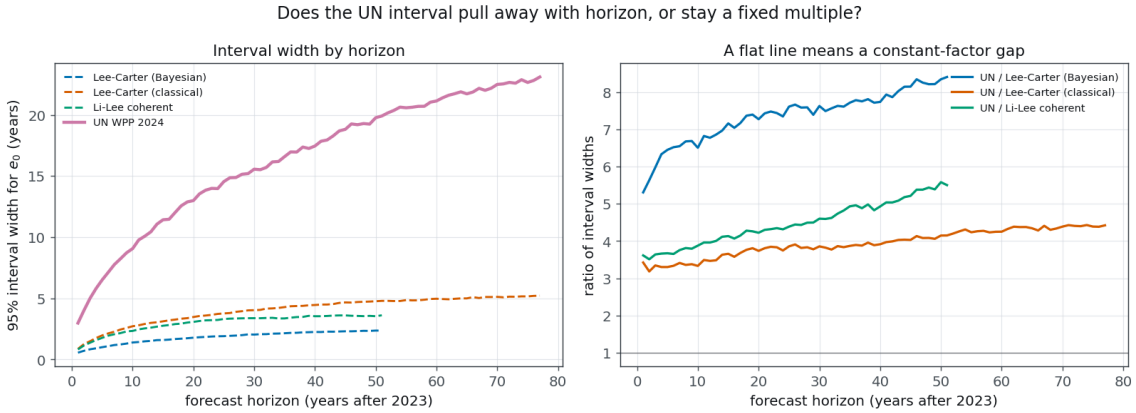

Figure 3: Interval width against horizon (left) and the ratio of the UN width to each model's width (right). A horizontal line in the right panel would indicate a constant-factor gap.

The prediction is partly borne out. Averaged across the three models, the ratio rises 1.27-fold between short horizons ($h \leq 15$) and long ones ($h \geq 40$). The ratio is already 4.1 at $h = 1$, however, where a long-run ceiling cannot yet be operative. On a multiplicative accounting, horizon growth explains about 14% of the total gap and a horizon-independent offset explains the remaining 86%.

Fitting $\log(\text{width}) = a + b \log h$ gives exponents of 0.34 for classical Lee–Carter, 0.36 for the Bayesian version, 0.29 for Li–Lee and 0.46 for the UN projection. All fall below the value of 0.5 implied by a pure random walk, because e_0 is a saturating transform of the death rates. The exponents are also close to one another, indicating that the methods accumulate uncertainty at broadly similar rates and differ mainly in level. A difference in level that is multiplicative and independent of horizon points to a scale parameter rather than to a mechanism operating at long range.

4.6 The innovation variance

In Lee–Carter the width of the forecast interval is scaled by σ, the innovation standard deviation of k_t, which is estimated from the residuals of the fitted series. Reconstruction interpolates sparse survey and census observations through model life tables, which suppresses short-term variation. If that variation is what σ is intended to capture, then σ estimated from a reconstruction will be too small, and the resulting interval too narrow by a roughly constant factor at every horizon.

We fit the identical model to WPP 2024 rates for 1950–2019 for fifteen countries with long-standing complete civil registration and fifteen without. The fitting window stops before 2020 so that COVID-19 mortality, which is a genuine shock in both groups, does not drive the comparison.

Table 5: Estimated innovation standard deviation by data regime, thirty countries, WPP 2024, 1950–2019.

	Registration ($n = 15$)	No registration ($n = 15$)	Ratio
Median σ	2.95	1.02	2.90×
Interquartile range	[2.40, 3.45]	[0.82, 1.29]	
Median $\sigma/\lvert\text{drift}\rvert$	1.31	0.84	1.57×

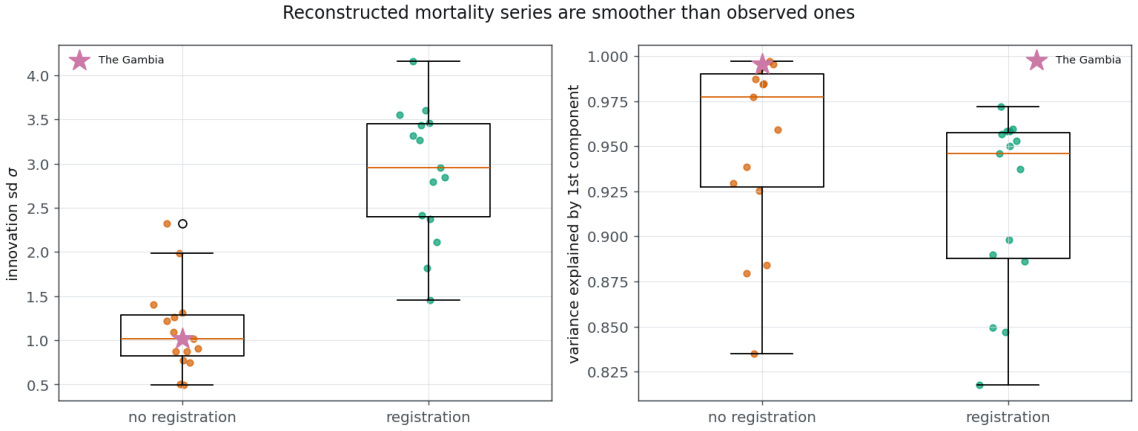

Figure 4: Estimated σ by data regime (left) and variance explained by the first singular component (right). The Gambia is marked.

The two groups separate clearly (Mann–Whitney test of registration exceeding no registration, $p < 0.0001$). The Gambia sits at $\sigma = 1.02$, a factor of 2.90 below the registration median, which is of the same order as the 4.1-fold horizon-independent gap reported in Section 4.5.

Three qualifications bound this result. The second is the most serious.

First, countries are not randomly assigned to the two groups. Registration countries are richer and demographically older, and an older population can express mortality shocks, such as influenza epidemics and heat waves, that a young population cannot. Data regime and age structure are therefore confounded.

Second, k_t is identified only up to the normalisation $\sum_x b_x = 1$, so σ is not strictly on a common scale across countries. The normalisation-invariant ratio $\sigma/\lvert\text{drift}\rvert$ separates the groups in the same direction but by a factor of 1.57 rather than 2.90. The effect is therefore bounded between roughly 1.6 and 2.9, and we do not claim the point estimate.

Third, several countries in the no-registration group experienced substantial HIV/AIDS and conflict mortality that survives into the reconstruction. Uganda ($\sigma = 2.32$) and Sierra Leone

($\sigma = 1.99$) are the two highest values in their group. This inflates the no-registration median and makes the comparison conservative.

Variance explained by the first singular component barely separates the groups (median 0.977 without registration against 0.946 with). Smoothness measured in that way does not distinguish the two regimes; the magnitude of σ does.

4.7 Engine validation

Supplied with the UN inputs, the cohort-component engine reproduces the UN published medium projection to within 0.3% to 0.9% across 2030–2074 (Table 6). Reproducing the benchmark from the benchmark's own inputs establishes that the engine can be trusted to carry alternative inputs.

Table 6: Engine validation against the WPP 2024 medium projection.

Year	This engine	WPP 2024	Difference
2030	3,149,576	3,160,605	−0.3%
2040	3,738,105	3,762,462	−0.6%
2050	4,300,578	4,328,827	−0.7%
2074	5,298,317	5,346,053	−0.9%

4.8 Census-based projection

Table 7 reports the projection re-based on the 2024 census and driven by our mortality uncertainty. Re-basing lowers the 2074 figure by approximately 0.7 million relative to the UN medium variant.

Table 7: Projected population of The Gambia, census-based.

Year	Median	95% interval	UN medium
2050	3.74M	3.59–3.90M	4.33M
2074	4.66M	4.35–4.98M	5.35M

4.9 Age structure

The total dependency ratio falls from 77 per 100 working-age adults in 2023 to approximately 49 by 2074. Its composition changes substantially over the same period: child dependency falls from 71.7 to 31.8, while old-age dependency more than triples, from 5.4 to 17.8.

5 Discussion

Of the three explanations tested for the narrowness of the forecast intervals, two are rejected and the third is supported but small. The intervals are not narrow because different reconstructions of The Gambia's mortality history disagree in ways that go unpropagated, since they agree closely on both trend and level. They are not narrow because an unstable drift is held constant, since the drift moves no more than its own standard error. The absence of a long-run asymptote in the random walk does contribute, but accounts for roughly a seventh of the gap.

The evidence instead points to the innovation variance. The residual gap is multiplicative and present at a one-year horizon, which is the behaviour of a scale parameter, and σ is the parameter that scales the interval. Countries whose mortality series are reconstructed have systematically smaller σ than countries whose series are close to observed, by a factor between roughly 1.6 and 2.9 depending on normalisation.

This interpretation should be held with appropriate caution. The country comparison is observational, and data regime is confounded with income and age structure. What the comparison establishes is an association of the right sign and approximately the right magnitude to account for the residual gap; it does not establish that reconstruction causes the reduction in σ. A matched design, comparing countries at similar levels of life expectancy and age structure that differ only in registration completeness, would be required to separate these. Directly observed mortality data for The Gambia, such as the Farafenni and Basse surveillance series, would provide a stronger test still, since it would allow σ to be estimated for the same population under two data regimes.

Subject to that caution, two implications follow. The first concerns practice. Where a mortality series is reconstructed rather than registered, estimating σ from that series will understate forecast uncertainty, and an externally informed value should be considered instead. Model life table systems already borrow an age pattern from comparable populations, and borrowing a volatility parameter on the same logic would be a modest extension. Our estimates suggest a correction factor between roughly 1.6 and 2.9 for The Gambia, and we would not narrow that range without the matched design described above.

The second implication is more general. Any forecasting model whose interval is scaled by a residual variance estimated from the fitted series will understate uncertainty when that series is model output rather than observation. Much applied forecasting in low- and middle-income settings relies on interpolated or model-based series because no alternative exists. The concern in such cases is not necessarily bias in level, which the reconstructions examined here do not display relative to one another, but suppressed short-term variation, which propagates directly into interval width.

For The Gambia specifically, the projection has two consequences for planning. The population will continue to grow substantially, but plans calibrated to UN figures are working from a base approximately 13% too high and a 2074 total approximately 0.7 million too high. At the same time the country is entering a period of falling total dependency, while old-age dependency triples against pension and geriatric care capacity that is currently minimal.

6 Limitations

The mortality inputs are partly reconstructions and the central estimates inherit their assumptions. Comparison against the Farafenni and Basse surveillance life tables, and against Global Burden of Disease estimates, requires manual extraction and has not been completed. The census base uses the published broad age structure combined with the UN within-group shape; detailed single-age census tables would refine it. Migration is treated as a central scenario, which is a strong assumption for The Gambia given its volatility, and scenario analysis is a priority for revision. The country comparison in Section 4.6 is observational and its limitations are set out in that section and in the discussion.

7 Conclusion

Mortality forecasts for countries without death registration report prediction intervals several times narrower than those of the United Nations for the same quantity. Three candidate explanations were examined here. Disagreement between successive reconstructions of the same history is small relative to model uncertainty, and propagating it widens the interval by at most 6%. Drift estimated on rolling windows is stable to within its own standard error. The absence of a long-run asymptote in random-walk extrapolation accounts for roughly 14% of the gap. The residual is multiplicative and independent of horizon, which locates it in the innovation standard deviation rather than in any of the three.

Across thirty countries that parameter is roughly three times larger where civil registration is complete than where it is not, and The Gambia sits at the low end. The evidence is consistent with reconstruction suppressing the short-term variation from which the parameter is estimated, so that a series smoothed during its construction yields a model confident in proportion to the smoothing. A normalisation-invariant comparison gives a smaller separation, and the two country groups differ in income and age structure as well as in registration completeness, so the mechanism is bounded rather than established. Separating it would need a matched design.

The practical consequence is that an interval reported for a country without death registration should not be read at face value, and the narrower it is, the less it should be trusted. For The Gambia specifically, re-basing a validated cohort-component projection on the 2024 census lowers the projected 2074 population from 5.35 to 4.66 million, with a 95% interval of 4.35 to 4.98 million.

8 Reproducibility

All code, pinned data extracts and figures are available at `https://github.com/Balisa50/gambia-population-projection`, which is private during review and will be made public on publication; access is available to editors and reviewers on request. Data retrieval is pinned to a specific upstream commit, so re-running the pipeline reproduces identical inputs. Each numerical result in this paper is produced by a named script in that repository.

References

Abril-Pla, O., Andreani, V., Carroll, C., Dong, L., Fonnesbeck, C. J., Kochurov, M., Kumar, R., Lao, J., Luhmann, C. C., Martin, O. A., Osthege, M., Vieira, R., Wiecki, T., and Zinkov, R. (2023). PyMC: A modern, and comprehensive probabilistic programming framework in Python. *PeerJ Computer Science*, 9:e1516.

Alkema, L., Raftery, A. E., Gerland, P., Clark, S. J., Pelletier, F., Buettner, T., and Heilig, G. K. (2011). Probabilistic projections of the total fertility rate for all countries. *Demography*, 48(3):815–839.

Booth, H. and Tickle, L. (2008). Mortality modelling and forecasting: A review of methods. *Annals of Actuarial Science*, 3(1–2):3–43.

Brouhns, N., Denuit, M., and Vermunt, J. K. (2002). A Poisson log-bilinear regression approach to the construction of projected lifetables. *Insurance: Mathematics and Economics*, 31(3):373–393.

Czado, C., Delwarde, A., and Denuit, M. (2005). Bayesian Poisson log-bilinear mortality projections. *Insurance: Mathematics and Economics*, 36(3):260–284.

Ezeani, E. S., Gollam, S., Mohammed, N., Roca, A., Hossain, J., Hossain, I., Saine, A. K., D'Alessandro, U., and Mackenzie, G. (2025). Cohort profile: Basse health and demographic surveillance system, the Gambia. *International Journal of Epidemiology*, 54(2).

Gambia Bureau of Statistics (2024). The Gambia population and housing census 2024. Technical report, Gambia Bureau of Statistics, Banjul.

Gambia Bureau of Statistics and ICF (2021). The Gambia demographic and health survey 2019–20. Technical report, GBoS and ICF, Banjul, The Gambia and Rockville, Maryland, USA.

Jasseh, M., Gomez, P., Greenwood, B. M., Howie, S. R. C., Scott, S., Snell, P. C., Bojang, K., Cham, M., Corrah, T., and D'Alessandro, U. (2015). Health and demographic surveillance system profile: Farafenni health and demographic surveillance system in The Gambia. *International Journal of Epidemiology*, 44(3):837–847.

Lee, R. D. and Carter, L. R. (1992). Modeling and forecasting U.S. mortality. *Journal of the American Statistical Association*, 87(419):659–671.

Li, N. and Lee, R. (2005). Coherent mortality forecasts for a group of populations: An extension of the Lee–Carter method. *Demography*, 42(3):575–594.

Preston, S. H., Heuveline, P., and Guillot, M. (2001). *Demography: Measuring and Modeling Population Processes*. Blackwell, Oxford.

Raftery, A. E., Li, N., Ševčíková, H., Gerland, P., and Heilig, G. K. (2012). Bayesian probabilistic population projections for all countries. *Proceedings of the National Academy of Sciences*, 109(35):13915–13921.

Ševčíková, H. and Raftery, A. E. (2016). bayesPop: Probabilistic population projections. *Journal of Statistical Software*, 75(5).

United Nations, Department of Economic and Social Affairs, Population Division (2024). World population prospects 2024. Technical report, United Nations, New York.

Vehtari, A., Gelman, A., Simpson, D., Carpenter, B., and Bürkner, P.-C. (2021). Rank-normalization, folding, and localization: An improved $\widehat{R}$ for assessing convergence of MCMC (with discussion). *Bayesian Analysis*, 16(2):667–718.

Wilmoth, J., Zureick, S., Canudas-Romo, V., Inoue, M., and Sawyer, C. (2012). A flexible two-dimensional mortality model for use in indirect estimation. *Population Studies*, 66(1):1–28.