Composition Beats Collapse: Insights from the Bisin–Verdier Model on Endogenous Fertility Reversal

By Sebastian Galiani and Raul A. Sosa*

Fertility rates have fallen below replacement in most countries, fueling predictions of demographic collapse. We show these forecasts overlook a crucial fact: societies are not homogeneous. Using the Bisin–Verdier model of cultural transmission with endogenous fertility and direct socialization, calibrated to U.S. and global data, we find that high-fertility, high-retention groups persist, gain share, and lead the total population to grow. Even if fertility remains below replacement in every country, extinction is unlikely. Simulations imply continued growth with pronounced compositional change, driven especially by religious communities with high fertility. In our ten-generation world calibration, Muslims become the largest tradition. (JEL J11, J13, Z12)

Across much of the world, fertility has fallen below replacement. From this fact, researchers often draw a stark conclusion: if fertility remains below replacement, population would eventually converge to zero. That inference is arithmetically correct for a homogeneous, stationary population with fixed behavior. It is a poor guide, however, for heterogeneous societies in which cultural traits are transmitted across generations and correlate with repro-

* Galiani: Department of Economics, University of Maryland (sgaliani@umd.edu); Sosa: Department of Economics, Universidad de San Andrés (rsosa@udesa.edu.ar). We thank Alberto Bisin and Thierry Verdier for their insightful comments and suggestions.

ductive behavior.

Our starting point is a straightforward evolutionary observation—familiar since Darwin (1859)—that traits linked to above-replacement reproduction, when they are successfully transmitted, tend to persist and gain population share. We apply that intuition to *cultural* traits—most saliently, religious affiliation—that parents transmit to children and that shape fertility choices. In such environments, the relevant question is not whether *aggregate* fertility sits below replacement, but whether some groups sustain above-replacement reproduction while transmitting identity strongly to their descendants. If so, composition-driven dynamics can overturn extinction logic that rests on aggregates.

We build on the economics of cultural transmission (Bisin and Verdier, 2000, 2001) and on the cultural-evolution tradition modeling vertical, oblique, and horizontal channels (Cavalli-Sforza and Feldman, 1981; Boyd and Richerson, 1985). On religion and fertility, we connect to demographic evidence linking religiosity, identity retention, and higher fertility (McQuillan, 2004; Lehrer, 2004); to empirical work on how families invest in transmitting identity (Bisin, Topa and Verdier, 2004); and to formal treatments that embed religious transmission into fertility dynamics (Bar-El et al., 2013). Our contribution is to show, through unified U.S. and world calibrations, how composition-driven dynamics can overturn the extinction logic even when sub-replacement fertility is widespread across countries.

We operationalize the mechanism within the Bisin–Verdier framework (Bisin and Verdier, 2001), where a child's identity reflects *direct socialization* in the family and *oblique* exposure to the social environment. We adopt the *cultural substitution* property emphasized there: as a group becomes

more common, the optimal intensity of within-family transmission tapers off. This preserves stable heterogeneity because minorities invest more in transmission and majorities invest less. We then couple the transmission block to group-specific reproduction and endogenize fertility through the same within-family choice, so group sizes and shares coevolve within a single mechanism. As high-fertility, high-retention groups gain weight, overall population growth becomes positive—even if many groups, or many countries, remain below replacement at a point in time. If no group is above replacement, the aggregate declines, but heterogeneity still reallocates shares toward relatively higher-fertility, higher-retention groups, slowing the decline relative to a homogeneous benchmark.

We evaluate this mechanism in two empirical settings. First, for the United States we run identical simulations under two alternative partitions. Grouping by race/ethnicity (White non-Hispanic, Black non-Hispanic, Hispanic, Asian non-Hispanic, American Indian/Alaska Native non-Hispanic, and Native Hawaiian/Other Pacific Islander non-Hispanic) yields reproduction profiles that cluster tightly below replacement, so the model projects an aggregate decline. Regrouping the very same population by religion (Evangelical Protestant, Mainline Protestant, Catholic, Black Protestant, Jewish, Muslim, Mormon (LDS), Other non-Christian, and Unaffiliated) reveals markedly greater dispersion in both reproduction and direct-socialization schedules; high-fertility, high-retention affiliations gain share and put the population on a growth path. A guiding principle follows: the finest partition that preserves coherent group traits best reveals the behaviorally meaningful *heterogeneity* that drives composition-driven dynamics.

Second, we apply the model to the world, allocating the population across

seven broad religious traditions (Christian, Muslim, Unaffiliated, Hindu, Buddhist, Folk/Other, Jewish). Dispersion across traditions in both fertility and identity transmission is even greater globally. Traditions that combine high fertility with strong transmission—most notably Muslims—expand persistently, and the aggregate grows even though fertility has fallen below replacement in most countries. The message is the same at a larger scale: long-run population trajectories are governed by compositional change rather than by under-replacement averages.

Our objective is not to forecast population paths. Parameters are held fixed by design, and the baseline abstracts from homophily, migration, age structure, and changing mortality. These simplifications isolate the mechanism: once at least one group remains above replacement and transmits identity effectively, aggregate collapse is not the dominant trajectory even when fertility lies below replacement. What *does* change is composition. Over generational time, high-fertility groups that transmit identity become increasingly represented, and the overall trajectory drifts upward toward the pace set by those groups. In extensions that add homophily and migration, the qualitative message strengthens: homophily amplifies within-group transmission by reducing out-group exposure, and migration shifts composition in proportion to the inflow, but the core result remains.

The rest of the paper proceeds as follows. Section I presents the framework based on Bisin and Verdier (2001). Section II describes the data and the calibration of initial shares, reproduction proxies, and direct-socialization targets. Section III reports the baseline results (United States by race/ethnicity and by religion; world by religion) and the extensions with homophily and migration. Section IV concludes.

I. Framework

We study a culturally heterogeneous population partitioned into groups $i = 1, \ldots, K$ (e.g., religious traditions or ethnoracial categories). Let $N_i(t)$ denote the number of *adult women* in group i at generation $t \in \{0, 1, \ldots\}$, and let

$$q_i(t) \equiv \frac{N_i(t)}{\sum_j N_j(t)} \tag{1}$$

be the group's population share (female line).

Identity transmission

Following Bisin and Verdier (2001), a newborn's identity is shaped by two channels: direct (vertical) socialization in the family and oblique exposure to the social environment. Group-i parents choose an effort $d_i(t) \in [0, 1)$; with probability $d_i(t)$ the child retains the parental identity, and with probability $1 - d_i(t)$ she imitates a role model drawn from the environment faced by group-i families.

We model the oblique environment using a nonnegative mixing (contact) kernel $G = \{G_{ij}\}_{i,j=1}^{K}$. From the perspective of a group-i family, the probability to encounter a role model of type j is

$$\pi_{i\to j}(t) = \frac{G_{ij}\, q_j(t)}{\sum_k G_{ik}\, q_k(t)}, \qquad \sum_j \pi_{i\to j}(t) = 1, \tag{2}$$

and the probability that the oblique draw matches the parent's own identity

is

$$q_{\mathrm{env},i}(t) \;\equiv\; \pi_{i\to i}(t) \;=\; \frac{G_{ii}\, q_i(t)}{\sum_k G_{ik}\, q_k(t)} \in [0,1]\,. \tag{3}$$

The induced parent-origin transition matrix $P(t)$ is row-stochastic with elements

$$P_{i\to j}(t) = \begin{cases} d_i(t) \;+\; (1-d_i(t))\,\pi_{i\to i}(t), & j = i, \\ (1-d_i(t))\,\pi_{i\to j}(t), & j \neq i. \end{cases} \tag{4}$$

This decouples environmental exposure (G, hence $\pi_{i\to\cdot}$) from the within-family retention choice d_i and nests the random-matching benchmark as $G_{ij} \equiv 1$ (then $\pi_{i\to j} = q_j$ and $q_{\mathrm{env},i} = q_i$). Under the Bisin–Verdier *cultural substitution* property—$d_i(\cdot)$ strictly decreasing in own-group exposure and $d_i(1) = 0$—the transmission block preserves heterogeneity (minorities optimally socialize more, majorities less).[1]

Aggregate population dynamics

Let $n_i(t)$ denote group i's per-generation replacement factor—the expected number of next-generation adult women produced by one adult woman[2]. Births from group i at time t equal $N_i(t)\, n_i(t)$. Newborns are assigned identities via $P(t)$, yielding the law of motion in adult-female lev-

[1]This formulation is closely connected to network-based models in which paternalistic vertical socialization and homophily jointly shape transition probabilities; see Panebianco and Verdier (2017). Our diagonal tilt $G_{ii} = \zeta_i$ can be read as a reduced-form "inbreeding–homophily" tilt in the sense of Currarini, Jackson and Pin (2009); the uniform case in the baseline nests as $\zeta_i \equiv \zeta$.

[2]So $n_i(t) = 1$ is replacement on the female line; units: "adult women per adult woman per generation".

els:

$$N(t+1) \;=\; \big[N(t) \odot n(t)\big]\; P(t) \;+\; M(t)\,, \tag{5}$$

where $N(t)$ and $n(t)$ are $1 \times K$ row vectors and $\odot$ denotes element-wise multiplication. For migration, we use

$$M(t) \;=\; \mu(t)\,\Big(\sum_{j} N_j(t)\Big)\, m(t)\,, \qquad m(t) \in \Delta^{K-1}\,, \tag{6}$$

with $\mu(t)$ net adult inflow per host adult per generation and $m(t)$ the arrivals' composition by trait. Total adult-female population therefore evolves according to

$$\sum_{j} N_j(t+1) \;=\; \Big(\bar{n}(t) + \mu(t)\Big) \sum_{j} N_j(t)\,, \qquad \bar{n}(t) \;\equiv\; \sum_{i} q_i(t)\, n_i(t)\,. \tag{7}$$

When $\bar{n}(t) > 1$ the aggregate of adult women grows; when $\bar{n}(t) < 1$ it declines. For display we map adult-female levels to *total population* with a fixed two-sex factor:

$$N^{\text{total}}(t) \;\approx\; \frac{N(t)}{\pi_f}\,, \tag{8}$$

where π_f is the initial female share among adults. This constant rescaling affects only the vertical scale in levels panels and leaves group shares and all comparative dynamics unchanged.

Within-family choice: direct socialization and fertility.

We endogenize both $d_i(t)$ and $n_i(t)$ via a within-family problem (units of $n_i(t)$: adult women per adult woman per generation). Let the identity "match value" per child be

$$E_i\big(d; q_{\text{env},i}\big) \;\equiv\; \big[d + (1-d)\, q_{\text{env},i}\big]\Delta V\,, \tag{9}$$

where $\Delta V > 0$ captures imperfect empathy/paternalistic altruism (parents evaluate the child's future actions with their own preferences), as in Bisin and Verdier (2001). Direct socialization entails cost $H_i(d) = \frac{1}{2}\kappa_i d^2$[3] and child quantity has convex cost $c_i(n) = \frac{1}{2} b_i n^2$. Parents of type i choose (n, d) to maximize

$$\max_{n\geq 0,\; d\in[0,1)} U_i(n, d; q_{\text{env},i}) \;=\; n(E_i(d; q_{\text{env},i}) - H_i(d)) \;-\; \tfrac{1}{2}\, b_i\, n^2\,. \tag{10}$$

The within-family first-order conditions (for interior solutions) are

$$H_i'(d) \;=\; (1 - q_{\text{env},i})\,\Delta V\,, \qquad c_i'(n) \;=\; E_i(d; q_{\text{env},i}) - H_i(d)\,. \tag{11}$$

With the quadratic baseline this gives the closed-form policies

$$\begin{aligned} d_i^\star(q_{\text{env},i}) &= \frac{(1 - q_{\text{env},i})\,\Delta V}{\kappa_i}\,, \quad \text{truncated to } [0, 1); \\ n_i^\star &= \frac{E_i(d_i^\star) - H_i(d_i^\star)}{b_i}\,, \quad \text{truncated at } 0. \end{aligned} \tag{12}$$

[3]Quadratic costs imply a linear direct-socialization rule in the Bisin–Verdier first-order condition—hence $d_i^\star$ falls linearly with own-group exposure. This is our baseline (and aligns with Bisin and Verdier 2001). For robustness, we also implemented the concave schedule in Bar-El et al. (2013), which makes $d_i^\star$ decreasing and concave in exposure; qualitative dynamics and headline results are unchanged (see Supplemental Appendix).

Two properties drive the dynamics. First, $d_i^\star$ is strictly decreasing in own-group exposure $q_{\text{env},i}$, implementing cultural substitution (minorities socialize more, majorities less). Second, for given $q_{\text{env},i}$, higher identity value ΔV or lower costs (κ_i, b_i) raise $n_i^\star$. Because next period's levels are $[\, N(t) \odot n^\star(t) \,] \, P(t)$, relative growth tilts toward groups with larger $n_i^\star(t)$, and $\bar{n}(t) = \sum_i q_i(t)\, n_i^\star(t)$ moves endogenously with composition. We set $n_i^\star = 0$ whenever $E_i(d_i^\star) \leq H_i(d_i^\star)$.

The framework nests the Bisin–Verdier transmission matrix, preserves their core stability logic for the transmission block under cultural substitution, and augments it with a demography block that tracks population levels alongside shares and an endogenous fertility margin tied to the same within-family decision that governs direct socialization. This coupling shows why, once any group sustains sufficiently high identity-weighted reproduction, the aggregate replacement factor can rise even if many groups lie below replacement at a point in time.[4]

II. Data and calibration

For each group i we specify (i) an initial adult share q_{i0}; (ii) a replacement factor on-impact target n_{i0} built from observed fertility with simple demographic adjustments; and (iii) a baseline direct-socialization target $d_{i0} \equiv d_i(q_{i0})$. Given a common ΔV (scale-normalization $\Delta V = 1$), we back out the cost parameters (κ_i, b_i) from the closed-form first-order conditions implied by the within-family problem. Parameters are held fixed over a ten-generation horizon by design; age structure and changing mortality are

[4]Bisin and Verdier (2001) analyze cultural transmission with endogenous direct socialization and show global stability of heterogeneous stationary distributions under cultural substitution; our contribution is to embed their matrix in a levels-tracking demography with an endogenous fertility margin and to study the induced aggregate dynamics.

omitted to isolate the transmission mechanism. The baseline sets random matching in the cultural pool ($G_{ij} \equiv 1$, hence $\pi_{i\to j} = q_j$) and no migration ($\mu(t) \equiv 0$).

Reproduction proxy.

We construct the replacement factor on-impact target as

$$n_{i0} \;=\; \text{Fertility}_i \times \phi_i \times s_i \,, \tag{13}$$

where Fertility_i is either total fertility rate (TFR) or completed fertility (CEB), ϕ_i is the female share at birth implied by the sex ratio at birth (SRB), and s_i is female survivorship to (approximately) the end of childbearing. For the U.S. religion partition we use CEB for women ages 40–59 and set $n_{i0} = \text{CEB}_i \times \phi$ (a GRR-style proxy with no survivorship term).[5]

United States: race/ethnicity.

We use the 2020 Census six-way partition: non-Hispanic White, non-Hispanic Black, Hispanic, non-Hispanic Asian, non-Hispanic American Indian/Alaska Native (AIAN), and non-Hispanic Native Hawaiian/Other Pacific Islander (NHPI). Group-specific fertility and SRB come from *National Vital Statistics Reports: Births—Final Data, 2023* (NCHS). Female survivorship to the end of childbearing is taken from *United States Life Tables, 2021* (NCHS), using the most disaggregated tables available by race/ethnicity; for NHPI, where female life tables are not published, we use the national

[5]The proxy follows the female line for tractability. Counting both sexes amounts to a fixed two-sex mapping that rescales levels but does not affect shares or dynamics. Our results hinge on relative differences across i, and the scale is absorbed in the calibration of b_i.

female proxy. Baseline direct-socialization targets d_{i0} are aligned to Pew Research Center measures of intermarriage among newlyweds (2015), with conservative proxies for very small groups where direct coverage is limited.

United States: religion.

We form nine buckets—Evangelical Protestant, Mainline Protestant, Catholic, Black Protestant, Jewish, Muslim, Mormon (LDS), Other non-Christian, and Unaffiliated. Initial shares are renormalized from Pew's *Religious Landscape Study* (2014) over these nine buckets, excluding small Christian categories (Orthodox, Jehovah's Witness, Other Christian). Fertility is *completed fertility* (CEB, children ever born) for women ages 40–59: RLS 2014/2015 for most groups ("non-Christian faiths" 1.8 used for Other non-Christian), Muslims from Pew 2017, and Jews from Pew 2021. We apply a common female-at-birth share $\phi = 0.488$ implied by the 2023 national SRB and, for this block, set $n_{i0} = \text{CEB}_i \times \phi$. Direct-socialization targets draw on Pew evidence on endogamy and on "raised versus current" religion, mapped consistently to the nine buckets.

World: religion.

Countries are aggregated into seven traditions—Christian, Muslim, Unaffiliated, Hindu, Buddhist, Folk/Other, and Jewish. Initial shares, group fertility (TFR 2010–2015), and retention proxies are taken from Pew's *The Future of World Religions*. To harmonize inputs across countries, we apply a common female-at-birth share $\phi = 0.488$ implied by the global SRB and a single female survivorship to age 50 $s = 0.964$, both from the UN *World Population Prospects 2024* life tables.

Calibrating κ_i and b_i.

We take one model period to be a generation (25–30 years) and simulate $T = 10$ generations. Baseline runs set random matching in the cultural pool ($G_{ij} \equiv 1$, so $\pi_{i\to j} = q_j$ and $q_{\text{env},i} = q_i$) and no migration ($\mu(t) \equiv 0$). For each application we calibrate (κ_i, b_i) *on impact* to match an initial direct-socialization target d_{i0} and a replacement factor n_{i0}. With $q_{\text{env},i}(0) = q_{i0}$ and a common scale ΔV (we set $\Delta V = 1$ for normalization), the within-family first-order conditions imply

$$\kappa_i \;=\; \frac{(1 - q_{i0})\,\Delta V}{d_{i0}} \tag{14}$$

$$b_i \;=\; \frac{[d_{i0} + (1 - d_{i0})\,q_{i0}]\,\Delta V - \frac{1}{2}\kappa_i d_{i0}^2}{n_{i0}} \tag{15}$$

Here, κ_i is pinned down by the intensity of direct socialization required to offset environmental exposure at q_{i0}; b_i matches the observed reproduction proxy once the identity benefit per child and the socialization cost per child are accounted for. After this on-impact match, the primitives $(\kappa_i, b_i, \Delta V)$ are held fixed, and $\{d_i(t), n_i(t)\}$ evolve endogenously with $q_{\text{env},i}(t)$ as determined by G and the path of $q(t)$. Extensions alter only exposure and composition: in homophily experiments we keep $G_{ij} = 1$ for $j \neq i$ and introduce group-specific diagonal weights $G_{ii} = \zeta_i^{\text{eff}} \in [1, 2]$ (nesting random matching at $\zeta_i^{\text{eff}} = 1$).

Cultural mixing kernel (homophily/segregation).

We calibrate the oblique environment using the mixing kernel G introduced in Section I. From the perspective of group i, the probability to encounter a role model of type j is given by equation (2), and the own–type exposure entering the family problem is $q_{\mathrm{env},i}(t) = \pi_{i\to i}(t)$ in equation (3). In homophily runs we set $G_{ij} = 1$ for $j \neq i$ and $G_{ii} = \zeta_i$, where the diagonal entries are group-specific and calibrated on impact to match a data anchor for own-type exposure: $q_{\mathrm{env},i}(0) \;=\; \tilde{q}_{\mathrm{env},i}$.

With $q_{\mathrm{env},i}(0) = \dfrac{\zeta_i\, q_{i0}}{\zeta_i\, q_{i0} + \sum_{k\neq i} q_{k0}}$, the implied inversion is

$$\widehat{\zeta}_i \;=\; \frac{\tilde{q}_{\mathrm{env},i}\left(1 - q_{i0}\right)}{q_{i0}\left(1 - \tilde{q}_{\mathrm{env},i}\right)}. \tag{16}$$

Anchors $\tilde{q}_{\mathrm{env},i}$ are chosen to be directly observed, one-number measures of own-type exposure: *U.S. religion* uses the share of currently married adults with a same-religion spouse for each bucket (Pew RLS; most recent published breakouts, with clearly flagged proxies for sub-traditions lacking separate publication), while the *World* partition uses the share of adherents who live in countries where their religion is the majority (2020). When a bucket lacks a published anchor (*U.S. Other non-Christian; World Folk/Other*), we set $\zeta_i = 1$ as a conservative default. We hold fixed the primitives $(\Delta V, \kappa_i, b_i)$ calibrated on impact; the paths $\{d_i(t), n_i(t)\}$ then evolve endogenously as $\{q_{\mathrm{env},i}(t)\}$ change with G and composition.

We do not use the odds-ratio tilt $\widehat{\zeta}_i$ directly. To keep the strength of homophily modest and comparable across partitions while preserving ranks, we apply a monotone, multiplicative (log) compression into the invariant

interval $[1, 2]$:

$$\zeta_i^{\text{eff}} \;=\; 1 \;+\; \frac{\log\big(\max\{1, \widehat{\zeta}_i\}\big)}{\log\Big(\max_j \max\{1, \widehat{\zeta}_j\}\Big)}\,\big(2-1\big) \;\in\; [1, 2]. \tag{17}$$

This mapping sends the random-matching benchmark to 1 and the largest observed $\widehat{\zeta}_i$ in the sample to 2, with all other groups placed proportionally in between on a log scale. When an anchor is missing or invalid we set $\widehat{\zeta}_i = 1$ (hence $\zeta_i^{\text{eff}} = 1$).

Migration.

On the adult-female line we model net inflows as in equation (6) of Section I. Here $\mu(t)$ is the net inflow per host adult woman per generation and m is the arrivals' composition by religion. We set the level so that the implied *two-sex* inflow at $t = 0$ equals 22.5 million per generation (0.9 million/year $\times$ 25): operationally, $M(0) = \pi_f \times 22.5$ million, with π_f the female share among adults used for our two-sex display rescaling. The composition m follows Pew's *Faith on the Move* (mapped to our nine buckets and renormalized; see Table 1).

Table 1 lists the empirical inputs used to construct the on-impact replacement factor n_{i0} and to calibrate (κ_i, b_i): initial shares q_{i0}; demographic inputs $\big(\text{Fertility}_i, \phi_i, s_i\big)$—TFR for U.S. race/ethnicity and the world; CEB for U.S. religion; and baseline d_{i0}. For the U.S. religion partition we also report the migration composition used in one of the migration scenarios.

III. Results

We apply the framework described in Section I to compare simulations across different partitions and scales. Each figure is organized into four panels: Panel A presents total population $N^{\text{total}}(t)$ (in logs), obtained by rescaling adult-female levels $N(t)$ by a fixed female share π_f; Panel B shows group populations (in logs); Panel C depicts group shares; and Panel D plots the aggregate replacement factor $\bar{n}(t)$. Across all exercises, parameters are held constant over ten generations by construction, allowing us to isolate how fertility choices and cultural transmission interact with population composition.

A. United States

Race/ethnicity

We begin with the United States to highlight how the choice of grouping dimension changes the aggregate forecast. Partitioning by race/ethnicity produces small cross-group dispersion in reproduction and in the direct-socialization schedules, so composition moves only modestly. The aggregate replacement factor $\bar{n}(t)$ remains below one throughout and $N^{\text{total}}(t)$ declines over the horizon (Figure 1, Panel A and Panel D). White non-Hispanics remain the largest group by share, Hispanics gain modestly, and the other groups shrink in share at varying speeds (Panel C), but the gaps in $n_i(t)$ are not large enough to overturn the under-replacement aggregate. In line with the model, direct-socialization intensity falls with a group's own exposure, further dampening divergence.

Religion

Regrouping the same U.S. population by religion reveals markedly greater heterogeneity. Groups differ both in their reproduction profiles and in their $d(q)$ schedules, so the endogenous response of direct socialization to group size becomes central. In the simulation, Mormons (LDS) expand rapidly in both headcount and population share. Evangelical Protestants and Catholics also grow in absolute numbers, but their shares remain broadly stable. By contrast, the Unaffiliated and Mainline Protestants not only fail to expand in levels but also experience substantial declines in share (Figure 2, Panels B and C). The aggregate replacement factor $\bar{n}(t)$ rises as weight shifts toward high-fertility, high-retention affiliations (Panel D). In short, race/ethnicity compresses heterogeneity and points to decline, whereas religion exposes heterogeneity and yields growth. This comparison underscores a key principle: the finest partition that preserves coherent group traits is the most informative, while coarser groupings mute composition-driven dynamics.

B. World

Aggregating countries into seven broad religious traditions delivers the same mechanics at larger scale, with starker heterogeneity. Muslims expand sharply in both levels and share; Christians and Hindus grow considerably in headcount but increase only slightly in share. The Unaffiliated and other low-fertility traditions, such as the Buddhist, remain roughly stable in headcount but decline in share, most markedly for the Unaffiliated (Figure 3, Panel B and Panel C). Because the spread in $n_i(t)$ and in $d(q)$ is wider globally, the share reallocation is more pronounced. The aggregate replacement

factor $\bar{n}(t)$ stays above one and climbs over time (Panel D), implying a world that drifts away from demographic collapse and experiences compositional change as the religiously affiliated share rises.

C. Extensions: homophily (cultural mixing) and migration

We extend the baseline along two blocks, *homophily* and *net migration*, holding the calibrated cost parameters (κ_i, b_i) and the common scale ΔV fixed. We overlay four scenarios—baseline, + homophily, + migration, and + both. All runs in this subsection use the U.S. religion partition; baseline primitives are unchanged, so differences reflect only the cultural pool and the inflow block. Levels panels map adult-female counts to total population with the same fixed two-sex factor as in the baseline.

Homophily

Oblique exposure is governed by the mixing kernel G with $G_{ij} = 1$ for $j \neq i$ and group-specific diagonals $G_{ii} = \zeta_i^{\text{eff}} \in [1, 2]$ calibrated from directly observed own-type exposure using (16) and (17). Because higher G_{ii} raises $q_{\text{env},i}$, the first-order conditions imply lower direct socialization $d_i^\star$ and higher fertility $n_i^\star$, so the aggregate replacement factor $\bar{n}(t) = \sum_i q_i(t)\, n_i^\star(t)$ lies above the baseline path.

Migration

Net inflows enter as $M(t) = \mu \left(\sum_j N_j(t)\right) m$, with constant μ and fixed composition m mapped from *Faith on the Move*. We calibrate μ so that the *two-sex* inflow at $t = 0$ equals 22.5 million per generation, i.e., $M(0) = \pi_f \times$ 22.5 million. The population growth factor on the female line decomposes

as

$$g(t) \;=\; \frac{\sum_j N_j(t+1)}{\sum_j N_j(t)} \;=\; \bar{n}(t) + \mu, \tag{18}$$

so migration shifts levels by μ and affects $\bar{n}(t)$ only through composition. In our U.S. religion runs (Figure 4) the inflow lifts levels (Panel A) but generates a negative composition effect on $\bar{n}(t)$ in subsequent generations (Panel D). The qualitative ranking of groups is unchanged.

IV. Conclusion

Aggregate under-replacement fertility need not imply population collapse. In a Bisin–Verdier framework with endogenous fertility and direct socialization, group sizes and reproduction coevolve; whenever at least one group remains above replacement on the female line and transmits identity effectively, its share rises and turns the aggregate path upward.

In the United States, the choice of partition is decisive. Grouping by race/ethnicity, where reproduction factors cluster around a common benchmark, yields an aggregate decline. Grouping by religion, where fertility and retention differ more, reallocates shares toward high-fertility, high-retention groups and places the aggregate on a growth path.

At the global scale, dispersion is wider still. Traditions with both high fertility and strong identity transmission expand while low-fertility groups lose share, so the global composition becomes more religious. Although the objective is not to forecast outcomes for particular groups, our world simulations imply not only a more religious composition but also that, within the horizon we study, Muslims become the largest tradition by share.

Extensions reinforce the mechanism. Allowing homophily raises the aggre-

gate replacement factor and accelerates reallocation; migration shifts levels and shares but does not overturn the qualitative dynamics.

Across settings, the lesson is the same: partition choice is first-order. Using the finest partition that preserves coherent group traits best reveals the heterogeneity that drives composition-driven dynamics. Long-run outcomes reflect aggregate growth with pronounced compositional change—*composition beats collapse*.

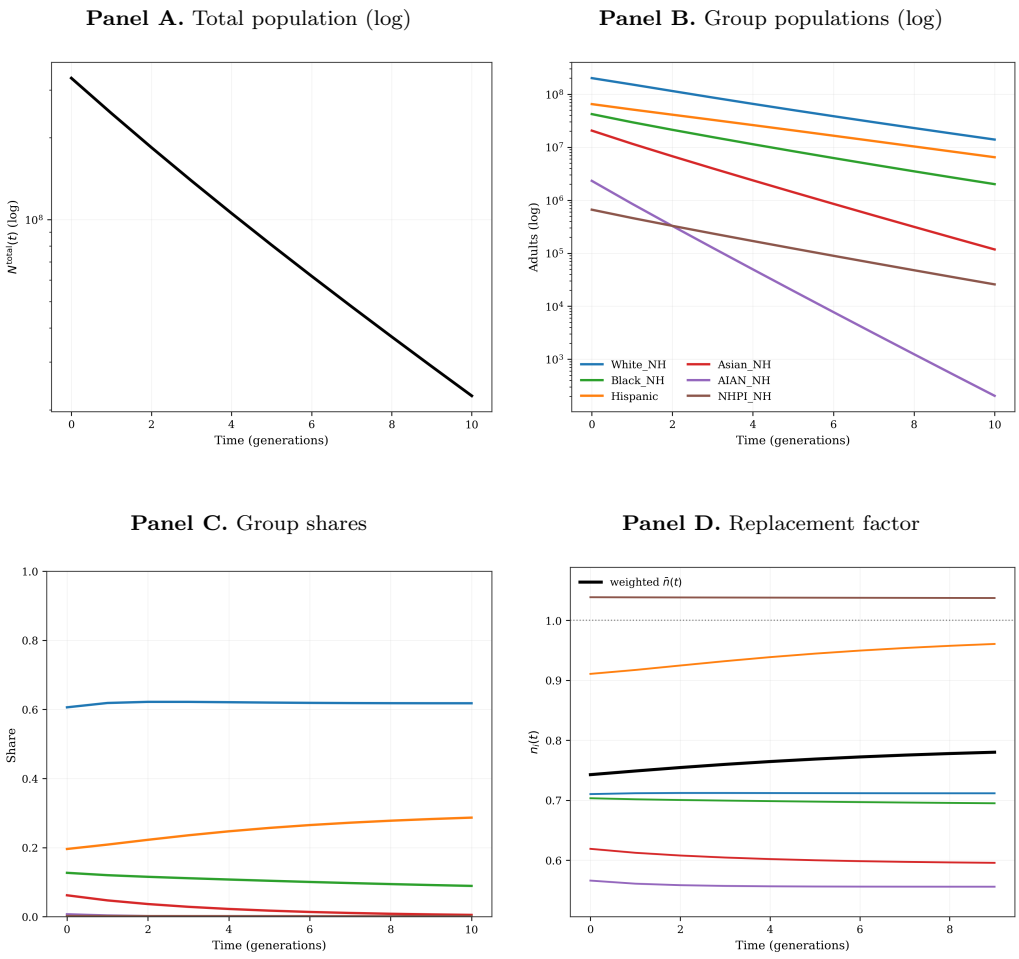

Figure 1. United States by Race/Ethnicity

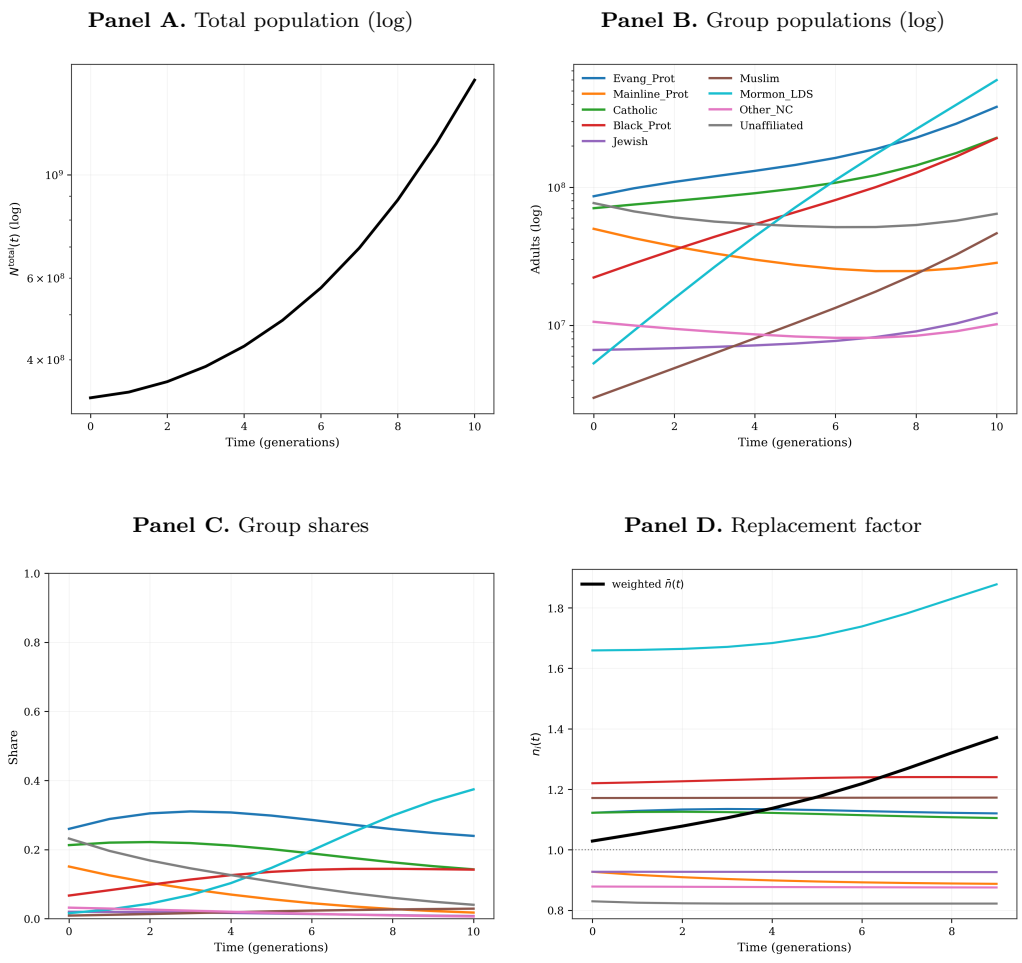

Figure 2. United States by Religion

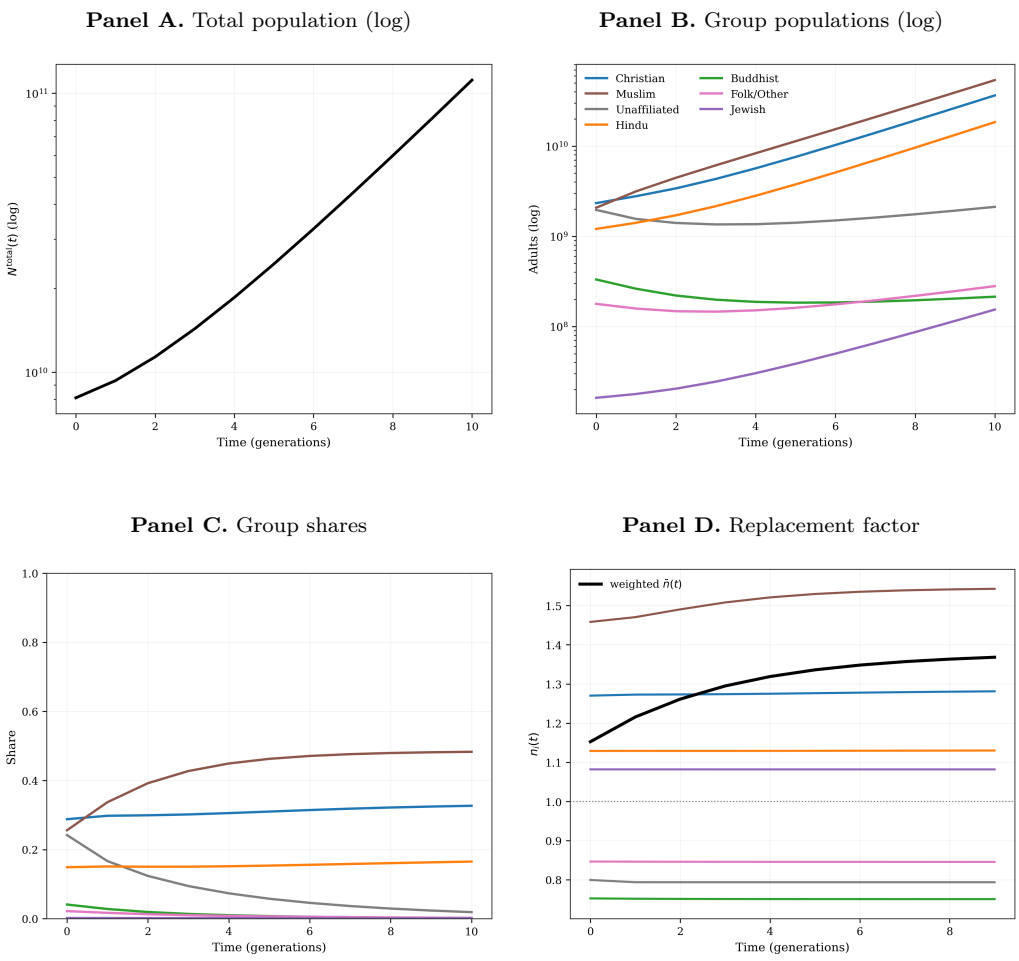

Figure 3. World by Religion

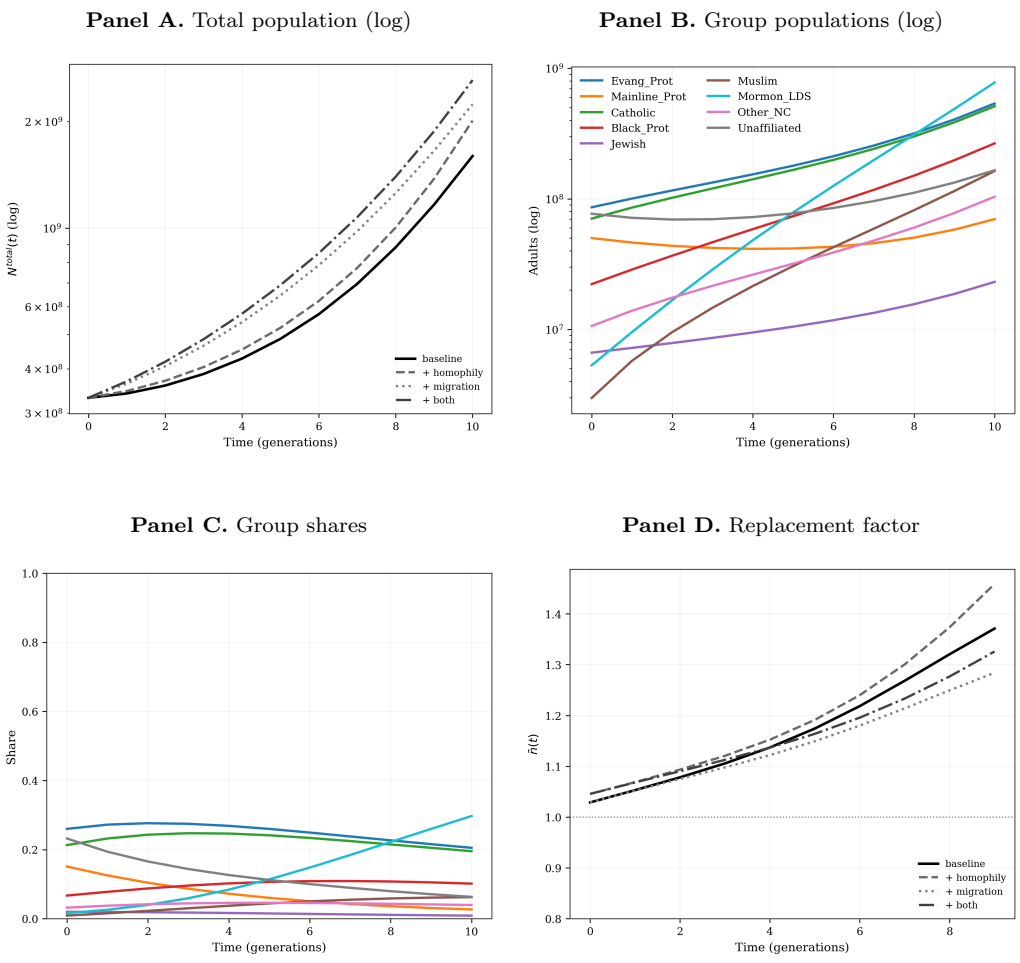

Figure 4. United States by Religion — Extensions

Table 1—Calibration Inputs

Group	q_{i0}	Fertility$_i$	ϕ_i	s_i	n_{i0}	d_{i0}	Anchor $\tilde{q}_{\text{env},i}$	m_i
Panel A: *United States by race/ethnicity*								
White (NH)	0.606	1.532	0.488	0.950	0.710	0.89	—	—
Black (NH)	0.127	1.581	0.491	0.906	0.703	0.82	—	—
Hispanic	0.196	1.946	0.490	0.955	0.911	0.73	—	—
Asian (NH)	0.062	1.309	0.483	0.979	0.619	0.71	—	—
AIAN (NH)	0.007	1.411	0.491	0.817	0.566	0.43	—	—
NHPI (NH)	0.002	2.218	0.497	0.942	1.038	0.55	—	—
Panel B: *United States by religion*								
Evangelical Protestant	0.260	2.3	0.488	—	1.122	0.90	0.81	0.130
Mainline Protestant	0.151	1.9	0.488	—	0.927	0.78	0.59	0.060
Catholic	0.213	2.3	0.488	—	1.122	0.83	0.75	0.410
Black Protestant	0.067	2.5	0.488	—	1.220	0.93	0.81[a]	0.010
Jewish	0.020	1.9	0.488	—	0.927	0.95	0.65	0.010
Muslim	0.009	2.4	0.488	—	1.171	0.98	0.79[b]	0.080
Mormon (LDS)	0.016	3.4	0.488	—	1.659	0.96	0.87	0.010
Other non-Christian	0.032	1.8	0.488	—	0.878	0.92	—[c]	0.150
Unaffiliated	0.232	1.7	0.488	—	0.830	0.89	0.68	0.140
Panel C: *World by religion*								
Christian	0.288	2.7	0.488	0.964	1.270	0.88	0.84	—
Muslim	0.256	3.1	0.488	0.964	1.458	0.99	0.79	—
Unaffiliated	0.242	1.7	0.488	0.964	0.800	0.90	0.77	—
Hindu	0.149	2.4	0.488	0.964	1.129	0.97	0.97	—
Buddhist	0.041	1.6	0.488	0.964	0.753	0.95	0.47	—
Folk/Other	0.022	1.8	0.488	0.964	0.847	0.96	—[c]	—
Jewish	0.002	2.3	0.488	0.964	1.082	0.95	0.46	—

Notes: This table reports inputs used in the calibrations. We construct the on-impact replacement factor as follows: when Fertility$_i$ is TFR (U.S. race/ethnicity; World), $n_{i0} = \text{TFR}_i \times \phi_i \times s_i$; when Fertility$_i$ is CEB (U.S. religion; women ages 40–59), $n_{i0} = \text{CEB}_i \times \phi$. Here, ϕ_i is the female share at birth implied by the sex ratio at birth, and s_i is female survivorship to about the end of childbearing. The remaining inputs are: q_{i0} initial adult share; d_{i0} direct-socialization/identity-retention target. We also report the *homophily anchor* $\tilde{q}_{\text{env},i}$, a directly observed own-type exposure used to calibrate group-specific diagonals G_{ii} of the mixing kernel in homophily runs (entries are left blank where no published value exists). Raw tilts $\widehat{\zeta}_i$ from these anchors are log-compressed into $\zeta_i^{\text{eff}} \in [1, 2]$ using (17); missing anchors imply $\zeta_i^{\text{eff}} = 1$. m_i is reported only for U.S. religion and gives the arrivals' composition by religion in the *fixed-composition* migration scenario; shares sum to 1 across rows. Inflow levels in migration extensions equal approximately 22.5 million per generation (0.9 million/year × 25 years). n_{i0} is used only on impact to pin down b_i; subsequent $n_i(t)$ are endogenous. [a] No separate endogamy data is published for *Black Protestant*; we use the *Protestant overall* data as a transparent proxy. [b] U.S. *Muslim* anchor from the RLS wave that reports this category (small-N in recent waves). [c] No published anchor for *U.S. Other non-Christian* and *World Folk/Other*.

Sources: U.S. race/ethnicity—NVSR: *Births—Final Data, 2023* (NCHS); female survivorship—*U.S. Life Tables, 2021* (NCHS; NHPI uses national female proxy); d_{i0} aligned to Pew intermarriage among newlyweds (2015). U.S. religion—Pew *Religious Landscape Study* (2014/2015) and follow-ups (Muslims 2017; Jews 2021); m_i from Pew *Faith on the Move* (mapped to nine buckets). World—Pew *The Future of World Religions* (2015) for shares and TFRs (2010–15); $\phi = 0.488$ (global SRB ≈ 105:100) and survivorship $s = 0.964$ from UN *World Population Prospects 2024* life tables; *Folk/Other* combines Pew's Folk (1.8) and Other (1.7). *Homophily anchors:* U.S. $\tilde{q}_{\text{env},i}$ from Pew RLS (same-religion spouse shares); World $\tilde{q}_{\text{env},i}$ from Pew Global Religious Futures (share of adherents living where their religion is the majority, 2020).

REFERENCES

Bar-El, Rebecca, Teresa García-Muñoz, Shoshana Neuman, and Yossef Tobol. 2013. "The Evolution of Secularization: Cultural Transmission, Religion, and Fertility—Theory, Simulations, and Evidence." *Journal of Population Economics*, 26(3): 1129–1174.

Bisin, Alberto, and Thierry Verdier. 2000. "Beyond the Melting Pot: Cultural Transmission, Marriage, and the Evolution of Ethnic and Religious Traits." *Quarterly Journal of Economics*, 115(3): 955–988.

Bisin, Alberto, and Thierry Verdier. 2001. "The Economics of Cultural Transmission and the Dynamics of Preferences." *Journal of Economic Theory*, 97(2): 298–319.

Bisin, Alberto, Giorgio Topa, and Thierry Verdier. 2004. "Religious Intermarriage and Socialization in the United States." *Journal of Political Economy*, 112(3): 615–664.

Boyd, Robert, and Peter J. Richerson. 1985. *Culture and the Evolutionary Process.* Chicago:University of Chicago Press.

Cavalli-Sforza, Luigi Luca, and Marcus W. Feldman. 1981. *Cultural Transmission and Evolution: A Quantitative Approach.* Princeton, NJ:Princeton University Press.

Currarini, Sergio, Matthew O. Jackson, and Paolo Pin. 2009. "An Economic Model of Friendship: Homophily, Minorities, and Segregation." *Econometrica*, 77(4): 1003–1045.

Darwin, Charles. 1859. *On the Origin of Species by Means of Natural Selection.* London:John Murray.

Lehrer, Evelyn L. 2004. "Religion as a Determinant of Economic and Demographic Behavior in the United States." *Population and Development Review*, 30(4): 707–726.

McQuillan, Kevin. 2004. "When Does Religion Influence Fertility?" *Population and Development Review*, 30(1): 25–56.

Panebianco, Fabrizio, and Thierry Verdier. 2017. "Paternalism, homophily and cultural transmission in random networks." *Games and Economic Behavior*, 105: 155–176.

A. Alternative direct-socialization schedule

This section re-implements our model with the direct-socialization cost used by Bar-El et al. (2013). In their simulations, the per-child cost of direct socialization takes a form that is increasing, convex, and diverges to infinity as $d \to 1$. Concretely, letting $d \in [0, 1)$ denote direct socialization intensity and suppressing time subscripts for clarity, we set (per child)

$$H_i(d) \;=\; \begin{cases} 0, & d = 0, \\ \dfrac{C_i}{\sqrt{1 - d^2}}, & d \in (0, 1), \end{cases} \tag{A1}$$

so $H_i(d)$ is convex, $H_i(0^+) = C_i$, and $H_i(d) \to \infty$ as $d \to 1$.

With all other primitives as in the main text, the FOC in d is

$$\Delta V \,(1 - q_{\text{env},i}) \;=\; H_i'(d) \;=\; \frac{C_i \, d}{(1 - d^2)^{3/2}} \,. \tag{A2}$$

implying $d_i^\star(q_{\text{env},i})$ is strictly decreasing and concave in own-group exposure. This preserves the cultural-substitution property highlighted by Bisin and Verdier (2001): minorities optimally socialize more, majorities less.

Relative to the quadratic-cost baseline, this specification increases the marginal cost steeply as d grows, thereby reinforcing the concavity of $d_i^\star(q)$. In our simulations partitioned by religion, the headline results remain unchanged (Figure A1 for the United States; Figure A2 for the world): high-fertility, high-retention groups gain share, and the aggregate replacement factor $\bar{n}(t)$ rises despite under-replacement averages elsewhere. The alternative cost schedule primarily affects the speed at which large groups taper direct socialization as their exposure expands.

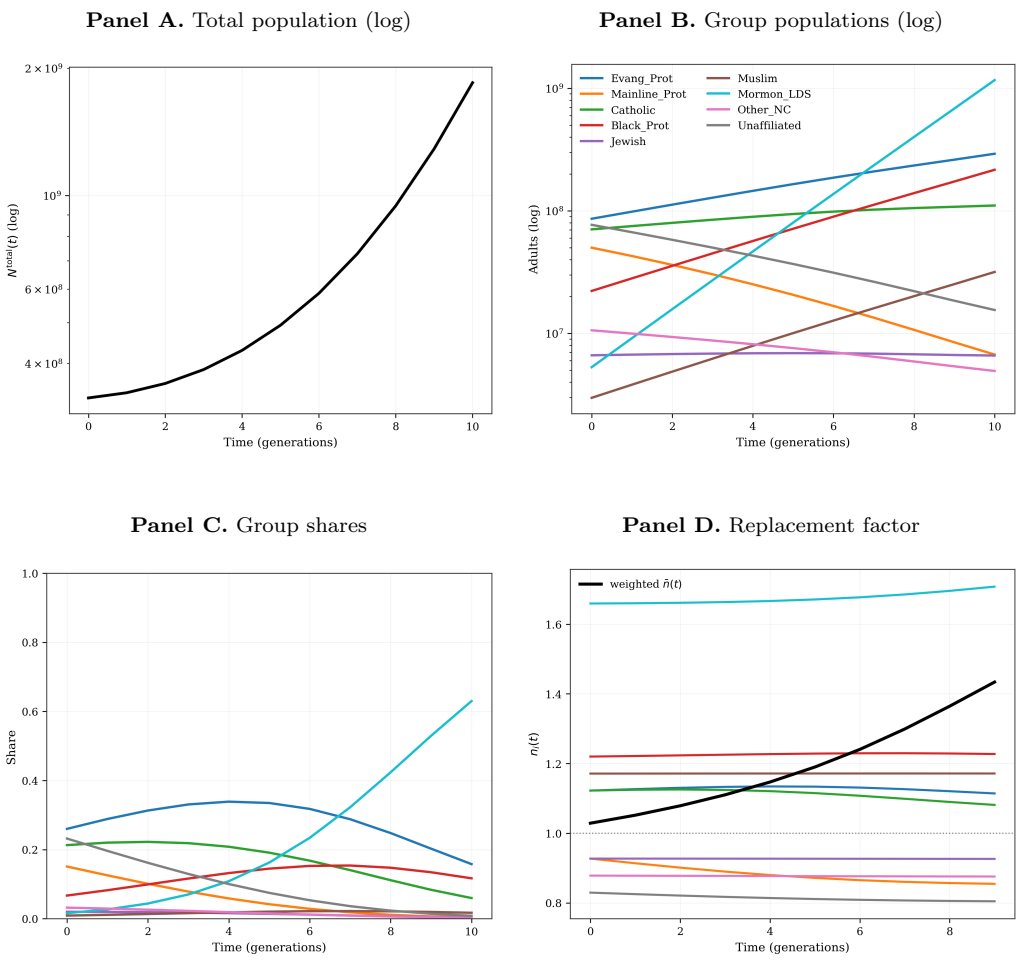

Figure A1. U.S. by Religion—Alternative socialization cost

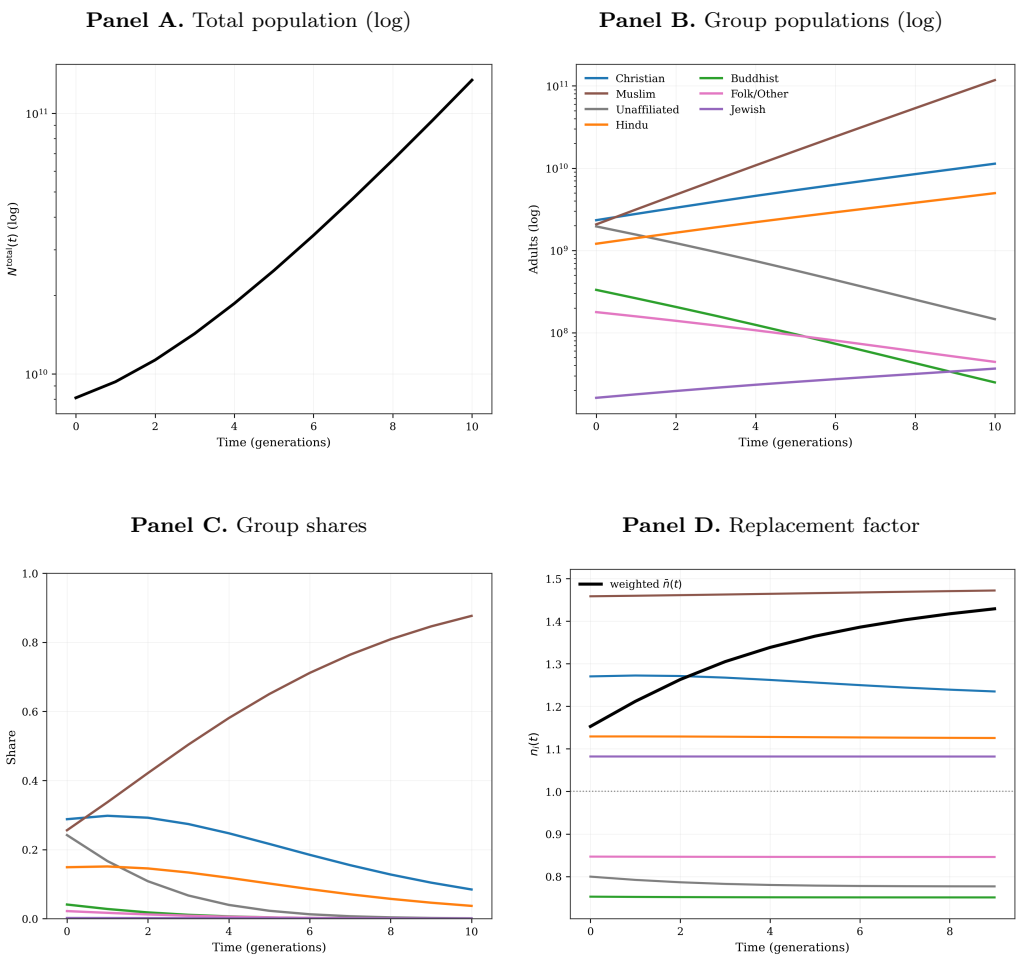

Figure A2. World by Religion—Alternative socialization cost

B. World extension

This section repeats the paper's extensions at the world scale, aggregating the population into seven religious traditions (Christian, Muslim, Unaffiliated, Hindu, Buddhist, Folk/Other, and Jewish), but now turning *on* homophily in the social environment. As in the paper, parameters are held fixed over ten generations; the exercise isolates how the direct-socialization choice interacts with composition when within-group exposure is greater than population exposure.

Figure B1 shows that homophily raises $q_{\text{env},i}$, which lowers optimal $d_i^\star$ and the per-child socialization cost; the freed resources raise $n_i^\star$, so the aggregate replacement factor $\bar{n}(t)$ shifts up relative to the baseline while preserving the ranking of groups by growth. Muslims expand fastest in both levels and share; the Unaffiliated and other low-fertility traditions decline in share. As in the U.S. exercise, homophily amplifies composition-driven dynamics without altering the qualitative message.

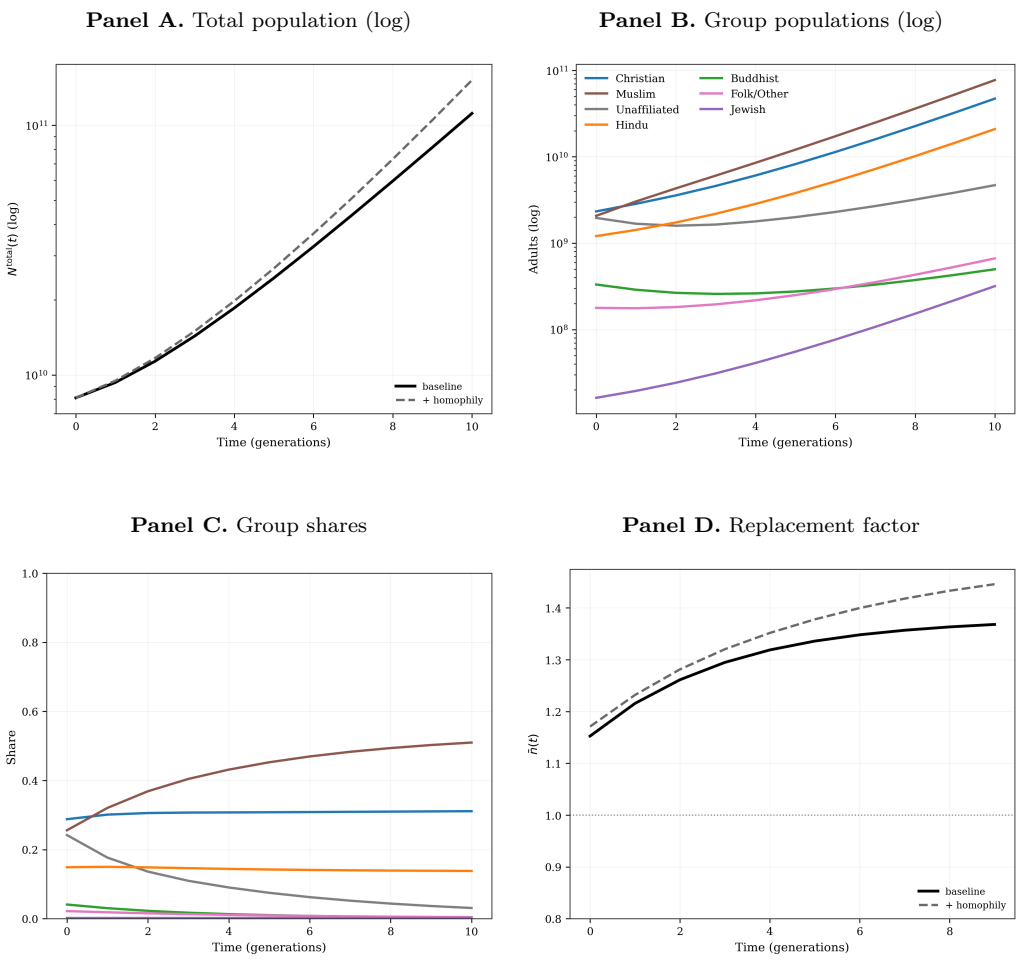

Figure B1. World by Religion — Extensions