

Supplementary material

Selection on Fertility, and the Environmental Decline That Would Cancel It

Dylan Slagh

Claude Opus 5

ChatGPT 5.6 Sol

Version 1.2.2 • 22 August 2026

This document holds the material a reader needs to check the main paper but does not need to read it: the engine’s validation in full, the implementation conventions, the mechanism parameter table with its sources, the historical backtest, and the reproduction path from a model run to a number on the page. Section references of the form §2 refer to the main paper.

Contents

S1 The engine, in full	1
S1.1 Validation against published outputs	1
S1.2 Implementation conventions	2
S1.3 Migration after 2100	2
S1.4 Fertility and longevity after 2100	3
S2 Mechanism parameters	3
S3 The historical backtest	5
S4 Reproducibility	5

S1 The engine, in full

S1.1 Validation against published outputs

The engine is a standard cohort-component projection: single year of age 0–100 with an open-ended final group, two sexes, 237 countries and territories, annual steps, populations dated 1 January, rates in force during the calendar year. It takes fertility, survival and net migration as given and does the bookkeeping, with no theory about the future inside it.

The typed engine of §2 is the same code with a composition axis. With every propensity set to one it reproduces the ordinary engine to a relative difference of 3×10^{-16} , and one generation of

Check at 2100, WPP 2024 zero-migration variant	Result	Tolerance
World population	0.001%	0.05%
Worst country above 10,000 people (COK)	0.13%	0.5%
Worst five-year age group below 100	0.006%	0.05%
Open-ended 100+ group	1.08%	2%
Constant fertility vs the UN’s own variant	0.05%	0.5%

Table S1: The zero-migration variant is the real test, because every input to it is published and nothing in it is estimated by us, so any discrepancy is our arithmetic being wrong. The medium variant agrees to 0.04% but is a diagnostic rather than a test: the United Nations does not publish net migrants by single year of age, so ours are backed out of its own medium path as a residual and absorb every procedural difference between us.

it matches the analytic parent–offspring covariance response to nine decimal places.

S1.2 Implementation conventions

The following source conventions are enforced explicitly because each can produce a numerically plausible but biased projection if misinterpreted.

- Ages run 0–100 where 100 is an open-ended group; sex 0 is female.
- Survival $s[a]$ is survival *into* age a , following the source’s own convention, so no translation step stands between the published table and the engine. $s[0]$ is birth survival, and $s[100]$ applies to the sum of the 99-year-olds and the existing 100+ group.
- Populations are people. The source publishes thousands; the multiplication happens exactly once, at ingest.
- Fertility is births per woman per year. The source publishes per thousand.
- Childbearing ages are 10–54. The single-age source file stops at 49 and the five-year file does not; ages 50–54 account for about 0.3% of world births and are included from the five-year file.
- Age-specific fertility uses mid-year women, so the denominator is the mean of the 1 January figure and the following one. Getting this wrong biases births by roughly half a year of cohort change — negligible in one year, not negligible after 126.
- The United Nations’ constant-fertility variant freezes 2023, not 2024: about 0.6% higher than the 2024 medium rates, compounding to 2.7% by 2100. The comparison reads the UN’s own frozen rates rather than assuming ours.
- World quantiles are computed from each draw’s own world total, never by summing country quantiles. Adding every country’s 5th percentile assumes all 236 land in their own bad tail in the same draw; at 2150 that shortcut gives 3.13 billion where the correct world 5th percentile is 6.97 billion.
- Archived WPP releases ship counterfactual scenarios beside the main projection — “no AIDS”, “instant replacement”, “zero migration”. They are excluded by name so that the backtest evaluates intended projections rather than counterfactual exercises.

S1.3 Migration after 2100

The published Bayesian migration trajectories [Azose and Raftery, 2015] run to 2100 and the public archive does not include the fitted posterior state, so the continuation refits the same

autoregressive form to the 2070–2100 portion of those trajectories and continues each path from its own 2100 rate. This is an emulator of a model’s output, not an official continuation.

Every draw-year is balanced so that world net migration is exactly zero, by shifting active-country rates by a common amount, which redistributes the discrepancy in proportion to current populations. The United Nations uses a more detailed staged procedure; this implements the essential accounting constraint and is labelled an approximation. Where a schedule would remove more people from a narrow age and sex cell than exist, the unmet removal is redistributed across populated cells rather than clipped, because clipping silently creates people.

The resulting extension gives 10.187 billion in 2100, 9.573 billion in 2125, and 8.725 billion in 2150, with a migration-only 90% range of 8.656 to 8.772 billion.

S1.4 Fertility and longevity after 2100

Fitted parameters for the continuation described in §7.2, over 2070–2100: the within-trajectory fertility autocorrelation has a median of 0.89 with a median innovation spread of 0.066 children, and the mean annual gain in female life expectancy is 0.114 years, ranging from 0.100 to 0.154 across countries.

Continued values are clipped only when they leave the full range occupied by the source archive, padded by one quarter of its width. Clipping affects $3.3\text{e-}6$ of generated values. These source-derived rails permit the observed female–male life-expectancy gap of -6.6 to $+20.2$ years and prevent ordinary source behaviour from being altered by an implicit model.

S2 Mechanism parameters

Table S2 is generated directly from the machine-readable table the model loads, so it cannot drift from the values the results were produced with. Each row carries an evidence field and a provenance field in the source file; the loader refuses to run if a row marked `sourced` carries no evidence, or if a row marked `knob` claims to have been verified.

Sources for the estimated rows

mainstream persistence

Parent–child correlation in completed family size. [Murphy \[1999\]](#) reports values mostly 0.15–0.20 in the larger recent British and American samples reviewed; [Murphy \[2013\]](#), covering 46 populations in 28 countries, reports unweighted regional means from 0.09 to 0.19. The interval 0.05–0.30 is a deliberately conservative cross-setting range rather than one study’s confidence interval. This is combined genetic and cultural continuity, not a heritability.

mainstream propensity cv

Computed from 43 pinned country files of the Cohort Fertility and Education database [[Wittgenstein Centre, 2017](#)]; among the 19 with mean completed fertility at or below 2.2 children, the country-median coefficient of variation is 0.570 and the observed range is 0.443–0.787. Independently cross-checked at 0.695 for United States cohorts 1947–1956 [[National Center for Health Statistics, 2024](#)]. See the main paper’s calibration section for the tail treatment and the declared inclusion rule.

Parameter	Kind	Value	Range	Unit
mainstream persistence	sourced	0.150	0.050–0.300	correlation
mainstream propensity cv	sourced	0.570	0.440–0.800	coefficient of variation
group retention haredi	sourced	0.867	0.735–0.950	proportion
group fertility haredi	sourced	6.400	5.800–7.000	children per woman
group share haredi	sourced	0.139	0.120–0.150	proportion of population
group retention amish	sourced	0.845	0.800–0.920	proportion
group fertility amish	sourced	6.100	5.000–6.500	children per woman
group share amish	sourced	0.00115	0.00110–0.00130	proportion of population
mainstream types	knob	3	—	count
defector high propensity weight	knob	0.600	0.300–0.900	proportion
development decline per decade	knob	0.040	0.000–0.080	proportional fall
development floor	knob	0.550	0.400–0.750	multiplier
group convergence per generation	knob	0.100	0.000–0.250	proportional fall

Table S2: All 13 mechanism parameters. Of these, 8 are estimated from published sources and checked against them, and 5 are scenario knobs with no independent support. The range column is propagated through the main paper’s parameter ensemble and boundary grid; it is a stated plausible interval across settings, not a confidence interval from a single study.

group retention/fertility/share, haredi

[Regev and Gordon \[2020\]](#) estimate that 13.3% of people aged 20–64 raised Haredi had left, implying 86.7% retention, with documented faction rates from 5.4% to 26.4% — which is what the wide interval reflects. Fertility 6.40 is the national statistical office’s own estimate of 6.38 for 2020–2022, which [Israel Democracy Institute \[2024\]](#) round to 6.4 and set against a fall from 7.5 in 2003–2005; the same report gives 1.392 million Haredim in 2024, or 13.9% of Israel’s population.

group retention/fertility/share, amish

[Anderson and Thiehoff \[2025\]](#) analyse more than 50,000 households and estimate an age-40+ retention of 84.5%, which the authors identify as the best available approximation to lifetime retention, and a total fertility rate of 6.10 from a 2002–2011 snapshot. Population share 0.115% from [Young Center \[2024\]](#), whose definition covers horse-and-buggy affiliations and excludes Beachy Amish and Amish Mennonites; the model inherits that boundary.

Why the knobs stay knobs

No literature check can convert a structural choice or an explicitly future path into an estimate. The number of mainstream types is a discretisation. The routing of group leavers has no direct measurement. The rate at which a group’s fertility advantage decays as it stops being unusual has documented instances but no established general value. And the two environment parameters describe a future that has not happened.

The last of these is the reason the main paper reports a threshold rather than a total. Fitting the environmental decline to observed fertility would consume exactly the evidence the mechanism is meant to explain, and setting it by judgement would let it control a headline it cannot support — at 2.79 billion against the mechanism’s 1.82 billion.

S3 The historical backtest

This backtest evaluates the historical performance of published United Nations projections rather than the compositional model. It is included because it informs the interpretation of published long-run intervals.

The same archive machinery that supplies the model’s inputs grades 8 archived United Nations revisions from 1992 to 2008 against the current estimates. World population was under-projected by -2.45% on average, and in the same direction in every revision since 1996. African fertility was projected -9.8% too low and East Asian fertility 14.9% too high; life expectancy was -1.32 years too low.

Two honesty constraints are enforced in the code rather than mentioned in prose. The comparison is against WPP 2024’s *estimates*, which are themselves a model output and not ground truth. And the low and high variants being compared against are fertility scenarios with no probability attached to them.

That second constraint limits what the coverage number can mean. Only 41 of 117 world-level projections (35%) fell inside the low-to-high band printed in the revision that made them. Because those bands never claimed a coverage rate, this is not a calibration failure in the technical sense, and it would be wrong to call it one. What it shows is narrower and still useful: the scenario envelopes demography has historically published around long-run population frequently failed to bracket what was later estimated, which is a reason to expect that intervals of that kind understate how much is unknown.

Grading is done at world and continental-region level rather than by country. The United Nations’ region codes mean the same thing in 1992 and now, whereas the USSR, Yugoslavia, Czechoslovakia, Sudan and Ethiopia all changed shape inside the window.

The archived releases are distributed as Excel workbooks under the “Archive” file type and use country codes compatible with the current series, permitting the automated comparisons reported here.

S4 Reproducibility

What produced each number in this paper

Every numerical value in the text, tables and figures is written by a script into a machine-readable result file, and from there into a L^AT_EX macro or a figure. No headline number is typed into the prose. The generator refuses to build if a result file is missing, if a key it expects is absent, or if a value has moved outside a stated sanity bound, on the reasoning that a paper compiling with a silently wrong number is worse than one that does not compile.

Data

Every downloaded file is checksummed and its manifest committed, so a source that is silently reissued fails loudly rather than changing a result. The inputs are World Population Prospects 2024 [[United Nations Population Division, 2024a](#)], the archived World Population Prospects revi-

Result	Produced by <code>scripts/</code>	Stored in
Engine validation	<code>validate_engine.py</code>	<code>out/validation_wpp2024</code>
Long-run diagnostics	<code>run_to_2150.py</code>	<code>out/run_to_2150</code>
Selection grid and ensemble	<code>run_phase5.py</code>	<code>out/phase5</code>
Ladder and boundary	<code>analyze_selection_</code> <code>break_even.py</code>	<code>reference/selection_</code> <code>break_even_sensitivity</code>
Paired-migration boundary	<code>analyze_paired_</code> <code>selection_boundary.py</code>	<code>reference/paired_selection_</code> <code>boundary_sensitivity</code>
Family-size dispersion	<code>analyze_cfe_dispersion.py</code>	<code>reference/mainstream_</code> <code>propensity_cv</code>
Project extension	<code>run_un_extension.py</code>	<code>reference/un_project_</code> <code>extension_summary</code>
Uncertainty decomposition	<code>decompose_uncertainty.py</code>	<code>out/uncertainty_decomposition</code>
Conventional comparator	<code>run_uw_ensemble.py</code>	<code>out/uw_ensemble</code>
Historical backtest	<code>run_backtest.py</code>	<code>out/backtest</code>
This paper’s macros and tables	<code>build_paper_results.py</code>	<code>paper/generated/</code>
This paper’s figures	<code>plot_paper_figures.py</code>	<code>paper/figures/</code>

Table S3: The path from a model run to a number on the page. Stored results are JSON; `reference/` abbreviates `data/reference/`, which is committed to the repository, while `out/` is regenerable and is not.

sions used for the backtest [United Nations Population Division, 2024b], the published Bayesian fertility, mortality and migration trajectories [BayesPop Project, 2024, Azose and Raftery, 2015], the Cohort Fertility and Education database [Wittgenstein Centre, 2017], and the national cohort fertility tables [National Center for Health Statistics, 2024].

Two data sets are not redistributed. The cohort fertility tabulations are used under terms that do not permit republication, so the repository commits the fetching script, the checksums, the derived per-country statistics and the analysis — everything needed to reproduce the number — but not the source tables. The large trajectory archives are likewise fetched rather than committed.

Implementation conventions

- Ages run 0–100 where 100 is an open-ended group; sex 0 is female.
- Survival $s[a]$ is survival *into* age a , following the source’s own convention, so that no translation step stands between the published table and the engine. $s[0]$ is birth survival, and $s[100]$ applies to the sum of the 99-year-olds and the existing 100+ group.
- Populations are people. The source publishes thousands; the multiplication happens exactly once, at ingest.
- Fertility is births per woman per year. The source publishes per thousand.
- Childbearing ages are 10–54. The single-age source file stops at 49 and the five-year file does not; ages 50–54 account for about 0.3% of world births and are included from the five-year file.
- Age-specific fertility uses mid-year women, so the denominator is the mean of the 1 January figure and the following one. Getting this wrong biases births by roughly half a year of cohort change — negligible in one year, not negligible after 126.
- World quantiles are computed from each draw’s own world total, never by summing country quantiles. Adding every country’s 5th percentile assumes all of them land in their own bad tail in the same draw; at 2150 that shortcut gives 3.13 billion where the correct world 5th

percentile is 6.97 billion.

Stored predictions

Projections intended to be scored later are written as immutable dated records with their own checksums; the writer raises rather than overwrites an existing record and has no override flag. Every quantity in the conventional comparator is marked as a comparator rather than as a claim of this paper.

References

- Cory Anderson and Stephanie Thiehoff. The population structure of the Amish, a rapidly growing ethnic religion in North America. *Population Studies*, 2025. doi: 10.1080/00324728.2025.2592576.
- Jonathan J. Azose and Adrian E. Raftery. Bayesian probabilistic projection of international migration. *Demography*, 52(5):1627–1650, 2015. doi: 10.1007/s13524-015-0415-0.
- BayesPop Project. BayesPop probabilistic population projection data downloads, 2024. URL <https://bayespop.csss.washington.edu/download/>. University of Washington Center for Statistics and the Social Sciences; accessed 9 August 2026.
- Israel Democracy Institute. Annual statistical report on Haredi society in Israel 2024. Technical report, Israel Democracy Institute, Jerusalem, 2024. URL <https://en.idi.org.il/media/27532/idi-annual-statistical-report-on-haredi-society-2024.pdf>.
- Michael Murphy. Is the relationship between fertility of parents and children really weak? *Social Biology*, 46(1–2), 1999. doi: 10.1080/19485565.1999.9988991.
- Michael Murphy. Cross-national patterns of intergenerational continuities in childbearing in developed countries. *Biodemography and Social Biology*, 59(2), 2013. doi: 10.1080/19485565.2013.833779.
- National Center for Health Statistics. Cohort fertility tables, 2024. URL https://www.cdc.gov/nchs/nvss/cohort_fertility_tables.htm. Tables 2 and 3; accessed 13 August 2026.
- Eitan Regev and Gabriel Gordon. Transitions between religious groups among Israeli Jews. Israel Democracy Institute, 2020. URL <https://en.idi.org.il/articles/32775>. Accessed 13 August 2026.
- United Nations Population Division. World Population Prospects 2024: Methodology of the United Nations population estimates and projections. Technical report, United Nations, New York, 2024a. URL https://population.un.org/wpp/Publications/Files/WPP2024_Methodology-Report_Final.pdf. Department of Economic and Social Affairs.
- United Nations Population Division. World Population Prospects: Download centre and revision archive, 2024b. URL <https://population.un.org/wpp/downloads>. Accessed 9 August 2026.
- Wittgenstein Centre. Cohort Fertility and Education database, 2017. URL <https://www.eurrep.org/database/database/>. Wittgenstein Centre for Demography and Global Human Capital; data accessed 13 August 2026 under the database’s terms of use.
- Young Center. Amish population, 2024, 2024. URL <https://groups.etown.edu/amishstudies/population-2024/>. Young Center for Anabaptist and Pietist Studies, Elizabethtown College; accessed 13 August 2026.