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.

Universal but Unequal: Bayesian Small-Area Estimation of District Fertility and Infant Mortality Transitions in Rwanda, 2006–2025

Guillaume Rudasingwa

guillaume92@sch.ac.kr

Soonchunhyang University

Maurice Tuyisenge

Seoul Hanyoung University

Nancy Clemence Ishimwe

Maximilians-Universität Würzburg

Frank Aimee Rodrigue Ndagijimana

Seoul National University

Research Article

Keywords: Small-area estimation, Fertility transition, Infant mortality, BYM2, Demographic and Health Surveys, Rwanda

Posted Date: November 13th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11552207/v1

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

Additional Declarations: No competing interests reported.

Universal but Unequal: Bayesian Small-Area Estimation of District Fertility and Infant Mortality Transitions in Rwanda, 2006–2025

Article type: Research Article

Word count: ~5,700 (main text, excluding abstract, references, tables and figure captions); 5 tables, 6 figures

Abstract

Rwanda's fertility and infant mortality fell rapidly after 2005, but national surveys do not show how these transitions unfolded across the 30 districts through which resources are planned. Using the 2010, 2014–15, 2019–20 and 2025 Rwanda Demographic and Health Surveys (2,044 clusters; 32,313 births), we estimated district total fertility rates (TFR) and infant mortality rates (IMR) for four five-year periods from 2006–2010 to 2020–2025. Direct estimates reproduced the official national series and were smoothed with Bayesian space–time small-area models with Besag–York–Mollié (BYM2) district levels and trends, jackknife design variances and outcome-specific survey likelihoods. The fertility transition was universal but unequal: TFR fell credibly in every district, by 0.72–1.43 children (18–27%), but declines were proportional, so relative dispersion and the core–periphery geography persisted. Fertility remained low in Kigali and central districts and high in the Lake Kivu, southern and eastern districts (Moran's I ≈ 0.4). IMR fell by 43–51% in every district, to 23–30 per 1,000. The surveys could not resolve district differences in infant mortality: no district differed credibly from the national level, so these estimates suit broad surveillance rather than district ranking. Both transitions stalled nationwide in 2015–2020 and resumed thereafter, and the district geographies of fertility and mortality were unrelated. Longer travel time to cities was associated with higher fertility; this ecological association does not establish individual-level effects of access. Rwanda's fertility decline reached every district without dissolving its spatial structure.

Keywords: Small-area estimation · Fertility transition · Infant mortality · BYM2 · Demographic and Health Surveys · Rwanda

1 Introduction

Rwanda's demographic change since the mid-2000s has been among the fastest in sub-Saharan Africa. The total fertility rate (TFR) fell from 6.1 children per woman in the 2005 Demographic and Health Survey (DHS) to 4.6 in 2010 (National Institute of Statistics of Rwanda et al., 2012). Analysts attributed this to rapid uptake of modern contraception within a strongly state-led family-planning programme (Muhoza et al., 2013; Westoff, 2013). Infant and child mortality fell just as sharply, in parallel with community-based health insurance, community health workers and the scale-up of malaria control and vaccination (Binagwaho et al., 2014; Farmer et al., 2013). The decline did not continue at an even pace. The national TFR barely changed between the 2014–15 and 2019–20 surveys (4.2 and 4.1), a pattern that resembles the fertility stalls documented elsewhere in the region (Bongaarts, 2008; Schoumaker, 2019). Infant mortality was also essentially unchanged between those two surveys (32 and 33 per 1,000) (National Institute of Statistics of Rwanda et al., 2016; National Institute of Statistics of Rwanda et al., 2021). The 2025 survey shows that both declines have resumed, to a TFR of 3.7 and an IMR of 27 per 1,000 (National Institute of Statistics of Rwanda et al., 2025).

National averages hide considerable within-country variation, and Rwanda allocates resources and sets targets through its 30 districts. District estimates of fertility and infant mortality are therefore of direct policy value. They also bear on central questions in spatial demography: whether a rapid national transition reaches all places at once or spreads outward from early-adopting cores, whether it narrows or preserves spatial inequality, and whether the geographies of fertility and mortality move together. Comparative transition theory predicts that mortality decline precedes and conditions fertility decline (Angeles, 2010; Bongaarts & Casterline, 2013). Spatial diffusion accounts predict that new reproductive behaviour spreads from urban cores and along social and geographic networks, which would produce spatially structured, rather than random, district differences. Whether district differences in Rwanda are narrowing (σ-convergence) and whether initially high-fertility and high-mortality districts are catching up (β-convergence) are open empirical questions.

DHS samples were not designed for district estimation. In Rwanda each district contains only 15–28 sampled clusters per survey and roughly 160–390 births in the five years before interview. Direct survey-weighted estimates of infant mortality are therefore highly unstable, with coefficients of variation that routinely exceed 30%. Small-area estimation (SAE) addresses this by borrowing

strength across space and time through hierarchical models while respecting the complex survey design (Rao & Molina, 2015; Wakefield et al., 2020). Area-level models that smooth design-based direct estimates have become the standard approach to subnational child mortality from DHS data (Li et al., 2019; Mercer et al., 2015; Wu et al., 2021). Comparable work on fertility is more recent (Wu & Wakefield, 2025). For Rwanda, Ren (2021) produced design-based district estimates of fertility and child mortality by pooling clusters, but only for the 2010 and 2014–15 surveys. To our knowledge, no study has produced model-based district estimates spanning the 2025 survey, quantified their uncertainty, and examined the convergence of fertility and mortality jointly.

This paper fills that gap. Using four Rwanda DHS surveys (2010, 2014–15, 2019–20 and 2025), we estimate district TFR and infant mortality rates (IMR) for four consecutive five-year periods from 2006–2010 to 2020–2025. We use direct estimators that reproduce the official national series, and Bayesian space–time models with BYM2 spatial structure (the Besag–York–Mollié model in the parameterisation of Riebler et al. (2016)) and design-consistent likelihoods. We ask three questions. (i) How much did district fertility and infant mortality change, and how certain are those changes? (ii) Did district differences narrow, and did the most disadvantaged districts catch up? (iii) Are district differences spatially structured, and how far are they associated with geospatial environmental and access covariates once spatial structure is accounted for? The contribution is primarily spatial-demographic. We test whether Rwanda's fertility decline was accompanied by spatial convergence or by a persistent core–periphery structure, whether the 2015–2020 stall was concentrated in particular districts or was nationwide, and whether district fertility and mortality geographies coincide, as transition theory would lead one to expect. We summarise the answer as a transition that was *universal but unequal*. By *universal* we mean that every district experienced a decline with high posterior probability. By *unequal* we mean that relative fertility disparities and the spatial ranking of districts persisted. The phrase describes fertility; for infant mortality, the surveys cannot resolve district differences with useful precision. Methodologically, we document the choices on which district estimates from DHS data depend, including censoring in births-based mortality estimators, design-effect adjustment and covariate harmonisation, because each of them materially changes the results.

2 Data and Methods

2.1 Data sources

We used the four most recent Rwanda Demographic and Health Surveys (RDHS): 2010 (fieldwork September 2010–March 2011), 2014–15 (November 2014–April 2015), 2019–20 (November 2019–July 2020) and 2025 (June–October 2025) (National Institute of Statistics of Rwanda et al., 2012; National Institute of Statistics of Rwanda et al., 2016; National Institute of Statistics of Rwanda et al., 2021, 2025, 2026). The 2025 dates follow the Key Indicators Report (5 June–25 October 2025) and agree with the interview dates recorded in the data; the NISR survey web page gives June–November 2025. A few 2010 interviews are dated April 2011 in the data. Each RDHS is a stratified two-stage cluster sample. Enumeration areas (clusters) are drawn first, with probability proportional to size, from strata formed by district and urban–rural residence. A fixed number of households is then sampled in each cluster. Every district is therefore a sampling domain with 15–28 clusters per survey (Table 1). The samples were nonetheless designed for national and provincial estimates (Ren, 2021). NISR district profiles report direct district values for selected indicators, including a five-year TFR for 2025, to one decimal and without measures of uncertainty; district infant mortality is not reported (National Institute of Statistics of Rwanda, 2026). Where those 2025 district TFRs are stated in the profile text, our direct estimates differ from them by 0.15–0.26 children per woman, less than one jackknife standard error (for example, Nyaruguru 4.54 vs. 4.8 and Nyamagabe 3.28 vs. 3.1).

We used the women's recode (IR) for exposure, the birth recode (BR, complete birth histories) for births and deaths, and the cluster-level geospatial covariate files (GC) (Mayala et al., 2018). The recode versions were RWIR61, RWIR70, RWIR81 and RWIR91 (IR); RWBR61, RWBR70, RWBR81 and RWBR91 (BR); RWGE61, RWGE72, RWGE81 and RWGE91 (GE); and RWGC62, RWGC72, RWGC81 and RWGC91 (GC), downloaded from The DHS Program on 11 September 2026. In total the four surveys contained 2,044 clusters and 32,313 births in the 60 months before interview. Districts were taken from each survey's own district variable and matched to the 30 second-level administrative units of the Database of Global Administrative Areas (GADM), version 4.1 (GADM, 2022). As a check, we also assigned clusters to districts by point-in-polygon overlay of the displaced GPS coordinates (GE files). The two assignments agreed for 99.6% (2010, 2014–15) and 100% (2019–20, 2025) of clusters (Online Resource 1, Table S1). DHS displaces cluster coordinates by up to 2 km in urban areas and 5 km in rural areas (Burgert et al., 2013; Warren et al.,

2016). Because districts come from the survey variable, displacement affects only the linkage of the raster covariates, not district membership.

2.2 Outcomes and direct estimates

All district estimates refer to the 60 months before interview. The four surveys therefore yield four almost contiguous five-year periods, which we label 2006–2010, 2010–2015, 2015–2020 and 2020–2025. Their reference-period midpoints, computed from the weighted mean interview date minus 2.5 years, are 2008.5, 2012.6, 2017.7 and 2023.1.

Total fertility. For each woman, exposure in each five-year age group from 15–19 to 45–49 was computed as the number of months lived in that age group within the window (interview month excluded), using her date of birth and interview date. Births from the birth histories were classified by the mother's age at birth. Age-specific rates are weighted births divided by weighted woman-years, and the TFR is five times their sum (Croft et al., 2018; Schoumaker, 2013). We used a five-year window to increase the number of births per district. As a check on the data processing, we also computed the national three-year TFR, which reproduced the published values to within 0.04 (Table 2).

Infant mortality. We estimated infant mortality with a period life table built from child-months of exposure within the 60-month window. Each child contributes the months of age it lived inside the window: 0, 1–2, 3–5 and 6–11 months, the DHS age segments (Croft et al., 2018). A death is counted in the segment and calendar month in which it occurred. The monthly hazard in each segment is the weighted number of deaths divided by the weighted child-months. The IMR is one minus the product of the segment survival probabilities. Only exposure observed before interview enters the estimator. This avoids the right-censoring that affects births-based ratios, in which children born in the last 12 months are counted as births but have not been observed for a full year. Nationally, our estimates reproduced the published synthetic-cohort IMR to within 1 per 1,000 in every survey (49.4, 31.9, 32.0 and 26.5 vs. 50, 32, 33 and 27; Table 2). We therefore did not calibrate the district estimates to the published values.

Design-based variances. Sampling variances, and the corresponding standard errors (SE), were estimated by the stratified delete-one-cluster jackknife (JKn) within each district–period, with the DHS strata (district × urban–rural) as jackknife strata (Rao & Molina, 2015). When a cluster is

deleted, the remaining clusters in its stratum are reweighted by $n_h/(n_h - 1)$, and the variance is $\sum_h \frac{n_h-1}{n_h}\sum_{i\in h}\left(\hat{\theta}_{(hi)} - \hat{\theta}\right)^2$. Strata with a single cluster in a district were merged with the other stratum of that district. Variances were computed directly for log TFR and for the IMR. Kish-type design effects for infant mortality were computed as deff = V_{JK} / $\{\hat{p}(1 - \hat{p})/n\}$, where n is the unweighted number of births in the window.

2.3 Small-area model

Let $d = 1, \ldots, 30$ index districts and $t = 1, \ldots, 4$ periods. We fitted area-level Bayesian hierarchical models that treat the direct estimates as noisy measurements of latent district–period quantities (Fay & Herriot, 1979; Mercer et al., 2015; Wakefield et al., 2020).

Fertility (Fay–Herriot model). With $y_{dt} = \log \hat{TFR}_{dt}$ and design variance V_{dt},

$$y_{dt} \mid \eta_{dt}, e_{dt} \sim N(\eta_{dt} + e_{dt}, V_{dt}), \qquad e_{dt} \sim N(0, \sigma_e^2).$$

Infant mortality (effective-sample-size binomial). A Fay–Herriot model on the logit scale breaks down when a district records very few deaths, because the jackknife replicates can equal zero; one district–period recorded a single death (Table 1). We therefore used the working likelihood of Chen et al. (2014):

$$p(y^*_{dt} \mid p_{dt}) \propto p_{dt}^{y^*_{dt}}(1 - p_{dt})^{n^*_{dt} - y^*_{dt}}, \qquad n^*_{dt} = n_{dt}/\overline{\text{deff}}_t, \quad y^*_{dt} = n^*_{dt}\,\hat{p}_{dt},$$

where $\hat{p}_{dt}$ is the weighted life-table estimate, n_{dt} the unweighted number of births and $\overline{\text{deff}}_t$ the median design effect of survey t (1.03–1.25). The effective sample size carries the information loss from clustering and weighting into the likelihood. Here logit $p_{dt} = \eta_{dt} + e_{dt}$.

Linear predictor. For both outcomes,

$$\eta_{dt} = \beta_{0t} + \mathbf{x}_{dt}^{\top}\boldsymbol{\beta} + u_d + v_d\,\tau_t,$$

where β_{0t} are period-specific intercepts, u_d is a district level and v_d a district-specific linear trend in τ_t. τ_t is the reference-period midpoint, centred and expressed in decades. $\mathbf{x}_{dt}$ are covariates (models M4–M5 only), and e_{dt} captures district–period deviations not explained by the smooth components. Both u and v follow the BYM2 model, a reparameterisation of the Besag–York–Mollié model (Besag et al., 1991) proposed by Riebler et al. (2016):

$$u = \sigma_u\left(\sqrt{\rho_u}\, s_u^* + \sqrt{1-\rho_u}\, \theta_u\right), \qquad \theta_u \sim N(0, I),$$

where s_u^* is an intrinsic conditional autoregressive (ICAR) field on the queen-contiguity graph of the 30 districts (69 edges, no islands). The field is scaled so that the geometric mean of its marginal variances is one (Sørbye & Rue, 2014). σ_u is then the total standard deviation of the district effects, and ρ_u is the share that is spatially structured. The specification is a discrete-space analogue of the linear space–time interaction model of Bernardinelli et al. (1995), with an additional unstructured interaction term (Knorr-Held, 2000). With four periods, fixed period intercepts are preferable to a random walk. They impose no smoothness on the national trajectory and use only four parameters.

Priors. Following Simpson et al. (2017), all standard deviations received penalised-complexity (PC) priors with $\Pr(\sigma > 1) = 0.01$. The mixing parameters ρ_u and ρ_v received the PC prior of Riebler et al. (2016) with $\Pr(\rho < 0.5) = 2/3$, computed from the eigenvalues of the scaled ICAR covariance matrix. Period intercepts received $N(\mu_0, 5^2)$ priors centred near the national level on the model scale, and covariate coefficients received N(0, 1) priors.

Model sequence. We fitted five nested models for each outcome:

- M1: period intercepts, iid district effects and e_{dt}.
- M2: BYM2 district effects in place of the iid district effects.
- M3: M2 plus BYM2 district trends.
- M4: M3 plus covariates.
- M5: a non-spatial analogue of M3 plus covariates.

We specified M3 in advance as the primary small-area model. The aim was to describe district levels and trends without conditioning on covariates whose comparability across surveys is uncertain. M4 and M5 were used to describe covariate associations and to show how they change when spatial random effects are added (Hodges & Reich, 2010).

Covariates. We selected six DHS geospatial covariates in advance. For each survey we used the GC column closest in year to the survey (Table S2):

- enhanced vegetation index (EVI);
- land-surface temperature;

- aridity index;
- travel time to the nearest city (log);
- insecticide-treated net (ITN) coverage;
- *Plasmodium falciparum* incidence (log).

The DHS missing-value code (−9,999) was recoded to missing, and EVI was rescaled to its 0–1 range where the extract stores it multiplied by 10,000. District values are means over sampled clusters. The underlying raster products change between DHS releases, so we standardised each covariate within survey round. β therefore measures the cross-sectional association of a one-standard-deviation difference between districts, and the period intercepts absorb any level shifts. Night-time lights and population density were excluded because of their collinearity with travel time. Human footprint and wet days were excluded because they were correlated with travel time and with malaria incidence. All retained covariates had variance inflation factors below 5 (maximum 4.3, land-surface temperature).

2.4 Computation, model assessment and derived quantities

Models were fitted in Stan 2.32 through rstan (Carpenter et al., 2017) with four chains of 3,000 iterations (1,500 warm-up), giving 6,000 posterior draws per model. The ICAR fields used the pairwise-difference formulation with a soft sum-to-zero constraint (Morris et al., 2019). For the Fay–Herriot model, e_{dt} was integrated out analytically and drawn afterwards from its conditional posterior, which stabilises leave-one-out cross-validation. Convergence was assessed with the rank-normalised $\hat{R}$ and the bulk effective sample size of the Markov chain Monte Carlo draws (bulk ESS) (Vehtari et al., 2021); this is distinct from the survey effective sample size n^*_{dt} above. Models were compared by the expected log pointwise predictive density (ELPD), estimated with Pareto-smoothed importance-sampling leave-one-out (LOO) cross-validation (PSIS-LOO) (Vehtari et al., 2017). Calibration was assessed through the coverage of LOO posterior predictive intervals for the direct estimates and through LOO probability integral transform (PIT) values.

From the posterior draws we computed the following quantities:

- district medians with 95% credible intervals (CrI);
- the probability that each district exceeds the national value of the same period;
- changes between 2006–2010 and 2020–2025, with the probability of decline;

- posterior ranks;
- convergence measures, computed per draw: the between-district coefficient of variation (σ-convergence) and the slope of change on initial level (β-convergence);
- the across-district correlation between TFR and IMR.

Global Moran's I was computed for direct estimates and for posterior medians with row-standardised queen weights and 9,999 permutations (Bivand & Wong, 2018; Moran, 1950). Because smoothing itself induces similarity between neighbours, Moran's I of posterior medians can overstate clustering. We therefore also computed Moran's I for each posterior draw of the latent district surface, which gives a posterior distribution for spatial autocorrelation.

We ran three sensitivity analyses. First, we refitted the M3 mortality model with a naive binomial likelihood on unweighted counts, which ignores clustering. Second, we compared the model-based national IMR with the published series. Third, we compared the district estimates and probabilities of decline from M3 with those from the other four models.

All analyses were run in R 4.3.3 with haven, sf, spdep, rstan and loo (Bivand & Wong, 2018; Pebesma, 2018; Vehtari et al., 2017), and all figures were produced with the ggplot2 and patchwork packages in R. A single script reproduces every table and figure from the DHS recodes (Online Resource 2).

Use of AI tools. Claude (Anthropic) was used to help synthesise the literature related to the study and the methods used, and to polish the language of the manuscript. The authors reviewed and edited all content and take full responsibility for it.

3 Results

3.1 Sample and direct estimates

The four surveys contributed 2,044 clusters, 32,313 births and 1,149 infant deaths in the 60-month windows (Table 1). District samples were small: 164–387 births and 1–26 infant deaths per district–period. As mortality fell, the median coefficient of variation (CV) of the direct district IMR rose from 0.29 in 2006–2010 to 0.48 in 2020–2025. In the latest period the direct CV exceeded 0.3 in 28 of 30 districts and 0.5 in 13. Across district–periods the direct IMRs ranged from 0.4 to 86.5 per

1,000 (Online Resource 1, Table S7), an implausible spread that illustrates the instability small-area estimation is meant to address. Direct TFR was much more precise (median CV 0.06–0.07). Design effects for infant mortality were modest (1.03–1.25), as expected for a rare outcome with small cluster sizes.

Table 1 Sample sizes, precision of direct estimates and design effects by survey

Survey	Reference period	Clusters (min–max per district)	Births (min–max per district)	Infant deaths (min–max)	Median CV, direct TFR	Median CV, direct IMR	Design effect, IMR
2010	2006–2010	492 (16–20)	9,095 (238–359)	446 (7–26)	0.065	0.291	1.16
2014–15	2010–2015	492 (16–20)	7,883 (185–334)	248 (2–14)	0.071	0.365	1.03
2019–20	2015–2020	500 (16–21)	8,121 (208–369)	264 (3–15)	0.064	0.394	1.20
2025	2020–2025	560 (15–28)	7,214 (164–387)	191 (1–12)	0.063	0.477	1.25

Note: Counts are unweighted. CV is the jackknife standard error divided by the direct estimate. The IMR design effect is the median over districts with at least five deaths.

Nationally, the direct five-year TFR fell from 4.92 to 3.75 children per woman, and the direct IMR fell from 49.4 to 26.5 per 1,000 (Table 2). The three-year TFR (4.56, 4.16, 4.13, 3.68) and the life-table IMR both reproduced the official RDHS figures closely. This validates the processing of exposure, births and deaths before any modelling. The five-year TFR is higher than the three-year TFR in 2010 (4.92 vs. 4.56) because the longer window reaches back into the faster decline before 2010. The model-based national aggregates were close to the direct national estimates (Fig. 2), so smoothing did not distort the national series. The model-based national IMR (49.3, 32.4, 32.2 and 26.4) lay within 2.3% of the published values. Benchmarking would therefore change district levels by less than 3%.

Table 2 National estimates: direct, model-based (aggregated from districts) and published RDHS values

Survey (period)	TFR 5-yr direct (SE)	TFR 5-yr model (95% CrI)	TFR 3-yr direct	TFR 3-yr published	IMR direct (SE)	IMR model (95% CrI)	IMR published (0–4 yr)
2010 (2006–2010)	4.92 (0.06)	4.90 (4.79–5.02)	4.56	4.6	49.4 (2.5)	49.3 (44.5–53.9)	50
2014–15 (2010–2015)	4.11 (0.06)	4.12 (4.02–4.23)	4.16	4.2	31.9 (2.2)	32.4 (28.7–36.6)	32
2019–20 (2015–2020)	4.16 (0.05)	4.19 (4.09–4.29)	4.13	4.1	32.0 (2.3)	32.2 (28.2–36.7)	33
2025 (2020–2025)	3.75 (0.05)	3.77 (3.68–3.86)	3.68	3.7	26.5 (2.3)	26.4 (22.5–30.9)	27

Note: The five-year estimates (0–59 months before interview) are the basis of the district analysis. The three-year TFR (0–35 months) and the life-table IMR are shown for comparison with the

published RDHS figures (TFR: 3 years before the survey; IMR: 0–4 years before the survey). Jackknife SEs in parentheses.

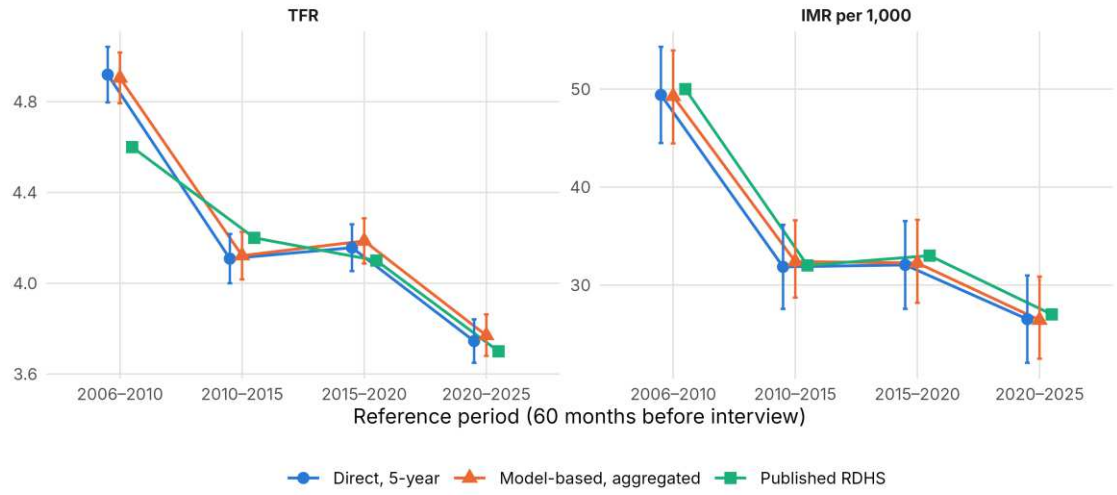


Fig. 2 *National trajectories of TFR and IMR: direct five-year estimates (95% CI), model-based estimates aggregated from districts (95% CrI) and published RDHS values (TFR: 3-year; IMR: 0–4 years)*

3.2 Model comparison and diagnostics

All models converged ($\hat{R} \leq 1.007$, minimum bulk ESS 746, no divergent transitions; Table S3). Among the models without covariates, differences in PSIS-LOO were small: at most 2.4 ELPD units (about two standard errors) for TFR, and within one standard error for IMR (Table 3). Adding covariates worsened predictive performance for TFR (ΔELPD = −5.1, SE 2.7, for M4 versus M2) and left it unchanged for IMR. LOO predictive intervals were well calibrated: under M3, empirical coverage of the 95% intervals was 0.94 for TFR and 0.96 for IMR. The LOO-PIT values did not depart from uniformity (Kolmogorov–Smirnov $p \geq 0.50$ for all models), and at most 7 of 120 Pareto k values exceeded 0.7.

PSIS-LOO thus slightly favoured M2 for TFR and M4 for IMR, by margins of one to two standard errors or less. We nonetheless retain M3, specified in advance, as the primary model for three reasons. It is the most parsimonious model that gives each district its own spatially smoothed trend, which the convergence questions require. It does not condition on covariates whose comparability across surveys is uncertain. And the choice does not change the substantive conclusions. District TFR medians under M3 correlated at 0.995 or more with those of every other model (largest absolute difference 0.18 children), and IMR medians correlated at 0.999 with M2 (largest difference 1.5 per

1,000) and at 0.93 with the covariate models M4 and M5 (largest difference 11 per 1,000). Under every model, every district had a posterior probability of decline between 2006–2010 and 2020–2025 above 0.99 for TFR and above 0.96 for IMR (Online Resource 1, Table S5).

Table 3 Model comparison by PSIS-LOO and LOO predictive calibration

Outcome	Model	Structure	ELPD-LOO (SE)	ΔELPD (SE)	p_loo	Pareto k > 0.7	LOO 95% coverage	PIT KS p
TFR	M1	Period + iid district + iid district–period	118.1 (8.9)	−2.4 (1.1)	30.7	4	0.95	0.65
TFR	M2	Period + BYM2 level + iid	120.5 (8.8)	0.0 (0.0)	28.5	2	0.94	0.81
TFR	M3	M2 + BYM2 district trends (primary)	119.3 (8.9)	−1.2 (0.6)	31.9	3	0.94	0.66
TFR	M4	M3 + covariates	115.4 (9.1)	−5.1 (2.7)	35.2	6	0.94	0.52
TFR	M5	Non-spatial M3 analogue + covariates	113.8 (8.9)	−6.8 (2.7)	35.8	7	0.94	0.50
IMR	M1	Period + iid district + iid district–period	−295.2 (8.1)	−1.0 (3.0)	14.6	0	0.96	0.97
IMR	M2	Period + BYM2 level + iid	−294.8 (8.0)	−0.7 (2.9)	14.8	0	0.95	0.98
IMR	M3	M2 + BYM2 district trends (primary)	−295.1 (7.9)	−1.0 (2.9)	15.6	0	0.96	0.98
IMR	M4	M3 + covariates	−294.1 (7.5)	0.0 (0.0)	16.2	0	0.96	0.80
IMR	M5	Non-spatial M3 analogue + covariates	−294.3 (7.5)	−0.2 (0.2)	16.3	0	0.95	0.85

Note: ΔELPD is relative to the best model for each outcome (higher ELPD is better). Coverage is the share of the 120 direct estimates inside their LOO 95% posterior predictive intervals. PIT KS p is the Kolmogorov–Smirnov test of uniformity of the LOO-PIT values.

Under M3, the district-level standard deviation of log TFR was 0.098 (95% CrI 0.072–0.136). This was large relative to the district–period deviation ($\sigma_e = 0.034$) and to the variation in district trends ($\sigma_v = 0.018$ per decade). About 70% of the level variation was spatially structured ($\rho_u = 0.71$, 95% CrI 0.12–0.98). For infant mortality all variance components were small and poorly identified ($\sigma_u = 0.106$, 0.008–0.213; $\rho_u = 0.28$, 0.01–0.94). Beyond the period effects, the data carry little information about persistent district differences in infant mortality (Table S4).

Smoothing reduced the median CV of district IMR from 0.29–0.48 (direct) to 0.13–0.16 (model), and that of TFR from 0.06–0.07 to about 0.04 (Table S6). After smoothing, no district–period had a CV above 0.2 for either outcome. Model-based TFR estimates stayed close to the direct values, whereas IMR estimates were pulled strongly towards the period means, as their much larger sampling variances imply (Fig. 5).

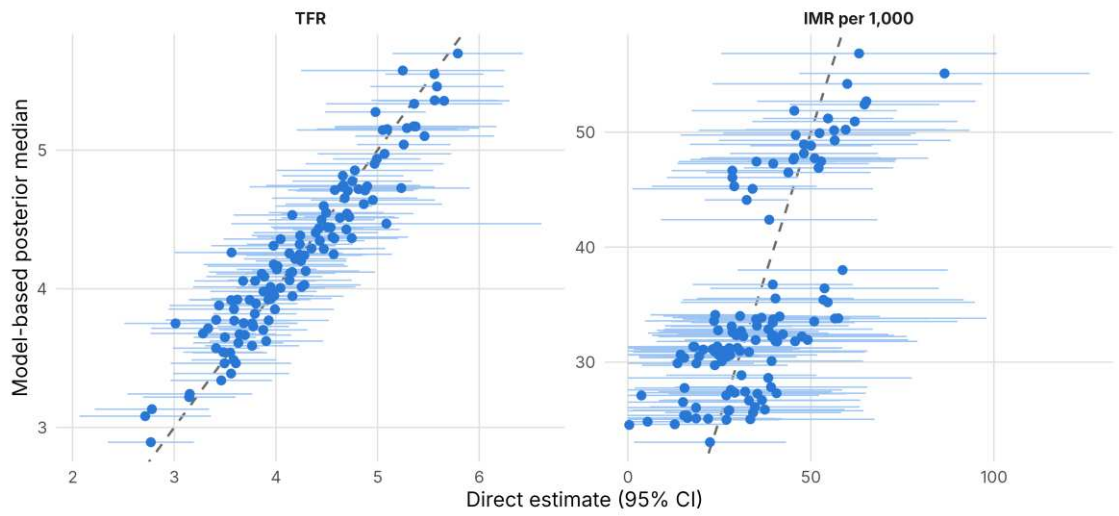


Fig. 5 *Direct estimates (with 95% CIs) against model-based posterior medians for all 120 district–periods. The dashed line is the identity*

3.3 District fertility: levels, transition and the 2015–2020 stall

District TFR ranged from 3.75 (Kicukiro) to 5.70 (Nyaruguru) in 2006–2010 and from 2.89 (Kicukiro) to 4.52 (Rusizi) in 2020–2025 (Fig. 1a; Table 4). The geography is persistent and clearly structured:

- **Lowest fertility:** the three City of Kigali districts (Kicukiro, Nyarugenge, Gasabo) and the adjacent central and northern districts (Muhanga, Rulindo, Kamonyi).
- **Highest fertility:** a periphery formed by the Lake Kivu districts of the Western Province (Rusizi, Nyamasheke, Rubavu), southern Nyaruguru and Gisagara, and the eastern districts (Nyagatare, Bugesera, Ngoma, Kayonza).

Global Moran's I of the posterior median TFR was 0.39–0.46 in every period (permutation $p \leq 0.0005$; Table S9). Computed draw by draw, the posterior Moran's I of the latent TFR surface was 0.33–0.40, with 95% CrIs excluding zero in every period (for 2020–2025, 0.37, 95% CrI 0.22–0.50). In 2020–2025, nine districts had a posterior probability above 0.9 of exceeding the national TFR, and seven had a probability below 0.1 (Fig. 4a).

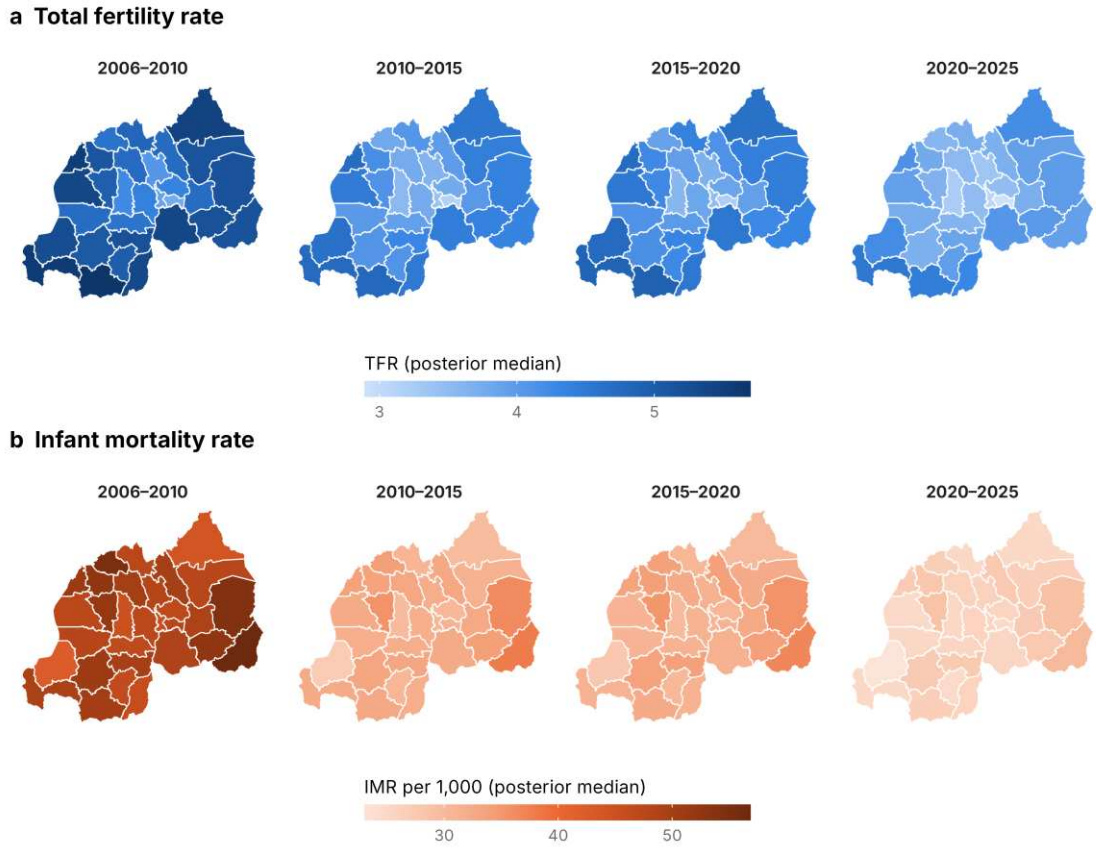


Fig. 1 *Posterior median district (a) total fertility rate and (b) infant mortality rate per 1,000 live births, by five-year period (model M3)*

Table 4 District TFR and IMR (per 1,000) in 2020–2025 and change since 2006–2010 for the five lowest- and five highest-fertility districts, primary model M3

District	Province	TFR 2020–2025	P(TFR > national)	ΔTFR since 2006–2010	IMR 2020–2025	P(IMR > national)	ΔIMR since 2006–2010
Kicukiro	Kigali	2.89 (2.62–3.18)	0.00	−0.86 (−1.23 to −0.50)	24.5 (17.2–32.4)	0.31	−22.9 (−37.9 to −12.1)
Nyarugenge	Kigali	3.08 (2.76–3.44)	0.00	−0.80 (−1.18 to −0.36)	25.9 (18.8–35.0)	0.44	−19.5 (−30.8 to −4.7)
Muhanga	South	3.24 (2.93–3.58)	0.00	−0.99 (−1.44 to −0.60)	25.0 (17.8–33.9)	0.36	−19.9 (−32.2 to −7.0)
Rulindo	North	3.34 (3.06–3.68)	0.01	−0.72 (−1.07 to −0.26)	26.1 (19.3–34.8)	0.47	−21.9 (−35.6 to −9.5)
Gasabo	Kigali	3.46 (3.15–3.82)	0.04	−0.90 (−1.31 to −0.44)	25.6 (19.2–34.4)	0.42	−21.1 (−34.1 to −8.2)
…							
Nyagatare	East	4.20 (3.88–4.54)	1.00	−1.25 (−1.77 to −0.78)	25.0 (18.0–34.1)	0.36	−19.1 (−31.1 to −5.3)
Nyamasheke	West	4.22 (3.88–4.57)	0.99	−1.06 (−1.50 to −0.55)	23.0 (15.4–31.8)	0.21	−19.1 (−31.9 to −7.3)
Gisagara	South	4.29 (3.99–4.60)	1.00	−1.06 (−1.58 to −0.52)	25.8 (18.9–35.1)	0.44	−20.3 (−32.5 to −5.6)
Nyaruguru	South	4.44 (4.08–4.82)	1.00	−1.26 (−1.76 to −0.75)	26.7 (19.6–36.1)	0.53	−23.4 (−38.6 to −10.9)

District	Province	TFR 2020–2025	P(TFR > national)	ΔTFR since 2006–2010	IMR 2020–2025	P(IMR > national)	ΔIMR since 2006–2010
Rusizi	West	4.52 (4.18–4.88)	1.00	−1.06 (−1.59 to −0.43)	25.1 (18.0–33.4)	0.36	−23.7 (−39.1 to −12.2)

Note: Posterior medians with 95% credible intervals. P(· > national) is the posterior probability that the district exceeds the national value for 2020–2025. Districts are ordered by TFR; estimates for all 30 districts are given in Online Resource 1, Table S8.

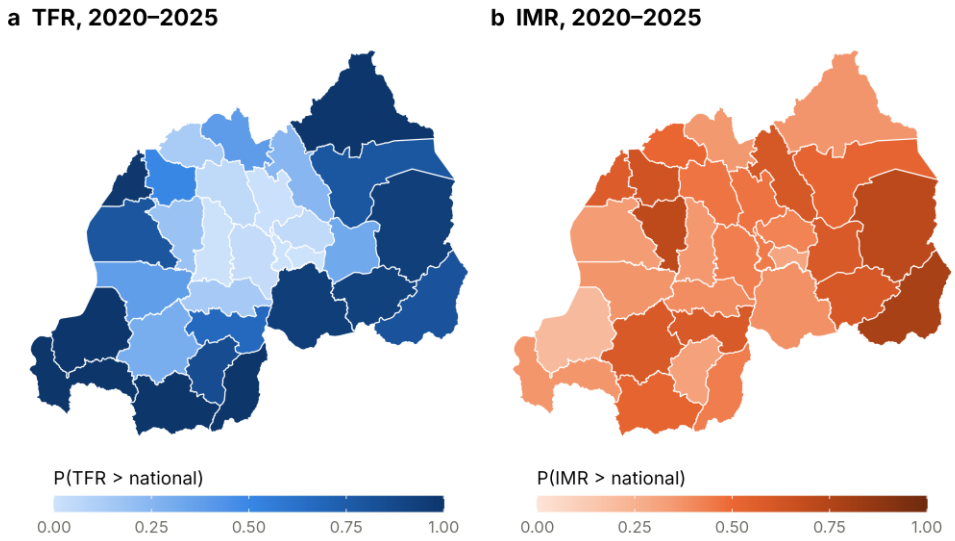


Fig. 4 *Posterior probability that the district (a) TFR and (b) IMR exceed the national value, 2020–2025*

Every district's TFR fell between 2006–2010 and 2020–2025 (posterior probability of decline ≥ 0.998). Absolute declines ranged from 0.72 to 1.43 children per woman (relative declines 18–27%; Fig. 3a; Table 4). The largest declines were in the west and south (Rutsiro −1.43, Rubavu −1.40, Nyamagabe −1.36, Nyabihu −1.32). The smallest were in Kigali and its neighbours (Rulindo −0.72, Nyarugenge −0.80, Kicukiro −0.86, Gasabo −0.90, Kamonyi −0.90).

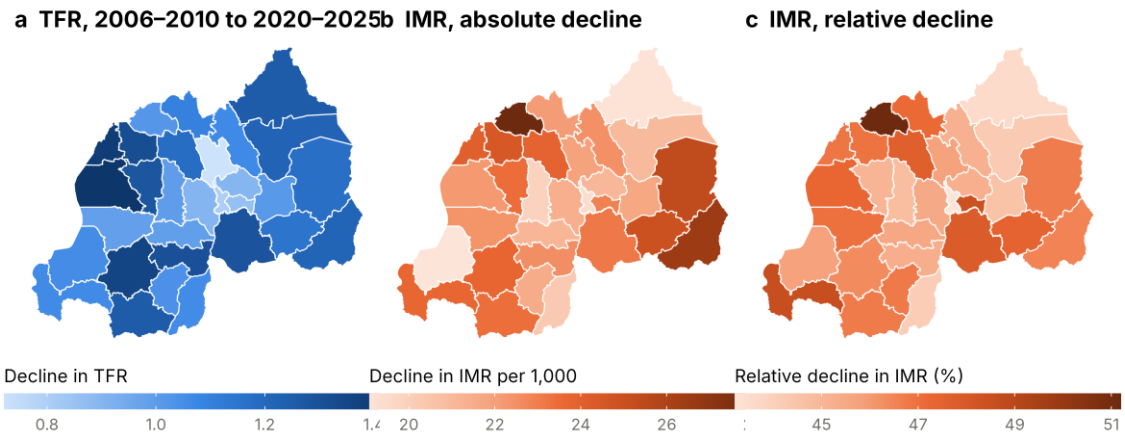


Fig. 3 *Posterior median decline between 2006–2010 and 2020–2025 in (a) TFR, (b) IMR per 1,000 and (c) IMR in relative terms. The posterior probability of decline exceeded 0.99 in every district for both outcomes*

The national stall between 2010–2015 and 2015–2020 appears in every district:

- **2010–2015 to 2015–2020:** no district had a posterior probability of decline above 0.9, and in 26 of 30 districts it was below 0.5.
- **2015–2020 to 2020–2025:** the decline resumed, with probabilities of decline of 0.85–1.00 (above 0.95 in 24 districts).

The stall was therefore a nationwide pause in the transition rather than a local phenomenon.

The absolute spread of district TFRs narrowed: the between-district standard deviation fell from 0.54 to 0.41 (Table S10). Relative dispersion did not change. The CV was 0.111 in 2006–2010 and 0.108 in 2020–2025, and the posterior probability that it fell was 0.63. Declines were roughly proportional, and the rank order of districts changed little. Districts with higher initial TFR had larger absolute declines (β-convergence slope −0.35, 95% CrI −0.53 to −0.21; Table S11). Because the initial level is itself estimated with error, this slope partly reflects regression to the mean and should be read together with the stable CV.

3.4 District infant mortality

District IMR ranged from 42.4 to 56.8 per 1,000 in 2006–2010 and from 23.0 to 30.1 in 2020–2025 (Fig. 1b; Table 4). Every district's IMR declined, with posterior probability above 0.99. The median decline was 22.4 deaths per 1,000 (range 19.1–28.1), or 43–51% in relative terms (Fig. 3b, c).

Mortality decline paused over the same interval as fertility. Between 2010–2015 and 2015–2020 the probability of decline was 0.42–0.61 in every district, so these periods are statistically indistinguishable. Between 2015–2020 and 2020–2025 it was 0.87–0.96.

District differences in infant mortality were much less certain than those in fertility:

- No district in any period had a posterior probability above 0.9 of exceeding the national IMR. The highest was Kirehe, at 0.85 in 2006–2010 and 0.78 in 2020–2025 (Fig. 4b).
- The 95% credible intervals of the posterior IMR ranks for 2020–2025 span 26–29 of the 30 positions, against a median of 14–15 positions for TFR (Table S13).
- Spatial clustering of infant mortality is not established. Moran's I of the posterior medians was 0.19–0.25 (p = 0.012–0.035), but this partly reflects the smoothing itself. The posterior Moran's I of the latent IMR surface was only 0.05–0.07, with 95% CrIs from about −0.2 to 0.4 in every period (Table S9). The direct estimates show no clustering after 2010 (I = −0.13 to 0.07).

In 2020–2025, relatively high posterior IMR was concentrated in the east (Kirehe 30.1, Kayonza 28.9, Ngoma 27.6) and in the north-western highlands (Ngororero 28.6, Nyabihu 27.8). It was lowest in Nyamasheke (23.0), Kicukiro (24.5) and Huye (24.6). These contrasts are of the order of 5–7 deaths per 1,000 and lie well within the credible intervals. The data therefore cannot resolve district IMR differences with the available sample sizes: the posterior intervals are wide, and the model-based estimates are dominated by smoothing towards the period means. This is not evidence that district mortality is identical. In practice, the district IMRs are suited to broad surveillance of levels and trends, not to fine-grained ranking of districts or performance targeting.

The across-district correlation between TFR and IMR was weak in every period (0.02–0.10), with 95% CrIs from about −0.3 to 0.4 (Table S12). Fertility and infant mortality declined in parallel nationally, but their district geographies were essentially unrelated.

3.5 Covariate associations

Of the six covariates, only travel time to the nearest city had a credible association with fertility. A one-standard-deviation longer (log) travel time was associated with a 6.7% higher TFR (rate ratio 1.067, 95% CrI 1.025–1.108; Table 5; Fig. 6).

Table 5 Covariate associations: rate ratios (TFR) and odds ratios (IMR) per one-standard-deviation difference between districts within a survey

Outcome	Covariate (per SD, within survey)	M4 spatial: ratio (95% CrI)	M4 P(>1)	M5 non-spatial: ratio (95% CrI)
TFR	Vegetation index (EVI)	0.989 (0.962–1.019)	0.238	0.989 (0.961–1.019)
TFR	Land surface temperature	1.008 (0.954–1.067)	0.605	1.013 (0.958–1.070)
TFR	Aridity index	1.020 (0.965–1.080)	0.745	1.021 (0.970–1.076)
TFR	Travel time to nearest city (log)	1.067 (1.025–1.108)	0.999	1.076 (1.037–1.114)
TFR	ITN coverage	1.010 (0.991–1.032)	0.840	1.009 (0.989–1.029)
TFR	*P. falciparum* incidence (log)	0.997 (0.977–1.019)	0.393	0.998 (0.977–1.020)
IMR	Vegetation index (EVI)	1.043 (0.955–1.143)	0.830	1.040 (0.948–1.136)
IMR	Land surface temperature	0.825 (0.708–0.961)	0.007	0.822 (0.710–0.947)
IMR	Aridity index	0.859 (0.744–0.989)	0.018	0.858 (0.749–0.977)
IMR	Travel time to nearest city (log)	0.953 (0.866–1.049)	0.159	0.954 (0.866–1.058)
IMR	ITN coverage	1.029 (0.950–1.116)	0.761	1.029 (0.952–1.115)
IMR	*P. falciparum* incidence (log)	1.079 (0.988–1.180)	0.955	1.080 (0.987–1.180)

Note: M4 is the spatial model (BYM2 level and trend) and M5 its non-spatial analogue. P(>1) is the posterior probability that the ratio exceeds one. Land-surface temperature and the aridity index are strongly negatively correlated across districts (r = −0.80), so their IMR coefficients describe a conditional contrast, not two independent effects. All associations are ecological (between districts) and do not establish individual-level effects.

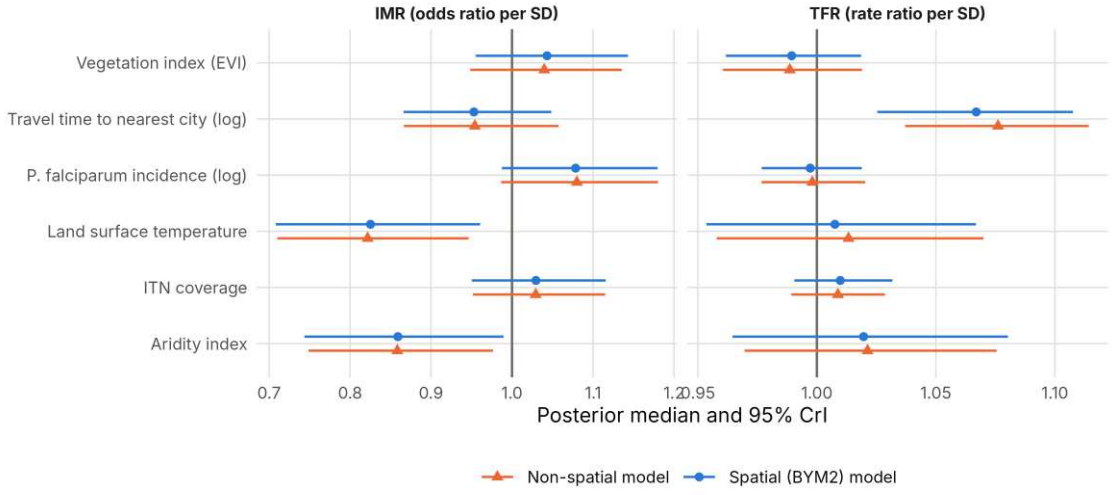


Fig. 6 *Covariate associations in the spatial (M4) and non-spatial (M5) models: posterior medians and 95% CrIs per one-standard-deviation difference*

For infant mortality, land-surface temperature (odds ratio 0.83, 0.71–0.96) and the aridity index (0.86, 0.74–0.99) had credible negative coefficients. These two covariates are strongly negatively correlated with each other (r = −0.80), so their coefficients describe a conditional contrast rather than two independent effects. *P. falciparum* incidence had a positive but uncertain association (odds ratio

1.08, 0.99–1.18; posterior probability of a positive effect 0.95). ITN coverage was not associated with IMR (1.03, 0.95–1.12).

Coefficients from the spatial (M4) and non-spatial (M5) models were nearly identical, so spatial confounding did not affect these associations. Together with the LOO results, this indicates that the covariates explain little of the district variation that the space–time structure does not already capture.

3.6 Sensitivity analyses

Ignoring the survey design in the mortality likelihood (a naive binomial on unweighted counts) gave district medians that correlated at 0.998 with the effective-sample-size model. The 95% interval widths differed by a median of 2% (range −4% to +10%; Online Resource 1, Section S4), reflecting the modest design effects in these surveys. Ratio-benchmarking the district IMRs to the published national series would multiply them by 1.015, 0.988, 1.023 and 1.022 in the four periods. Rankings and exceedance probabilities would be unchanged.

4 Discussion

This study produced design-consistent district estimates of fertility and infant mortality for Rwanda across four five-year periods from 2006–2010 to 2020–2025. The direct estimators reproduce the official national series, and the model-based district estimates come with explicit uncertainty. Four findings stand out.

First, the fertility transition was universal across districts but has not levelled them geographically. Every district reduced its TFR, typically by about a fifth to a quarter. Because the declines were roughly proportional, the spatial ranking was preserved. Kigali and the central–northern districts led throughout. The Lake Kivu, southern and eastern districts lagged. In 2020–2025 the highest-fertility districts (Rusizi, Nyaruguru, Gisagara, Nyamasheke, Nyagatare; TFR 4.2–4.5) were at about the national level of 2010–2015 (4.1). As an illustrative benchmark rather than a formal convergence measure, this places them roughly a decade behind the national trajectory, and further still behind Kigali. Absolute gaps narrowed while relative gaps did not, and national monitoring of district targets should make this distinction explicit.

The spatial autocorrelation of district fertility was strong and stable (posterior Moran's I 0.33–0.40, with credible intervals excluding zero in every period), and a large share of the district variance was spatially structured. Together with the positive association with travel time to cities, this is consistent with diffusion along an urban core–periphery gradient, as described for fertility transitions in sub-Saharan Africa (Bongaarts & Casterline, 2013). Our ecological design cannot distinguish diffusion from the spatial clustering of compositional factors, such as education, contraceptive use and livelihoods, which were deliberately not included in the model.

Second, the national fertility stall between the 2014–15 and 2019–20 surveys (National Institute of Statistics of Rwanda et al., 2021; Schoumaker, 2019) was nationwide, not concentrated in particular districts. In no district do the data support a decline between 2010–2015 and 2015–2020. In almost every district they support a decline afterwards. A stall that is uniform across very different districts points to national drivers (programme, supply or policy factors) rather than to local compositional change, although we can only raise this as a hypothesis for further work. Infant mortality was also flat over the same interval. One candidate explanation is the malaria resurgence of 2012–2016, when reported cases rose roughly eightfold (Karema et al., 2020). Consistent with this, district *P. falciparum* incidence had a positive, though uncertain, association with IMR.

Third, the surveys cannot resolve district differences in infant mortality with the available sample sizes. In 2020–2025 the model-based district range was 23–30 per 1,000, the posterior intervals were wide, and no district could be distinguished from the national value with high posterior probability. This is not evidence that district mortality is uniform; real differences of a few deaths per 1,000 may exist but cannot be detected. A narrow true dispersion would nonetheless be consistent with the near-universal coverage of Rwanda's community-based health system (Binagwaho et al., 2014; Farmer et al., 2013). It contrasts with the much wider within-country dispersion of child mortality reported for many African countries (Burstein et al., 2019; Golding et al., 2017; Li et al., 2019). For users of the estimates, the practical implication is that district IMRs from DHS data are suited to broad surveillance of levels and trends, not to ranking districts, allocating resources by district mortality or assessing district performance. Direct estimates are especially misleading: in our data they ranged from 0.4 to 87 per 1,000, almost entirely because of sampling noise. The model-based surface hints at relatively higher mortality in the eastern districts and the north-western highlands, but the posterior Moran's I of infant mortality was close to zero, so this pattern is not established.

Fourth, the fertility and mortality geographies were essentially unrelated across districts. Classic transition accounts treat mortality decline as a precondition for fertility decline (Angeles, 2010). At the district scale in Rwanda, however, fertility variation appears to be organised by an urban–rural and access gradient that infant mortality does not share. A plausible reason is that mortality decline in Rwanda was driven by nationally uniform health-system programmes, whereas fertility responds to more slowly changing local social and economic conditions.

4.1 Extensions

The framework extends naturally in three directions. First, splitting infant mortality into neonatal and post-neonatal components, or extending the life table to under-five mortality, would help to separate health-system and environmental influences; the same child-month estimator and effective-sample-size likelihood apply, although neonatal deaths are rarer still at the district level. Second, a joint model of fertility and mortality, with correlated district effects, would estimate the fertility–mortality association directly rather than from posterior draws of separate models. Third, routine data from Rwanda's health management information system (HMIS) record facility births and deaths every month in every district. Combined with the survey estimates in a data-integration model, with the survey anchoring levels and the routine data informing timing and district contrasts, they could supply the annual district series that four surveys cannot.

4.2 Methodological implications

District small-area estimates from DHS data depend materially on several choices that are easy to get wrong. We document them here.

1. **The mortality estimator.** Births-based ratios that count deaths before age one among births in the last five years are right-censored, because births in the final year have not been fully observed. In the same surveys, such ratios understated the official IMR by 3–6%. A period life table built from child-months, as used here, or a discrete-hazards formulation (Croft et al., 2018; Mercer et al., 2015), removes this bias. With such an estimator, benchmarking to the official series becomes a check rather than a correction.
2. **The likelihood for rare outcomes.** Fay–Herriot models on the logit scale fail when districts record one or two deaths, because the jackknife variance becomes undefined or degenerate.

The effective-sample-size binomial working likelihood (Chen et al., 2014) handles such cells naturally.

3. **Covariate harmonisation.** The DHS geospatial covariate files change their underlying products between releases. They also code missing values as −9,999, and older releases store EVI on a different scale. Pooling covariates across surveys without harmonisation can produce spurious covariate effects and temporal artefacts. Standardising within survey round avoids level artefacts but limits interpretation to cross-sectional associations.
4. **Temporal resolution.** Four surveys identify four period levels. They do not identify annual dynamics, and annual district series built by interpolating between surveys should not be presented as estimates.

4.3 Limitations

The most important limitation is that the analysis is ecological. The covariate associations describe districts, not women or children, and should not be interpreted causally; in particular, the association between travel time and fertility does not show that poorer access raises an individual woman's fertility. They are also conditional on a spatial field and on period effects, and with 30 areas the BYM2 mixing parameters are only weakly identified. District membership came from the survey design variable, but the geospatial covariates were extracted at displaced coordinates (Burgert et al., 2013; Warren et al., 2016). This adds measurement error to the covariates and probably attenuates their associations. Our jackknife merges strata that contain a single cluster in a district, which makes those variances slightly conservative. The life-table IMR assumes a constant hazard within each age segment, and it inherits the age heaping of reported ages at death (notably at 12 months) that affects all DHS mortality estimates. The five-year TFR is not directly comparable with the published three-year TFR, which we report separately. Finally, our national benchmarks for 2025 (TFR 3.7, IMR 27 per 1,000) are those of the Key Indicators Report (National Institute of Statistics of Rwanda et al., 2025), which the final report published in August 2026 supersedes (National Institute of Statistics of Rwanda et al., 2026). The district estimates are computed from the released microdata and do not depend on either report, but the comparison in Table 2 refers to the preliminary values.

4.4 Conclusions

Between 2006–2010 and 2020–2025, fertility fell by about a fifth to a quarter, and infant mortality by almost half, in every district of Rwanda. Both transitions paused nationwide in 2015–2020 and resumed afterwards. The fertility transition was universal but unequal: universal because every district experienced a decline with high posterior probability, and unequal because relative disparities and the spatial ranking of districts persisted along a core–periphery gradient. For infant mortality, the surveys cannot resolve district differences, and the district estimates should be used for broad surveillance rather than for ranking or targeting. The association of fertility with travel time to cities is ecological and does not establish individual-level effects of access. District planning should therefore prioritise the high-fertility periphery. Claims about district mortality disparities should rest on model-based estimates with explicit uncertainty rather than on direct survey values.

Supplementary Information

Online Resource (ESM_1.pdf): supplementary methods, tables ,figures and R codes

References

Angeles, L. (2010). Demographic transitions: Analyzing the effects of mortality on fertility. *Journal of Population Economics*, *23*(1), 99–120. https://doi.org/10.1007/s00148-009-0255-6

Bernardinelli, L., Clayton, D., Pascutto, C., Montomoli, C., Ghislandi, M., & Songini, M. (1995). Bayesian analysis of space-time variation in disease risk. *Statistics in Medicine*, *14*(21-22), 2433–2443. https://doi.org/10.1002/sim.4780142112

Besag, J., York, J., & Mollié, A. (1991). Bayesian image restoration, with two applications in spatial statistics. *Annals of the Institute of Statistical Mathematics*, *43*(1), 1–20. https://doi.org/10.1007/BF00116466

Binagwaho, A., Farmer, P. E., Nsanzimana, S., Karema, C., Gasana, M., de Dieu Ngirabega, J., Ngabo, F., Wagner, C. M., Nutt, C. T., Nyatanyi, T., & Gatera, M. (2014). Rwanda 20 years on: Investing in life. *The Lancet*, *384*(9940), 371–375.

Bivand, R. S., & Wong, D. W. S. (2018). Comparing implementations of global and local indicators of spatial association. *TEST*, *27*(3), 716–748. https://doi.org/10.1007/s11749-018-0599-x

Bongaarts, J. (2008). Fertility transitions in developing countries: Progress or stagnation? *Studies in Family Planning*, *39*(2), 105–110. https://doi.org/10.1111/j.1728-4465.2008.00157.x

Bongaarts, J., & Casterline, J. (2013). Fertility transition: Is sub-Saharan Africa different? *Population and Development Review*, *38*(s1), 153–168. https://doi.org/10.1111/j.1728-4457.2013.00557.x

Burgert, C. R., Colston, J., Roy, T., & Zachary, B. (2013). *Geographic displacement procedure and georeferenced data release policy for the demographic and health surveys. DHS Spatial Analysis Reports no. 7* (DHS Spatial Analysis Reports 7). ICF International.

Burstein, R., Henry, N. J., Collison, M. L., Marczak, L. B., Sligar, A., Watson, S., et al. (2019). Mapping 123 million neonatal, infant and child deaths between 2000 and 2017. *Nature*, *574*(7778), 353–358. https://doi.org/10.1038/s41586-019-1545-0

Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., & Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, *76*(1), 1–32. https://doi.org/10.18637/jss.v076.i01

Chen, C., Wakefield, J., & Lumley, T. (2014). The use of sampling weights in Bayesian hierarchical models for small area estimation. *Spatial and Spatio-Temporal Epidemiology*, *11*, 33–43. https://doi.org/10.1016/j.sste.2014.07.002

Croft, T. N., Marshall, A. M. J., & Allen, C. K. (2018). *Guide to DHS statistics*. ICF.

Farmer, P. E., Nutt, C. T., Wagner, C. M., Sekabaraga, C., Nuthulaganti, T., Weigel, J. L., et al. (2013). Reduced premature mortality in Rwanda: Lessons from success. *BMJ*, *346*, f65. https://doi.org/10.1136/bmj.f65

Fay, R. E., & Herriot, R. A. (1979). Estimates of income for small places: An application of James-Stein procedures to census data. *Journal of the American Statistical Association*, *74*(366a), 269–277. https://doi.org/10.1080/01621459.1979.10482505

GADM. (2022). *Database of global administrative areas, version 4.1*. https://gadm.org

Golding, N., Burstein, R., Longbottom, J., Browne, A. J., Fullman, N., Osgood-Zimmerman, A., et al. (2017). Mapping under-5 and neonatal mortality in Africa, 2000–15: A baseline analysis for the Sustainable Development Goals. *The Lancet*, *390*(10108), 2171–2182. https://doi.org/10.1016/S0140-6736(17)31758-0

Hodges, J. S., & Reich, B. J. (2010). Adding spatially-correlated errors can mess up the fixed effect you love. *The American Statistician*, *64*(4), 325–334. https://doi.org/10.1198/tast.2010.10052

Karema, C., Wen, S., Sidibe, A., Smith, J. L., Gosling, R., Hakizimana, E., et al. (2020). History of malaria control in Rwanda: Implications for future elimination in Rwanda and other malaria-endemic countries. *Malaria Journal*, *19*, 356. https://doi.org/10.1186/s12936-020-03407-1

Knorr-Held, L. (2000). Bayesian modelling of inseparable space-time variation in disease risk. *Statistics in Medicine*, *19*(17–18), 2555–2567. https://doi.org/10.1002/1097-0258(20000915/30)19:17/18<2555::AID-SIM587>3.0.CO;2-#

Li, Z. R., Hsiao, Y., Godwin, J., Martin, B. D., Wakefield, J., & Clark, S. J. (2019). Changes in the spatial distribution of the under-five mortality rate: Small-area analysis of 122 DHS surveys in 262 subregions of 35 countries in Africa. *PLOS ONE*, *14*(1), e0210645. https://doi.org/10.1371/journal.pone.0210645

Mayala, B., Fish, T. D., Eitelberg, D., & Dontamsetti, T. (2018). *The DHS program geospatial covariate datasets manual (second edition)*. ICF.

Mercer, L. D., Wakefield, J., Pantazis, A., Lutambi, A. M., Masanja, H., & Clark, S. (2015). Space-time smoothing of complex survey data: Small area estimation for child mortality. *The Annals of Applied Statistics*, *9*(4), 1889–1905. https://doi.org/10.1214/15-AOAS872

Moran, P. A. P. (1950). Notes on continuous stochastic phenomena. *Biometrika*, *37*(1/2), 17–23. https://doi.org/10.2307/2332142

Morris, M., Wheeler-Martin, K., Simpson, D., Mooney, S. J., Gelman, A., & DiMaggio, C. (2019). Bayesian hierarchical spatial models: Implementing the Besag York Mollié model in Stan. *Spatial and Spatio-Temporal Epidemiology*, *31*, 100301. https://doi.org/10.1016/j.sste.2019.100301

Muhoza, D. N., Rutayisire, P. C., & Umubyeyi, A. (2013). *Measuring the success of family planning initiatives in Rwanda: A multivariate decomposition analysis. DHS Working Papers no. 94* (DHS Working Papers 94). ICF International.

National Institute of Statistics of Rwanda. (2026). *Rwanda Demographic and Health Survey 2025: District profiles by province*. NISR. https://statistics.gov.rw/statistical-publications/health-nutrition-and-food-security

National Institute of Statistics of Rwanda, Ministry of Finance and Economic Planning, Ministry of Health, & ICF International. (2016). *Rwanda Demographic and Health Survey 2014–15*. NISR, MINECOFIN, MOH and ICF International. https://dhsprogram.com/pubs/pdf/FR316/FR316.pdf

National Institute of Statistics of Rwanda, Ministry of Health, & ICF. (2021). *Rwanda Demographic and Health Survey 2019–20: Final report*. NISR and ICF. https://dhsprogram.com/pubs/pdf/FR370/FR370.pdf

National Institute of Statistics of Rwanda, Ministry of Health, & ICF. (2025). *Rwanda Demographic and Health Survey 2025: Key indicators report*. NISR, MOH and ICF. https://dhsprogram.com/pubs/pdf/PR164/PR164.pdf

National Institute of Statistics of Rwanda, Ministry of Health, & ICF. (2026). *Rwanda Demographic and Health Survey 2025: Final report*. NISR and ICF. https://dhsprogram.com/pubs/pdf/FR401/FR401.pdf

National Institute of Statistics of Rwanda, Ministry of Health, & ICF International. (2012). *Rwanda Demographic and Health Survey 2010*. NISR, MOH and ICF International. https://dhsprogram.com/pubs/pdf/FR259/FR259.pdf

Pebesma, E. (2018). Simple features for R: Standardized support for spatial vector data. *The R Journal*, *10*(1), 439–446. https://doi.org/10.32614/RJ-2018-009

Rao, J. N. K., & Molina, I. (2015). *Small area estimation* (2nd ed.). Wiley. https://doi.org/10.1002/9781118735855

Ren, R. (2021). *Design-based small area estimation: An application to the DHS surveys. DHS Working Papers no. 180* (DHS Working Papers 180). ICF.

Riebler, A., Sørbye, S. H., Simpson, D., & Rue, H. (2016). An intuitive Bayesian spatial model for disease mapping that accounts for scaling. *Statistical Methods in Medical Research*, *25*(4), 1145–1165. https://doi.org/10.1177/0962280216660421

Schoumaker, B. (2013). A Stata module for computing fertility rates and TFRs from birth histories: tfr2. *Demographic Research*, *28*, 1093–1144. https://doi.org/10.4054/DemRes.2013.28.38

Schoumaker, B. (2019). Stalls in fertility transitions in sub-Saharan Africa: Revisiting the evidence. *Studies in Family Planning*, *50*(3), 257–278. https://doi.org/10.1111/sifp.12098

Simpson, D., Rue, H., Riebler, A., Martins, T. G., & Sørbye, S. H. (2017). Penalising model component complexity: A principled, practical approach to constructing priors. *Statistical Science*, *32*(1), 1–28. https://doi.org/10.1214/16-STS576

Sørbye, S. H., & Rue, H. (2014). Scaling intrinsic Gaussian Markov random field priors in spatial modelling. *Spatial Statistics*, *8*, 39–51. https://doi.org/10.1016/j.spasta.2013.06.004

Vehtari, A., Gelman, A., & Gabry, J. (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, *27*(5), 1413–1432. https://doi.org/10.1007/s11222-016-9696-4

Vehtari, A., Gelman, A., Simpson, D., Carpenter, B., & Bürkner, P.-C. (2021). Rank-normalization, folding, and localization: An improved $\hat{R}$ for assessing convergence of MCMC. *Bayesian Analysis*, *16*(2), 667–718. https://doi.org/10.1214/20-BA1221

Wakefield, J., Okonek, T., & Pedersen, J. (2020). Small area estimation for disease prevalence mapping. *International Statistical Review*, *88*(2), 398–418. https://doi.org/10.1111/insr.12400

Warren, J. L., Perez-Heydrich, C., Burgert, C. R., & Emch, M. E. (2016). Influence of demographic and health survey point displacements on point-in-polygon analyses. *Spatial Demography*, *4*(2), 117–133. https://doi.org/10.1007/s40980-015-0015-z

Westoff, C. F. (2013). The recent fertility transition in Rwanda. *Population and Development Review*, *38*(s1), 169–178. https://doi.org/10.1111/j.1728-4457.2013.00558.x

Wu, Y., Li, Z. R., Mayala, B. K., Wang, H., Gao, P., Paige, J., Fuglstad, G.-A., Moe, C., Godwin, J., Donohue, R. E., Croft, T. N., & Wakefield, J. (2021). *Spatial modeling for subnational administrative level 2 small-area estimation. DHS Spatial Analysis Reports no. 21* (DHS Spatial Analysis Reports 21). ICF.

Wu, Y., & Wakefield, J. (2025). *Small area estimation of fertility in low- and middle-income countries* (arXiv:2507.03584). arXiv. https://doi.org/10.48550/arXiv.2507.03584

Supplementary Files

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

- ESM1.pdf