Intergenerational Religious Persistence in the United States

Kristina Butaeva[1], Renato de Angelis[1], Steven N. Durlauf[1], and Weiqi Wang[1]

[1]*Harris School of Public Policy, University of Chicago*

September 1, 2026

Abstract

This paper studies the intergenerational dynamics of religious affiliation in the United States across birth cohorts using General Social Survey data from 1973 to 2024. We estimate cohort-specific transition matrices from religion of origin to current affiliation and apply Markov chain measures that distinguish mobility as movement, or how often people leave the religion in which they were raised, from mobility as memory, or how much origin shapes destinations across generations. We find that overall mobility has risen significantly, but almost entirely through structural reallocation to nonaffiliation, with circulation among traditions remaining steady across cohorts. The persistence of nonaffiliation increases substantially, and it becomes faster to reach from every origin, while Catholicism and mainline Protestantism become much harder to reach; steady-state proportions concentrate on nonaffiliation and evangelical Protestantism as two poles. The aggregate memory in the system also decreases, indicating that religious origin has become less predictive of adult affiliation. These processes differ by region, sex, and especially political party, with Democrats trending toward nonaffiliation and Republicans toward evangelical Protestantism.

Keywords: religion, intergenerational mobility, Markov chains.

This research was supported by the James M. and Cathleen D. Stone Foundation. We are grateful to participants in the Social Consequences of Religions (SCORE) 3 Strand: Social Consequences of Religions initiative research seminar at the University of Cambridge for their helpful comments. We used Claude Fable 5 (Anthropic) for assistance with coding our statistical analyses, and Claude Sonnet 5 (Anthropic) for copyediting the manuscript. Responsibility for all content rests solely with the authors. Direct correspondence to Renato de Angelis, Harris School, University of Chicago, 1307 E 60th St, Chicago, IL 60637, email: rdeangelis@uchicago.edu.

1 Introduction

Commentators since Tocqueville ([1835/1840] 2000) have noted the centrality of religion to American political life, with contemporary discussions emphasizing the roles played by evangelical Protestantism (FitzGerald, 2017; Hollinger, 2022) and increasing nonbelief (Campbell et al., 2020). The objective of this paper is to study the evolution of intergenerational persistence in religious affiliation over the last half-century. These estimates represent the background facts for substantive assessment of the changing interactions of (non-)religion and society. Changes in this persistence across birth cohorts may also give insights on how religious affiliations will evolve in the future.

The religious composition of the United States has changed significantly in the last half-century. At the highest level, this change is characterized by a decline in the share of Americans identifying as Christians (and particularly as mainline Protestants) and a corresponding increase in share reporting no religious affiliation (Smith et al., 2025). Voas and Chaves (2016) have argued that this decline has advanced through cohort replacement: each birth cohort is less likely to claim an affiliation, attend religious services, and profess religious belief than the previous. Schnabel and Bock (2017) counter that intense religiosity has persisted even as moderate religiosity has eroded, though this question remains disputed (Voas and Chaves, 2018; Schnabel and Bock, 2018). Past work on religious switching has found low Catholic persistence countered by immigration (Hout, 2016) and high retention for more conservative denominations and declining denominational loyalty for liberal denominations, though distances between religious origins and destinations remain relatively invariant across periods and cohorts (Sherkat, 2001).

Several mechanisms have been proposed to explain these trends: one strand of research places religious affiliation downstream of political identity, partially explaining the rise of nonaffiliation as motivated by a backlash against the religious right (Hout and Fischer, 2002, 2014; Margolis, 2018; Campbell et al., 2020; Braunstein, 2022). A second attributes decline in religion to a cumulative failure of intergenerational transmission (Voas, 2008; Bengtson et al., 2013; Puga-Gonzalez et al., 2022; Stolz et al., 2025), though there is also evidence that switching in adulthood changes overall composition more than transmission, and Schnabel et al. (2025) also find that disaffiliation is equally rapid among children raised by highly religious parents. Another explains it with reference to moral disagreements, particularly with respect to gender and sexuality (Woodell and Schwadel, 2020), or declining institutional credibility, especially around cases of sexual abuse (Hungerman, 2013; Bottan and Perez-Truglia, 2015).

Recent theoretical work has tended to understand secularization as a stadial and internally heterogeneous process, rather than a uniform one. The secular transition model likens religious decline to the demographic transition inasmuch as it is hard to reverse, varies geographically, and proceeds through cohort replacement and transmission failures rather than adult disaffiliation (Voas, 2008; Kasselstrand et al., 2023). Stolz et al. (2025) argue that the dimensions of religiosity are abandoned in rough order of cost: first public participation, then personal importance, and finally nominal belonging, which is the quantity of interest in the present study, which covers the United States, a country they identify as being in the mid-to-late stages of secularization (compared to northern European countries which are understood to be in the late stages of the process).

Watts and Wilkins-Laflamme (2026) argue that late-stage secularity is positively institutionalized and socialized rather than residual, as in mid-stage secularization, which can help explain the increasing persistence of nonaffiliation across generations. Both frameworks predict non-stationarity in transition rates and subpopulation heterogeneity within countries, and several accounts claim that persistence is greatest at the extremes of religiosity and weakest in the middle (Voas, 2009; Puga-Gonzalez et al., 2022; Bok, 2025).

Formal transition-based estimation remains scarce, with existing work estimating transition probabilities as inputs in projection models (Skirbekk et al., 2010; Kramer et al., 2022), fitting two-state Markov models Stinespring and Cragun (2015), or simulating transitions from assumed rules (Puga-Gonzalez et al., 2022). The mobility table analyses of Sherkat (2001, 2014) have not been extended over the most recent increase in nonaffiliation and, as they rely on log-multiplicative models of mobility, can only capture the strength of association between religious origin and current affiliation. The staging literature also disagrees on whether the United States conforms to or deviates from the European trajectory (Brauer, 2018; Stolz et al., 2025).

Our analysis contributes to this literature by developing a range of measures of intergenerational persistence. Our analysis employs Markov Chains to study persistence. The Markov Chain formalism is the natural formulation of the conditional probabilities that link parents and children given the categorical formulation we employ for affiliation. While Markov chains constitute a standard and longstanding methodology, we employ newly developed measures that facilitate the analysis of two related but different dimensions of persistence. The first is mobility as *movement*, which captures the probability that an individual claims a religious affiliation different from the tradition in which they were raised. The second is mobility as *memory*, which measures the extent to which being raised in a particular tradition affects adult affiliation. To illustrate the distinction, an adult may leave their given religious affiliation with high probability, yet the direction of the change is influenced by their starting point. Movement, in other words, characterizes the magnitude of switching, while memory captures the full effects of a particular origin on adult affiliation. We measure movement and memory in transitions from religious upbringing to current adult affiliation across cohorts born between 1925 and 1984 using the General Social Survey.

Our results produce a broad set of stylized facts that can guide theorizing about intergenerational religious dynamics in the United States.

For movement, we document that rising religious switching is not a uniform increase in fluidity but a reorganization around two poles. Overall mobility (OM) rises across cohorts born between 1925 and 1984, but this increase is driven almost entirely by structural mobility (SM), specifically the reallocation of the population toward None[1], while

[1]We use three related terms for each category, distinguishing the state in our category space, the tradition it corresponds to, and the adjective. Capitalized labels (Catholic, Evangelical, Black Protestant, Mainline, Other, None) denote the rows and columns of the transition matrix, arguments of the category-specific measures, and entries of π_0 and π_∞. Tradition names (evangelical Protestantism, mainline Protestantism, Catholicism, Black Protestantism, nonaffiliation) denote the traditions themselves. We use the corresponding adjectives (evangelical, mainline, Catholic, unaffiliated) for people and institutions, and *affiliated* for the five non-None categories taken together. Other, being a residual category rather than

exchange mobility (EM), the component that measures circulation among stable categories, stays flat. The relationship between overall mobility and steady-state mobility (SSM) changed direction over the period: in earlier cohorts, the long-run implication of that era's transition regime was more mobile than what was actually observed, whereas in the most recent cohorts the reverse holds. This trend is consistent with that of a system settling into an increasingly sorted, less mobile long-run structure. Category-specific passage times confirm the same pattern from a different angle. Reaching None becomes dramatically faster from every origin, reaching Evangelical stays easy throughout, and reaching Catholic or Mainline becomes markedly slower, while the expected time to leave Catholic or Mainline falls and the expected time to leave None rises to the longest of any category.

The memory measures point to the same reorganization, but reveal that it proceeds unevenly across traditions and unevenly across time. Aggregate one-step memory (AIM) is roughly flat through the middle of the period before declining sharply in the most recent cohorts, while mixing time (t_{mix}), the number of generations required for a cohort's transition regime to approach its long-run pattern within the specified numerical tolerance, falls by nearly half over the same span at a much steadier pace, suggesting the tail of the memory process began shortening well before any single generation started erasing more origin-state information. Underneath these aggregate patterns, Black Protestant origin retains a level of intergenerational memory well above every other tradition throughout, Catholic and Other converge downward from a once-higher tier, and Mainline and Evangelical diverge from one another rather than moving in step. None, the origin with the lowest memory in the earliest cohorts, by some measures, approaches the durability of established religious traditions.

This reorganization is not uniform across the population. The clearest split runs along party lines: Democrats' long-run trajectory increasingly points toward None and Republicans' toward Evangelical, a divide that widens across cohorts. Region and sex shape the pace and channel of the same shift rather than its basic direction: the South remains anchored by a durable Evangelical identity, the West is furthest along in the move toward None, and men and women arrive at the same two poles by different routes, men leaning more toward None and women more toward Evangelical.

The paper proceeds as follows. Section 2 describes the data and the Markov chain–based measures of movement and memory used throughout. Section 3 presents the estimated cohort-specific transition matrices and their implied stationary distributions. Section 4 reports the movement results, and Section 5 the memory results. Section 6 examines heterogeneity in these patterns by region, political party identification, and sex. Section 7 concludes.

a tradition in its own right, appears only as a state label.

2 Methods

2.1 Data

Our data come from the General Social Survey (GSS), a nationally representative cross-sectional survey that collects information on respondents' religious affiliations and other socio-demographic variables (Davern et al., 2025). The GSS was conducted via home interviews every year between 1972 and 1994 (save for 1979, 1981, and 1992), and every two years since. Our main analysis uses the unweighted GSS sample from 1973–2024, excluding the 2021 wave. We exclude 2021 because the COVID-induced shift from face-to-face to online interviewing generated differential non-response that disproportionately under-represented the most intensely religious respondents (Schnabel et al., 2024). Because the 2022 and 2024 waves were administered using a mixed-mode design that raises a similar concern, Appendix E reports a robustness check that additionally excludes these two waves, restricting the analysis to the pre-2020, exclusively in-person survey waves. In addition, since our baseline sample retains the 1982 and 1987 Black oversamples, to assess whether this sampling design affects our estimates, Appendix F re-estimates the full set of measures using the appropriate survey weights. Across both robustness checks, the substantive patterns in our main results remain essentially unchanged.

We restrict respondents to individuals born between 1925 and 1984 who were between 30 and 75 years old at the time of the survey. We set a lower bound at 30 because religious affiliations tend to be relatively unsettled and biased toward None in young adulthood and stabilize around the age of 30 (Pew Research Center, 2025); the upper bound limits the influence of differential mortality rates. Religious participation is associated with lower mortality risk (Hummer et al., 1999), and because our oldest cohorts are observed at disproportionately older ages, this could bias results. We group respondents into six decadal birth cohorts, from 1925–34 to 1975–84.

These restrictions jointly determine the cohort range. Given the 1973–2024 survey window and the 30–75 age window, the oldest cohort is first observable at age 39 (born 1934, interview in 1973) and the youngest cohort is last observable at age 49 (born 1975, interviewed in 2024). Thus, each cohort can be observed over ages 39–49, with each age in the range observed continuously for each cohort. Because the cohorts are nonetheless observed at different portions of the 30–75 age window, we re-estimate each transition matrix on the age-standardized 39–49 subsample in Appendix D. Levels and trends are qualitatively unchanged, which we take as evidence that our findings are not artifactual. We do note that the sample sizes for the band-restricted estimates are smaller, and that the restriction introduces cohort-period collinearity. We emphasize that we do not claim to separate cohort and period effects.

2.2 Religious Affiliations Classification

Across each GSS wave from 1973 to 2024, respondents were asked about their current religious preference and the religion in which they were raised, with detailed denominational classifications for Protestants. In 2024, the GSS began to classify the responses

to both of these questions according to RELTRAD, a widely used typology of American religious traditions that was first proposed by Steensland et al. (2000) and later modified by Stetzer and Burge (2016). We measure religious origin using the GSS variable RELTRAD16, which captures respondents' religious tradition of origin and is constructed from their reported religion in which they were raised (RELIG16), Protestant denomination (DENOM16), recoded open-ended responses for other Protestant denominations (OTH16), and race. Current religious affiliation is measured using RELTRAD, which is constructed analogously to RELTRAD16, and additionally includes attendance at religious services ATTEND[1]. RELTRAD divides respondents into seven groups: evangelical Protestant, mainline Protestant, Black Protestant, Catholic, Jewish, other, and unaffiliated. We retain Black Protestant as a distinct category throughout the main analysis, given its substantive and theoretical interest in the literature on religious persistence, and collapse Jewish into the residual Other category to minimize sampling zeroes in the transition matrices. We label the resulting six states Catholic, Evangelical, Black Protestant, Mainline, Other, and None.

Our unit of analysis is an individual-level transition from religious origin to current religious affiliation, rather than a parent–child dyad. The distinction is consequential inasmuch as a separate measure of parental affiliation would record a categorical answer and be silent on whether anything was transmitted, whereas a measure of religious upbringing corresponds more closely to actual (recalled) practice, filtering parental affiliation through intensity. We consider this an appropriate construct, especially considering the fuzzy fidelity literature (Voas, 2009), which suggests that many affiliations are only weakly or nominally held. Thus we refer throughout to religious origin rather than parental religion.

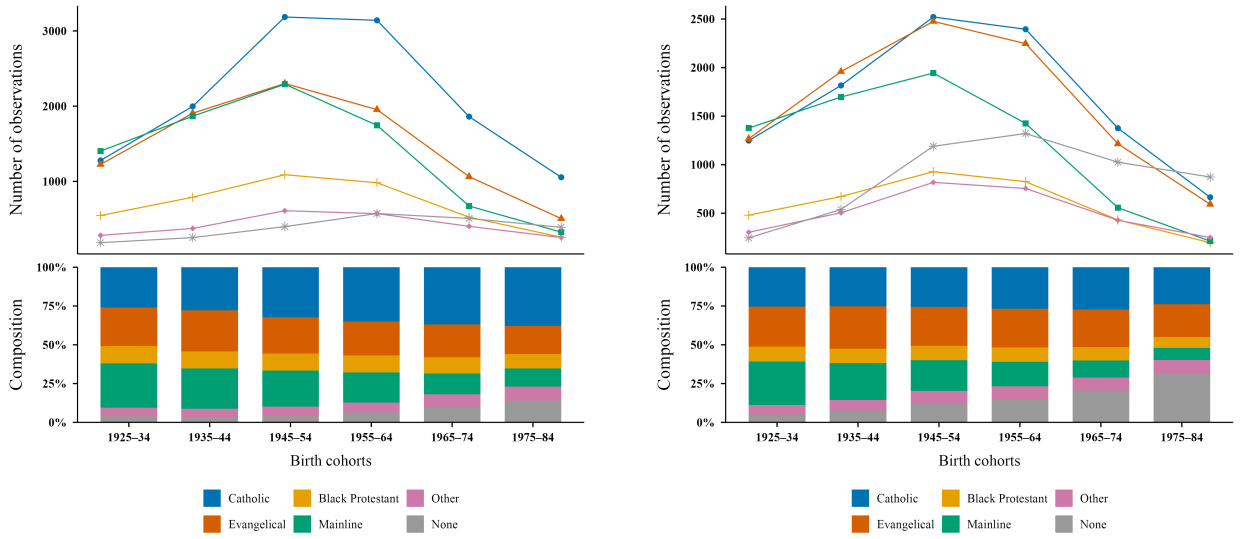
This approach carries a substantial weakness: namely, it is impossible to distinguish between an accurate retrospective self-report and one that re-interprets upbringing to be more consistent with one's current identity. For example, an individual who identifies as unaffiliated and reports a religious origin of None may be identifying an entirely secular upbringing, one in which, say, the Catholic identification of her parents was nominal (which is what we would like the measure to capture), or she may be re-interpreting a genuinely Catholic upbringing to match her current nonaffiliation, which would bias our persistence estimates. Given the survey design, it is impossible to tell the difference, and this is a limitation. Hout and Hastings (2016), using the GSS panels, find that religious origin is reported consistently by the same individual across waves, which indicates that the recall error is not substantially random, but it cannot identify biased reconstruction, which is our concern.

Separately, the destination category of None is internally heterogeneous, with a substantial portion of unaffiliated respondents moving between affiliation and nonaffiliation over time (Lim et al., 2010). If churning is stationary, our repeated cross-sections recover affiliation at an interview conditional on religious origin without bias. The fact of liminality, however, suggests that the diagonal cells of the transition matrix should not be

[1] In RELTRAD, ATTEND is used only to classify particular Protestant or nonspecific Christian cases with limited or ambiguous denominational information, rather than as a general measure of religiosity. Because the GSS does not collect a comparable measure of religious-service attendance at age 16, ATTEND does not enter RELTRAD16 (Sparkman et al., 2026).

read as entirely settled identities for affiliated and unaffiliated respondents alike.

Figure 1: Sample composition by religious origin and current affiliation.



(A) Religious origin (RELTRAD16)

(B) Current affiliation (RELTRAD)

Notes: Top plots report raw individual-level transition counts by religious category and 10-year birth cohort; bottom plots report the corresponding compositional shares. Panel A classifies respondents by *RELTRAD16*, the GSS variable recording the respondent's religion tradition at age 16; Panel B classifies the same respondents by *RELTRAD*, the GSS measure of respondents' religious tradition at the time of interview.

Figure 1 reports the raw counts and compositional shares underlying the transition matrix estimates that follow, separately for religious origin (Panel A) and current affiliation (Panel B). Cohort sample sizes rise from the 1925–34 cohort to a peak in the 1945–54 and 1955–64 cohorts before falling sharply in the two youngest cohorts; this non-monotonicity follows mechanically from the fixed 1973–2024 survey window intersected with the 30–75 age restriction, which gives the middle cohorts more calendar years of coverage than the oldest and youngest ones. Compositionally, the origin distribution remains dominated by Catholic, Evangelical, and Mainline origins throughout, with the share reporting a None origin rising only modestly, from roughly 3–4 percent in the earliest cohort to roughly 13 percent in the latest. The destination distribution shows a substantially steeper shift: the share currently unaffiliated rises from roughly 4 percent to roughly 30 percent across the same cohorts. Our final analytic sample comprises $n = 38,778$ respondent-level observations, with cohort sizes ranging from $n = 2,789$ to $n = 9,877$.

2.3 Analysis

We analyze religious persistence using the Markov chain-based measures proposed by Blume et al. (2025), Wodtke et al. (2026), Butaeva, Chen, Durlauf, and Park (2026), and

Butaeva, Durlauf, and Shapoval (2026). These measures take the empirical transition probabilities from religious origin to current affiliation for each cohort we observe in the GSS and act as if the resulting transition 'regime' were to remain stable for each subsequent generation, allowing us to estimate the theoretical long-run influence of religious origins on families conditional on indefinitely stable transition probabilities.[1]

Consider the category space $S = \{1, ..., 6\}$, where each element corresponds to a religious affiliation. Let P be a transition matrix in which each element p_{ij} is the probability of identifying with adult affiliation $j \in S$ conditional on religious origin $i \in S$. Applied to π_0, a 1×6 vector describing the distribution of religious origins, P returns the corresponding distribution of destinations $\pi_1 = \pi_0 P$. This can be generalized to any arbitrary generation t, yielding $\pi_t = \pi_{t-1} P = \pi_0 P^t$.[2]

Our measures follow from asking what P would imply if the transmission regime it summarizes were held fixed and iterated across generations. For a primitive[3] matrix P, iterating the sequence produces a limiting vector $\pi_\infty = \pi_\infty P$, with P^t converging to a limiting matrix P^∞ where each row equals π_∞. We refer to π_∞ as the steady-state or stationary distribution implied by the full transition system; the influence of origin disappears in the limiting matrix $P^\infty = \mathbf{1}\pi_\infty$, in which every row equals π_∞.

A few caveats: first, P remains fixed as we iterate, so our measures do not incorporate differential fertility or immigration, both of which are significant for religious demography. Adjusting π_∞ for observed fertility (see Appendix C) results in no qualitative differences. Second, this construction assumes that transition probabilities are time-homogeneous[4] and first-order Markovian.[5] Third, we emphasize that π_∞ should not be read as a forecast; as Wodtke et al. (2026) note, their analytic function (and that of the measures derived from π_∞) is analogous to that of period life tables in demography, which generate 'synthetic cohorts' by applying age-specific death rates from a single cross-section to a population. Like our measures, their value is comparative: we are interested less in the level implied by any single cohort's regime than in how the steady-state distribution differs across cohorts. The measures we use were originally developed to measure intergenerational occupational mobility (Prais, 1955; Shorrocks, 1978; Sommers and Conlisk, 1979; Blume et al., 2025; Butaeva et al., 2026; Wodtke et al., 2026), and so we adopt them with some reservations. Most importantly, whereas in occupational mobility studies π_∞ is understood as a state of equality of opportunity where descendants' occupational outcomes do not depend on parents' occupations, the analysis that follows will not interpret the stationary distribution normatively. Rather, π_∞ represents the limiting distribution implied by the transition regime. It is the common

[1] Recall that we do not observe a strict parent–child dyad, but rather a within-individual transition. In using generational language throughout this section, we assume that religion of upbringing stands in for the religion transmitted by the respondent's parents and that the transition from the parental household to adult religious affiliation is the relevant phenomenon. We regard this measure as more appropriate than relying on nominal parental affiliation, but we emphasize that this interpretation rests on an assumption.

[2] Throughout, P^t denotes the t-fold matrix product of P, and $(P^t)_{ij}$ denotes its (i, j) entry: the probability of reaching category j after t transitions conditional on starting in category i.

[3] That is, $\exists$ an integer $k > 0$ such that $P^k > 0$.

[4] The probability structure described by P does not change across generations.

[5] Affiliation in generation t depends only on the distribution in generation $t - 1$, and not on the distribution in any earlier generation.

row of the limiting matrix $P^{\infty} = \mathbf{1}\pi_{\infty}$, in which every origin has the same destination distribution. As such, it constitutes a substantive benchmark representing the complete decoupling between upbringing and affiliation. Another corollary of using these measures in a nominal, rather than ordinal, state space is that there is no analogue to upward or downward mobility. The directional measures we treat as substantively meaningful concern particular origin-destination pairs.

We inherit from Wodtke et al. (2026), Butaeva, Chen, Durlauf, and Park (2026), and Butaeva, Durlauf, and Shapoval (2026) a distinction between mobility as movement and mobility as memory.

2.3.1 Mobility as Movement

The first set of measures asks how often individuals leave the religious traditions in which they were raised. In our setting, this corresponds to the proportion of individuals whose adult affiliation differs from their origin, or, mathematically, the mass lying off the diagonal of P, which we call overall mobility (OM). We additionally calculate the proportion of each subsequent generation under a fixed transition probability structure:

$$OM_t = 1 - \sum_i \pi_{i,t-1} \cdot P_{ii}. \tag{1}$$

These sorts of departures have two sources: switching can occur because the marginal distribution of religious affiliations itself is changing, while the remainder happens net of such change. This distinction maps onto the standard distinction between structural (SM) and exchange (EM) mobility in the sociology literature on intergenerational mobility (Erikson and Goldthorpe, 1992; Featherman et al., 1975; Sobel et al., 1985). Following this standard decomposition, we measure the first as the total variation distance between successive marginals,

$$SM_t = \frac{1}{2} \cdot \sum_i |\pi_{i,t-1} - \pi_{i,t}|. \tag{2}$$

And we take the rest as the residual:

$$EM_t = OM_t - SM_t. \tag{3}$$

Structural mobility only picks up movement that shows up as a change in the overall mix of religious categories. If every tradition lost and gained members at similar rates, that mix would stay roughly the same across generations, and most switching would show up as exchange mobility instead. A large SM_t, then, points to an *asymmetry* in how traditions gain and lose members—some traditions losing more members than they are replaced, others holding onto members and absorbing most new converts. Structural mobility is thus best understood as the aggregate, compositional signature of this asymmetry in persistence across traditions, rather than a separate phenomenon from persistence itself.

Because the marginal distribution stops changing once the steady-state distribution

π_∞ is reached, structural mobility vanishes in the limit and overall mobility becomes equivalent to exchange mobility alone:

$$SSM = 1 - \sum_i \pi_{i,\infty} \cdot P_{ii}. \tag{4}$$

Steady-state mobility (SSM) is the stationary-composition-weighted off-diagonal movement rate: the probability of changing categories in one transition when the origin distribution equals the stationary distribution implied by the full transition matrix.

A second question concerns individual traditions rather than the system as a whole: for how long does a tradition persist in a lineage before leaving it? The mean time to exit (MTE) (Kemeny and Snell, 1976) answers this in generations t for each religious origin i:

$$MTE_i = (1 - P_{ii}) \cdot \sum_{t=1}^{\infty} t \cdot P_{ii}^{t-1}. \tag{5}$$

With MTE_i, higher values indicate greater intergenerational persistence, and vice versa for lower values.

We capture a directional aspect of movement not captured by aggregate measures through mean first passage time ($MFPT$) (Kemeny and Snell, 1976; Grinstead and Snell, 1997). $MFPT$ estimates the expected number of generations it would take to reach destination category j when starting from origin category i. If the transition occurs immediately, then the number of generations is 1. Otherwise, the process moves to some intermediate k, a generation elapses, and $MFPT_{kj}$ remain in expectation:

$$MFPT_{ij} = P_{ij} + \sum_{k \neq j} P_{ik} \cdot (MFPT_{kj} + 1). \tag{6}$$

We then solve a system of 30 equations in the off-diagonal unknowns. We are especially interested in asymmetries between $MFPT_{ij}$ and $MFPT_{ji}$, which describe the structure of religious switching directionally.

2.3.2 Mobility as Memory

While movement concerns a single step, we are also interested in how particular religious origins affect dynasties after multiple iterations of the stationary transition matrix. A tradition can leave a durable imprint even if most of its adherents leave in a single step as long as the destinations of its descendants remain distinct from those of individuals of different origins. Memory measures capture how long the influence of origin persists, and how quickly it dissipates. We are especially interested in how memory changes for specific traditions across cohorts.

The benchmark is the limiting matrix, in which each row equals π_∞. Aggregate intergenerational memory (AIM) measures how far a regime sits from that condition

after t generations (Blume et al., 2025):

$$AIM_t = \frac{1}{2} \cdot \sum_{i,j} |(P^t)_{ij} - \pi_{j,\infty}| \cdot \pi_{i,0}. \tag{7}$$

At $t = 1$ it summarizes the strength of origin dependence; its decline in t describes how that dependence attenuates. Dropping the origin weights gives the same quantity for each tradition separately,

$$IM_{i,t} = \frac{1}{2} \cdot \sum_{j} |(P^t)_{ij} - \pi_{j,\infty}|, \tag{8}$$

whose $\pi_{i,0}$-weighted average recovers AIM_t. The individual memory curves (IM) express how persistence differs across religions. Additional memory measures—λ_2 and Dobrushin's $d'(t = 1)$—are defined in Appendix B.

While AIM_t and $IM_{i,t}$ track how much origin-state information survives after a fixed number of generations t, we can equivalently ask how many generations it takes for that information to disappear entirely. We call this the mixing time,

$$t_{mix} = \min\{t : P^t = P^\infty\}, \tag{9}$$

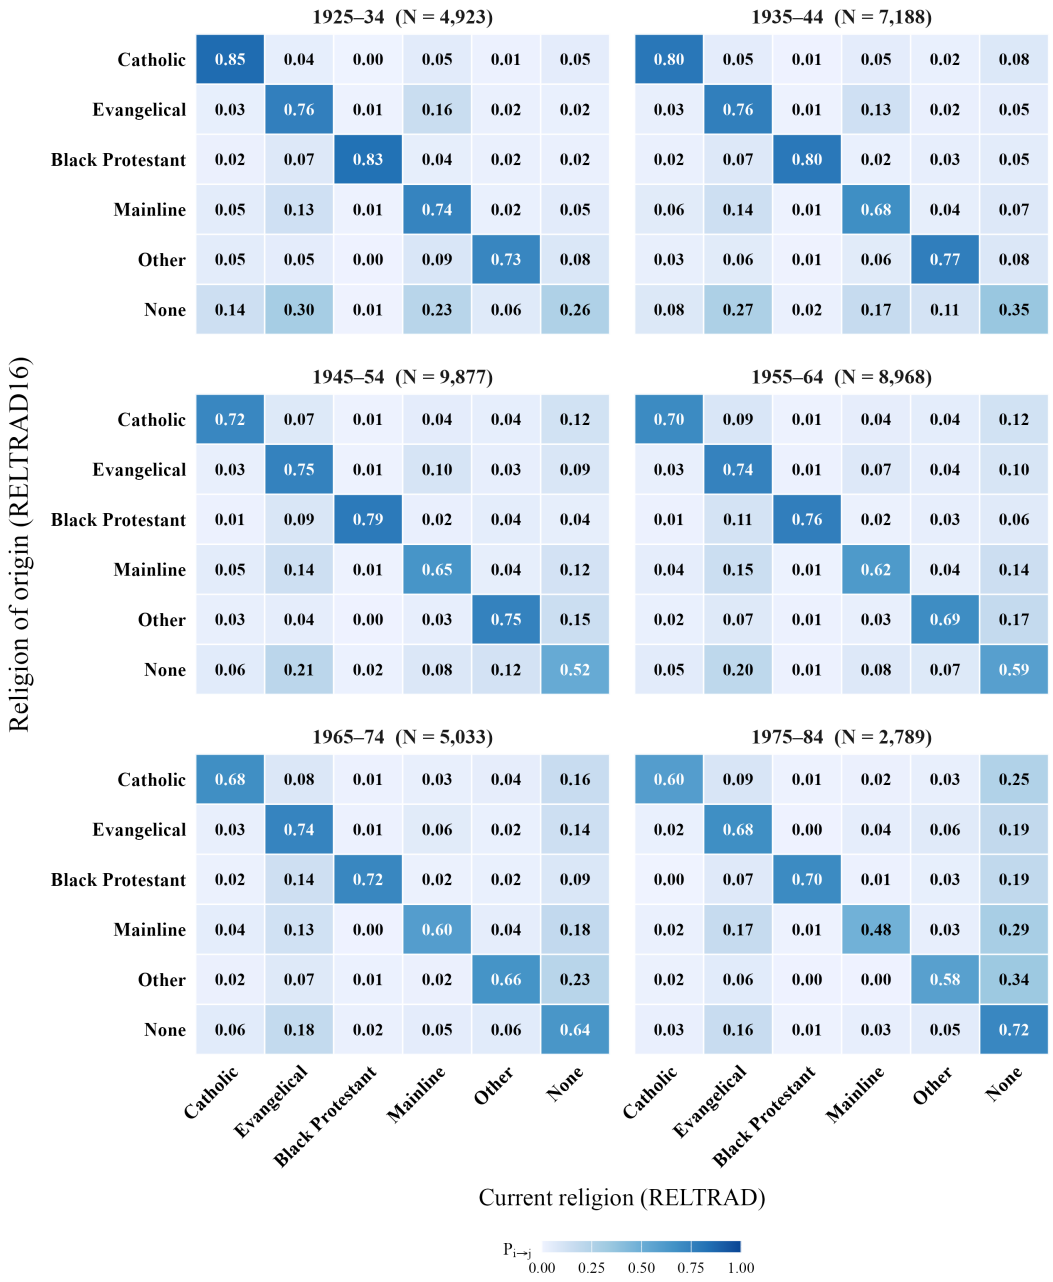
the smallest number of generations after which repeatedly applying a cohort's transition matrix produces a matrix in which every row equals π_∞, so that religious origin no longer carries any information about destination.[1] Where AIM_t and $IM_{i,t}$ describe the pace of convergence at a fixed horizon, t_{mix} summarizes the same underlying process in a single number: the point at which a cohort's transmission regime, if it held indefinitely, would have fully run its course. In Appendix A, we formally connect local persistence and memory, showing how each origin's diagonal retention probability and mean time to exit (MTE_i) relate to its individual memory profile ($IM_{i,t}$).

3 Transition Probabilities and Stationary Distributions

The cohort-specific transition matrices (Figure 2) display the state-level shifts that produce the aggregate trend. Reading across cohorts, the dominant change is the growth of the None column: outflows from every affiliated origin category into None rise substantially in each successive 10-year window, with the ordering Other > Mainline > Catholic > Evangelical ≈ Black Protestant preserved throughout. Outflows in the single digits for most affiliated origins in the 1925–34 cohort reach roughly one-quarter to one-third by the 1975–84 cohort for Other, Mainline, and Catholic, with Evangelical and Black Protestant somewhat lower throughout. Against this common trend, origin categories diverge

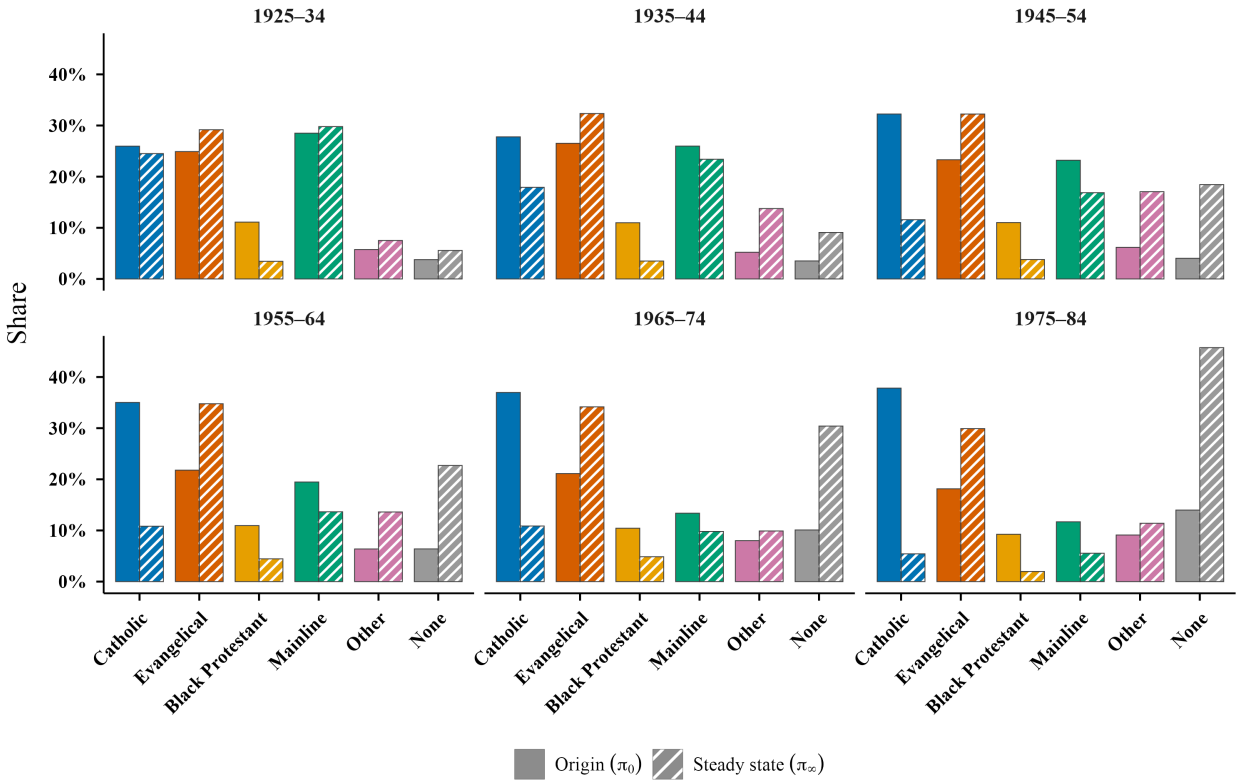
[1] In practice, convergence is assessed numerically rather than by requiring exact equality. We define the convergence time as the smallest t for which the maximum total-variation distance between any two rows of P^t is no greater than 0.01: $d'(t) \equiv \frac{1}{2} \max_{i,k} \sum_j \left|(P^t)_{ij} - (P^t)_{kj}\right| \leq 0.01$. The stationary distribution itself is computed independently as the normalized left eigenvector of P associated with eigenvalue 1, satisfying $\boldsymbol{\pi} P = \boldsymbol{\pi}$ and $\sum_j \pi_j = 1$.

Figure 2: Transition matrices by birth cohort.



Notes: This figure shows RELTRAD16 → RELTRAD transition matrices by 10-year birth cohort, 1925–1984. Rows denote religious origin at age 16 (RELTRAD16); columns denote current religious affiliation (RELTRAD). Cell values are row-conditional transition probabilities, p_{ij}.

Figure 3: Origin and implied stationary distributions.



Notes: This figure shows the share of each religious group in the origin distribution (π_0, top panel) and in the implied stationary distribution (π_∞, bottom panel), by 10-year birth cohort.

sharply in their diagonals. Catholic persistence falls steadily from 0.85 to 0.60 across the full range, and Black Protestant persistence falls from 0.83 to 0.70. Evangelical persistence is comparatively stable at 0.74–0.76 through the 1965–74 cohort before falling to 0.68 in the youngest cohort. Conversion into Catholic and Mainline is low and declines further in later cohorts; by the post-1950 cohorts, Evangelical is effectively the only non-None destination receiving meaningful inflow from the remaining traditions. Within the None row, outflow into Evangelical remains the largest non-diagonal destination in every cohort, even as it declines in absolute terms (from 0.30 to 0.16) alongside a much sharper drop in outflow toward Mainline (from 0.23 to 0.03); meanwhile None's own diagonal persistence rises from 0.26 to 0.72. This pattern is consistent with a stable subpopulation continuing to move between None and Evangelical even as both categories otherwise strengthen in place.

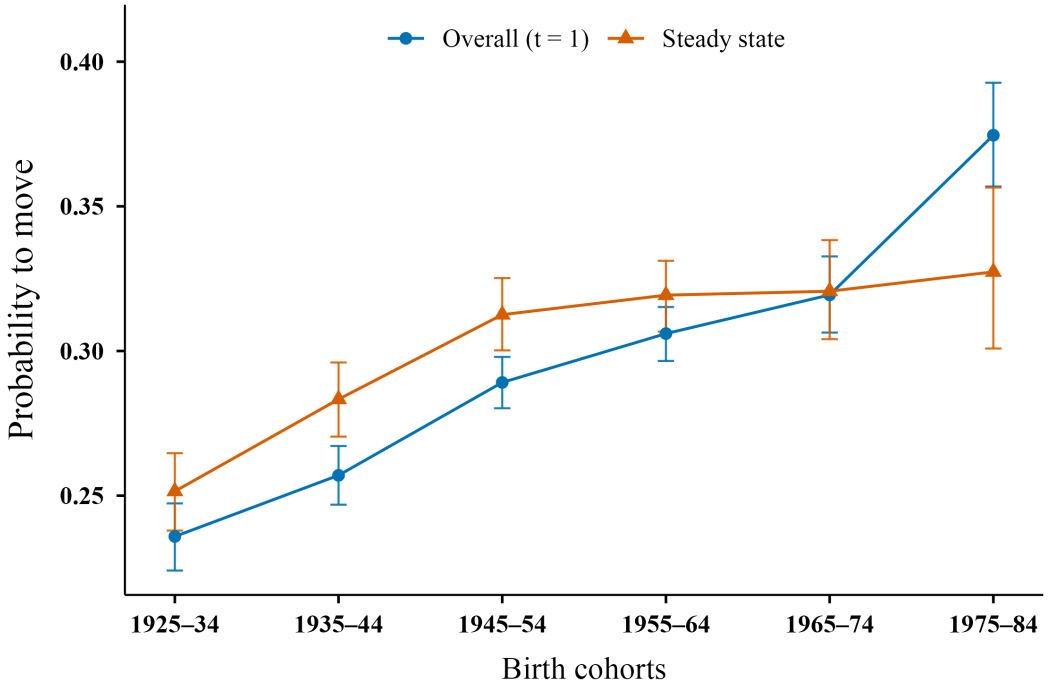
Figure 3 traces the implied origin and stationary distributions. The origin distribution (π_0, top panel) shows Catholic's share of origins rising steadily across cohorts (from roughly a quarter to over a third), while Mainline's origin share falls sharply (from roughly three-tenths to about one-tenth); Evangelical's origin share declines more modestly, and None and Other origin shares rise only slightly. The implied stationary distribution (π_∞, bottom panel) tells a different story: Evangelical's stationary share rises through the middle cohorts, peaking around 1955–74, before falling back in the youngest cohort, while None's stationary share grows steadily throughout and overtakes Evangelical to become the largest category in the 1975–84 cohort. Catholic and Mainline both decline toward small stationary shares across the full range, and Black Protestant's stationary share remains low and roughly flat. The overall picture is one in which persistence itself reorganizes the long-run system around two poles—None and Evangelical—even as the balance between them shifts toward None in the most recent cohort.

4 Movement Measures

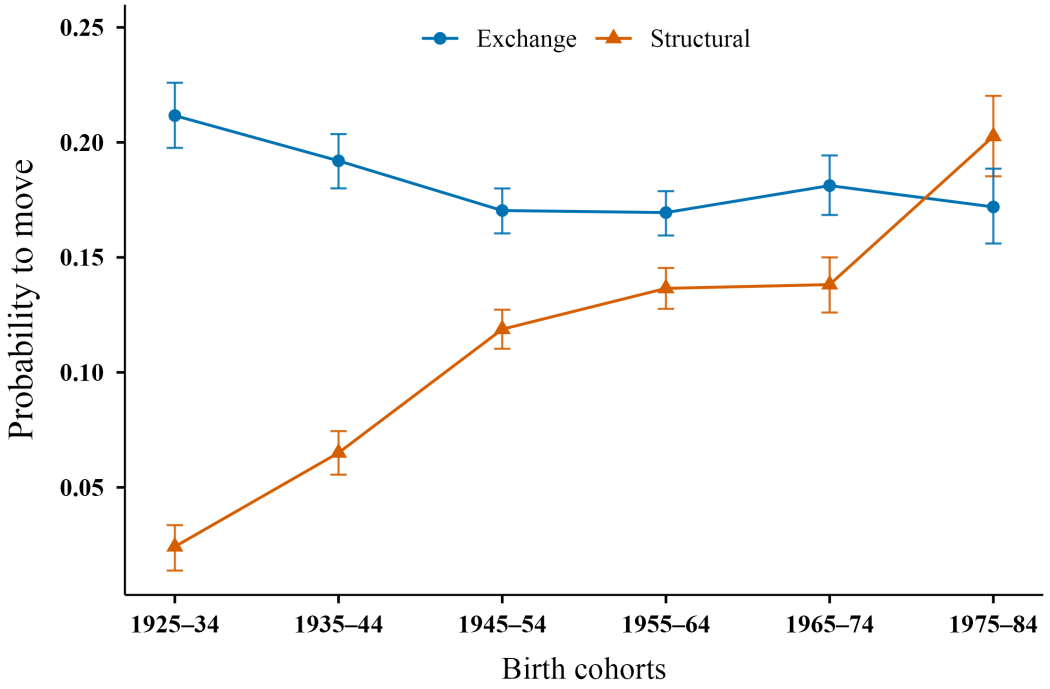
Overall mobility rises across cohorts, while steady-state mobility rises earlier and more modestly and then remains comparatively flat after the middle cohorts (Figure 4A). In the earlier cohorts, steady-state mobility exceeds mobility at $t = 1$: the long-run endpoint of that era's transition regime is more mobile than what was actually observed at the time. The two measures converge around the 1965–74 cohort, after which the relationship flips—observed mobility now exceeds the steady-state level, and the gap widens sharply in the 1975–84 cohort. This flip is consistent with the shift in π_∞ documented in Section 3: as the stationary distribution increasingly concentrates weight on None and Evangelical, the two categories with the highest and most stable diagonal persistence (Figure 2), the steady state comes to be weighted more heavily toward high-retention categories than the current cross-section is, making it, on net, less mobile than what is presently observed.

The increase in overall mobility is driven almost entirely by structural mobility (Figure 4B). Structural mobility rises in two stages: first from the 1925–34 through the 1945–54 cohorts, then holds roughly flat through the 1965–74 cohort, before rising sharply

Figure 4: Movement measures.



(A) Overall and steady-state mobility



(B) Structural and exchange mobility

Notes: Panel A reports overall mobility at $t = 1$ (OM_1) and steady-state mobility (SSM); Panel B decomposes OM_1 into its structural (SM_1) and exchange (EM_1) components. Bias-corrected bootstrap estimates with 95% confidence intervals (2,000 within-survey-year resamples).

again in the 1975–84 cohort. Exchange mobility, by contrast, is comparatively stable across the full range, apart from a higher estimate in the 1925–34 cohort, when switching more closely resembled circulation among traditions of roughly fixed size than a one-way drift toward new destinations. This pattern anticipates the category-specific results in Section 4.1: because Catholic and Mainline retention declines while conversion into these traditions stays negligible, and because Evangelical and None become the near-exclusive destinations for converts, the aggregate composition is pulled toward these two poles each generation almost mechanically. Structural mobility tracks the pace of this compositional reallocation; it is not an increase in reciprocal switching between stable categories, which instead would appear as exchange mobility—and which, consistent with this reading, remains comparatively flat.

Our stable exchange-mobility trend is broadly consistent with the stable origin–destination association that log-multiplicative models have found over a similar, though not identical, span of cohorts (Sherkat, 2001, 2014). The two are not the same quantity: those models capture association through odds ratios, which are mathematically invariant to the marginal distribution, whereas exchange mobility is a probability measure computed by subtracting out the observed shift in marginals between generations.

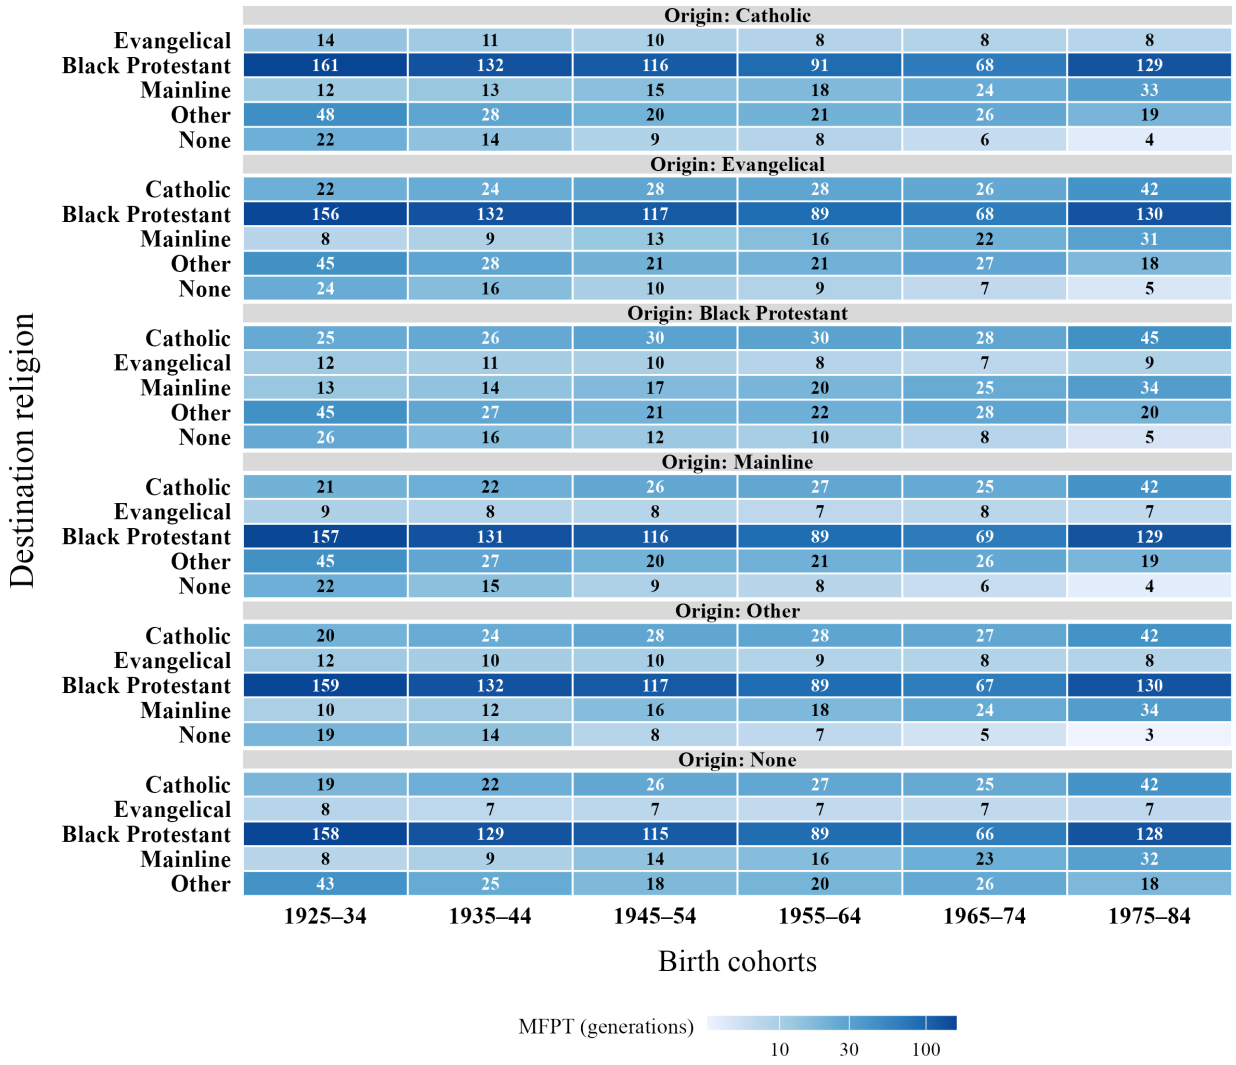
4.1 Category-specific Movement Measures

Figure 5 summarizes the category-level dynamics in passage-time units. Transit into None accelerates sharply for every affiliated origin: by the most recent cohort a None-bound passage takes roughly 3–5 generations from any origin, down from roughly 19–26 in the earliest cohort. Passage into Catholic and Mainline slows correspondingly, roughly doubling in the former case and tripling in the latter, consistent with these traditions' negligible convert inflow, while passage into Other declines by about half from every origin. Evangelical is the exception: passage into it stays low and largely flat throughout, making it the most reachable religious destination for almost every origin-cohort pair. Passage into Black Protestant declines through the 1965–74 cohort before jumping back up uniformly from every origin in 1975–84; given how small the Black Protestant category becomes in that cohort, we read this as sampling noise rather than a substantive reversal.

From every affiliated origin, the expected time to reach None falls sharply across cohorts, while the time to reach Catholic and Mainline rises correspondingly. For None-origin lineages specifically, Evangelical remains the closest affiliated destination throughout, Catholic and Mainline become markedly more distant, and Other becomes substantially easier to reach.

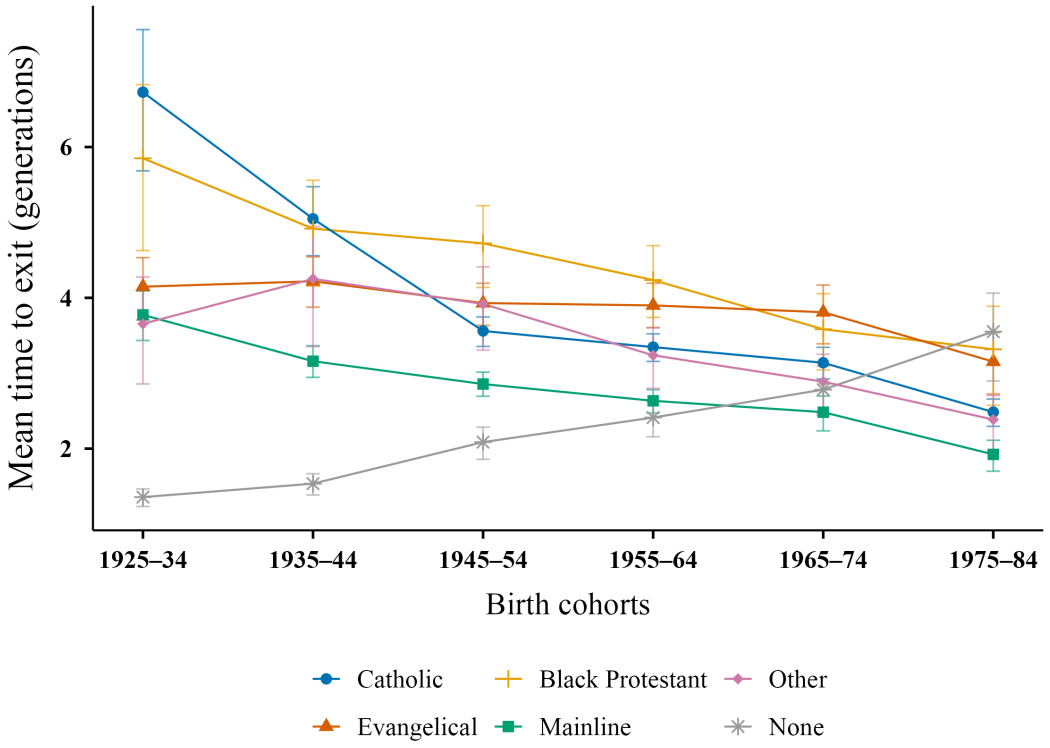
Mean time to exit (Figure 6) provides a complementary picture. Catholic and Black Protestant MTE_i both decline steadily across cohorts, with Black Protestant remaining somewhat higher throughout; Evangelical MTE_i is comparatively stable through the middle cohorts before dipping in the most recent one. Other MTE_i rises briefly before declining, unlike the remaining affiliated origins' more monotonic paths. MTE_i for None-origin lineages rises steadily across cohorts—an increase commensurate with the rise in None's diagonal persistence from 0.26 to 0.72—a trend that the memory-curve

Figure 5: Mean first passage time ($MFPT_{ij}$).



Notes: Each cell reports the bias-corrected bootstrap estimate of the expected number of transmission steps for a lineage starting in the origin category to first reach the destination category (2,000 within-survey-year resamples); darker cells indicate slower transit (log scale).

Figure 6: Mean time to exit from each origin category (MTE_i).



Notes: MTE_i is the expected number of transmission steps until a lineage's religion first differs from its origin state. Bias-corrected bootstrap estimates with 95% confidence intervals (2,000 within-survey-year resamples).

framing can obscure by conflating persistence with proximity to the stationary distribution.

The clearest MTE_i pattern is the reversal between Catholic and None: Catholic and Mainline become easier to exit, while None becomes increasingly persistent. By the latest cohort, None has the longest expected time to exit of any category, while Mainline has the shortest. This reversal is consistent with the argument that late-stage secularity is positively institutionalized rather than a residual, transitional category (Watts and Wilkins-Laflamme, 2026): None is no longer merely where people end up after leaving other traditions but has become, on this measure, the most durable identity in the system.

Taken together, the movement measures point to a common underlying change. Rising mobility is driven almost entirely by structural mobility–the reallocation of the population toward None–while exchange mobility remains flat. The $MFPT_{ij}$ results show why: reaching None becomes much faster from every origin, while reaching Catholic or Mainline becomes markedly slower. MTE_is reinforce this pattern: Catholic and Mainline become easier to leave, while None becomes increasingly persistent. Overall, movement is not simply increasing; it is being reorganized around None and Evangelical as two stable poles, at the expense of Catholic and Mainline.

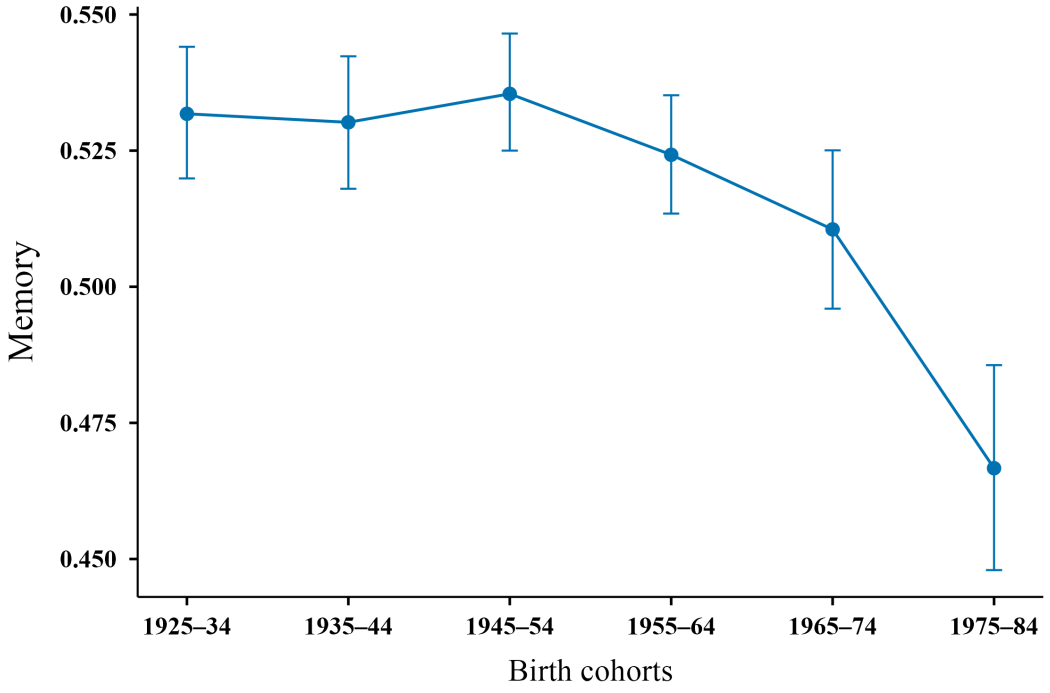
5 Memory Measures

Figure 7 summarizes memory in two complementary scalars per cohort. AIM (Panel A) is roughly flat across the first three cohorts before declining sharply in the youngest one. Mixing time t_{mix} (Panel B) is the more intuitive measure: it is the number of generations a lineage's religious composition would take to come within the specified tolerance (0.01) of its long-run distribution, if a cohort's transition rule were held fixed indefinitely. It fell by almost half over the period studied, from about 25 generations for the earliest cohort to about 13 for the most recent. Unlike AIM, this decline is fairly steady across the whole range rather than concentrated in the most recent cohorts, suggesting that the tail of the memory process was already shortening well before the one-step erosion captured by AIM picked up. This is broadly consistent with the eigenvalue results in Appendix B: since convergence to π_∞ proceeds at a rate governed by λ_2^t (Equation B.6), a smaller λ_2 is generally associated with a shorter mixing time. But mixing time declines more smoothly across the full period than λ_2 itself does, which shows that the staggered, two-stage pattern in that series is not the whole story, and that other eigenvalues and the constant c in the bound also shape how quickly the system as a whole converges.

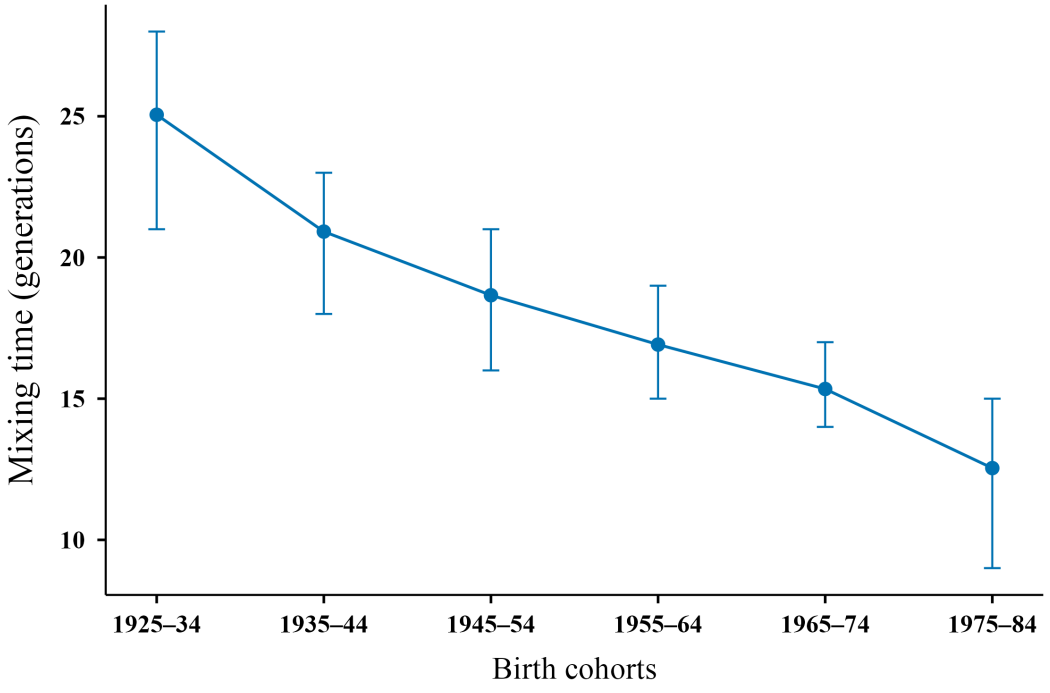
5.1 Category-specific Memory Measures

Figure 8 shows that Black Protestant has by far the highest memory of any origin throughout the period, well above every other category even as it declines from roughly 0.79 to 0.67. Catholic and Other form a second tier, both declining steadily from around 0.6–0.7 toward 0.5 by the most recent cohort. Mainline and Evangelical sit lower still and diverge over time: Mainline rises through the middle cohorts before falling back, while

Figure 7: Memory measures.



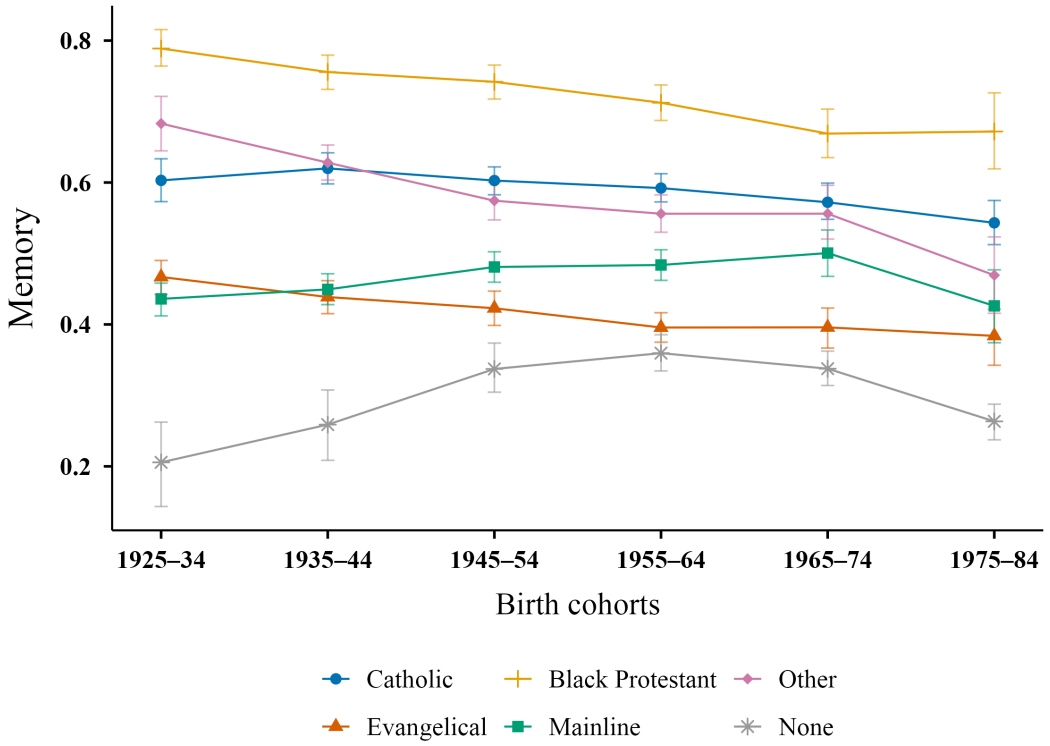
(A) Aggregate intergenerational memory (AIM, $t = 1$)



(B) Mixing time (t_{mix})

Notes: Panel A reports AIM, the π_0-weighted average of the category-specific individual memory measures at $t = 1$. Panel B reports the mixing time t_{mix}, the number of generations required for the transition matrix P^t to numerically converge to its steady state P^∞. Bias-corrected bootstrap estimates with 95% confidence intervals (2,000 within-survey-year resamples).

Figure 8: Intergenerational memory at $t = 1$ by origin ($IM_{i,1}$).



Notes: Bias-corrected bootstrap estimates with 95% confidence intervals (2,000 within-survey-year resamples).

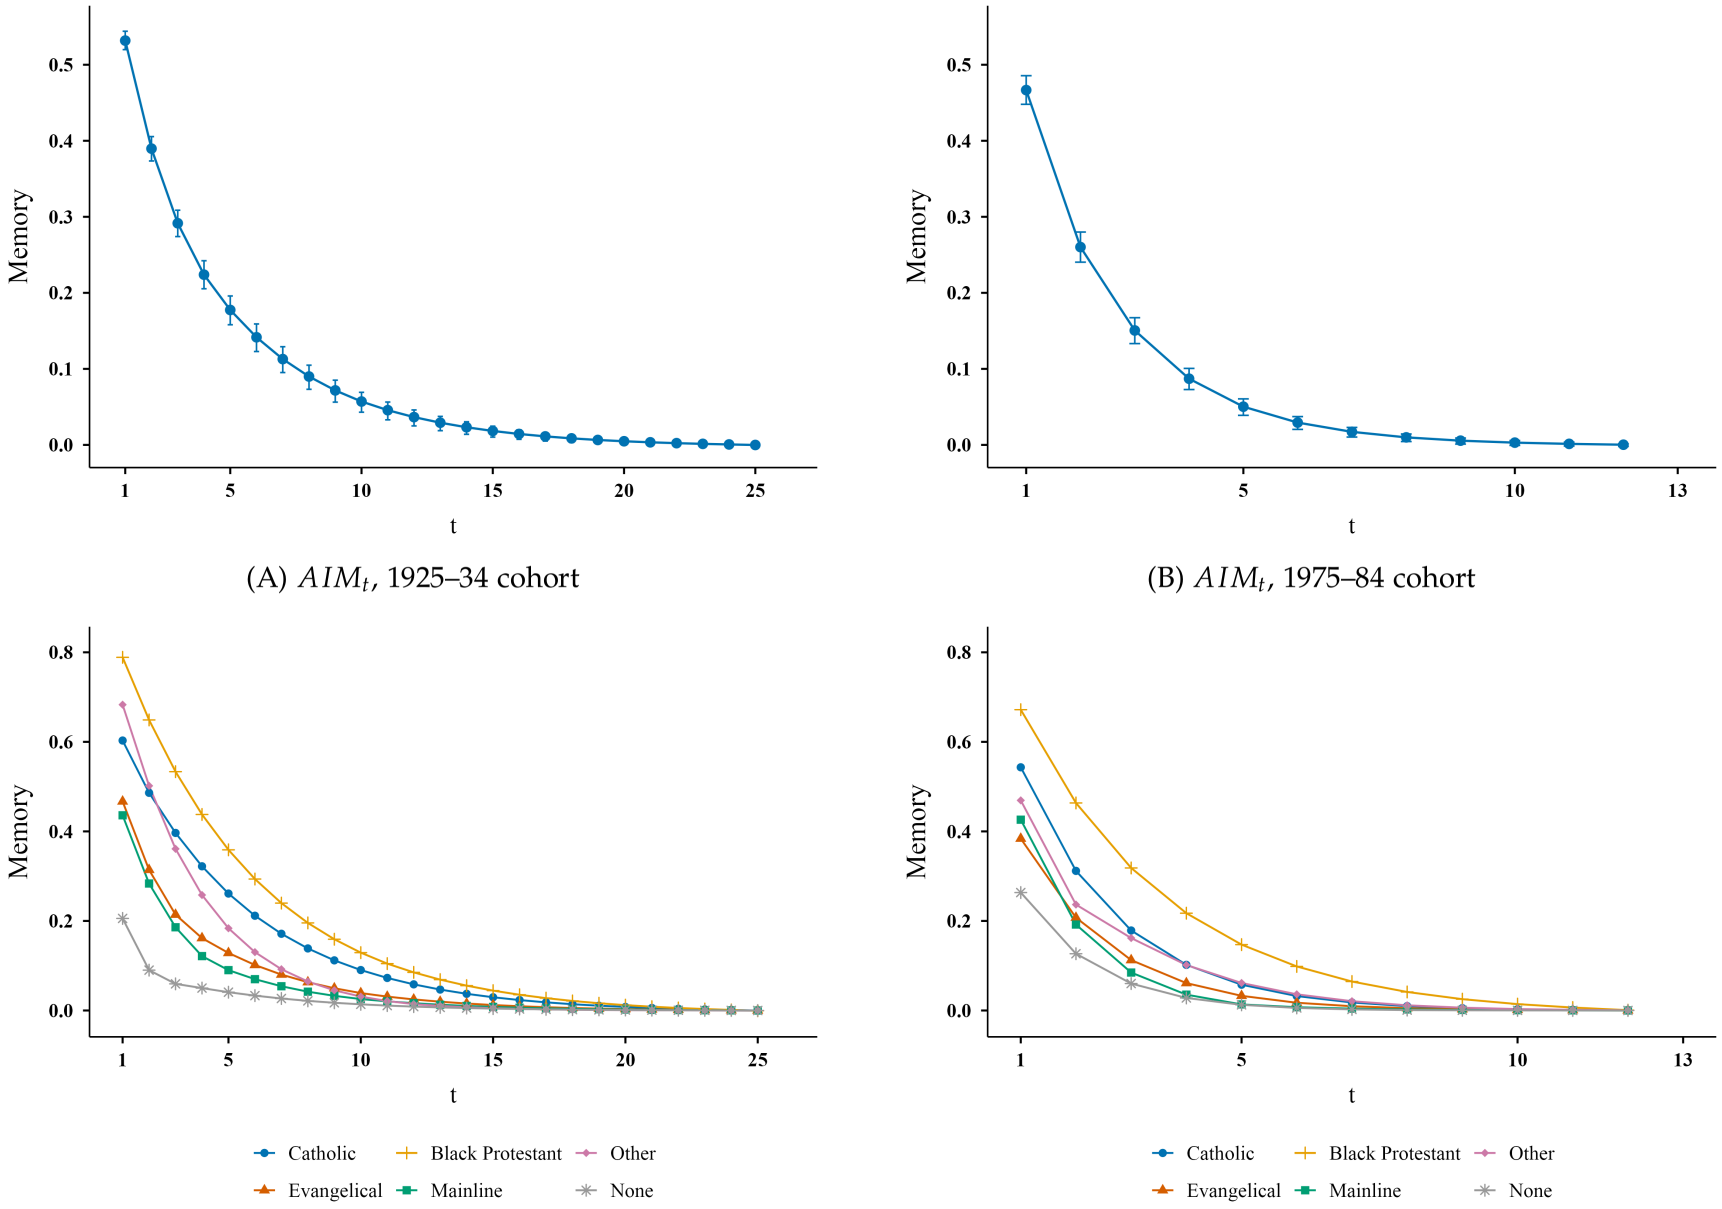
Evangelical declines steadily, so that by the final cohort Mainline sits above Evangelical rather than alongside it. None has the lowest memory in every cohort, but its path is non-monotonic, rising from about 0.21 in the earliest cohort to a peak near 0.36 around the middle of the period before falling back to about 0.27. Because low memory for None-origin lineages reflects proximity to π_∞, which itself assigns increasing weight to None, this pattern should be read as increasing closeness to a shifting target rather than as frequent exit from nonaffiliation.

The persistently high memory of Black Protestant origin is consistent with long-documented distinctiveness of Black Protestant religious identity in the United States, which has functioned historically as more than a matter of denominational preference; we do not pursue this connection further here. The convergence of the remaining affiliated categories toward lower, more similar memory levels, alongside rising memory for a None origin that is fast becoming the modal category, fits the broader claim that late-stage secularity in the United States is positively institutionalized rather than transitional (Watts and Wilkins-Laflamme, 2026; Stolz et al., 2025).

On Figure 9, we illustrate something implicit throughout this section: every point in Figures 7A and 8 is the $t = 1$ value of a full decay curve, one for each cohort's transition matrix. We plot that full curve for AIM_t and $IM_{i,t}$, comparing the earliest and latest cohorts directly. The contrast is visible in the horizontal axis alone: the 1925–34 curves extend to roughly 25 generations before flattening at zero, while the 1975–84 curves flatten by roughly 13, the same mixing times reported in Figure 7B. The category-level curves in Panels C and D show this compression is not uniform: Black Protestant's curve remains above every other category's for most of its range in both cohorts, while None's curve, already close to zero by $t = 2$ in the earliest cohort, is essentially flat throughout the latest one.

In sum, the memory measures point to a system in which aggregate persistence is falling, but unevenly: Black Protestant retains a distinctiveness that sets it apart from every other origin, Catholic and Other are converging downward from a once-higher tier, Mainline and Evangelical are pulling apart rather than moving together, and None remains the lowest-memory origin, with its memory falling back toward its early-cohort level even as its rising retention and mean time to exit indicate increasing durability.

Figure 9: Individual and aggregate memory across synthetic generations, earliest and latest cohorts.



(A) AIM_t, 1925–34 cohort

(B) AIM_t, 1975–84 cohort

(C) $IM_{i,t}$, 1925–34 cohort

(D) $IM_{i,t}$, 1975–84 cohort

Notes: This figure illustrates the evolution of AIM_t and $IM_{i,t}$ across generations for the 1925–34 and 1975–84 cohorts. Bias-corrected bootstrap estimates with 95% confidence intervals (2,000 within-survey-year resamples). Each curve flattens to zero at that cohort's mixing time, reported in Figure 7B.

5.2 Connecting Movement and Memory Results

Reading the category-specific movement and memory measures together clarifies how the same transition regime can produce apparently different forms of persistence. As shown in Appendix A, the implied stationary share of a tradition reflects the balance between the inflow it receives and the duration of the mass that enters it,

$$\pi_{j,\infty} = \mathrm{MTE}_j \sum_{i \neq j} \pi_{i,\infty} P_{ij},$$

while MFPT_{ij} measures the accessibility of that destination from a particular origin and

$$\mathrm{IM}_{i,t} = \mathrm{TV}((P^t)_{i,\cdot}, \pi_\infty)$$

measures the distance between an entire origin-specific destination distribution and the stationary benchmark. The measures need not move together because they describe, respectively, local duration, long-run prevalence, pairwise accessibility, and whole-row disequilibrium.

None illustrates how these dimensions of the transition process can reinforce one another. Its diagonal persistence rises sharply across cohorts, directly increasing $\mathrm{MTE}_{\mathrm{None}}$ and lengthening the expected duration of an uninterrupted chain in nonaffiliation. At the same time, inflow into None rises from every affiliated origin and the corresponding first-passage times fall dramatically, indicating that None becomes progressively easier to reach through the system's direct and indirect pathways. Under the stationary flow-balance relationship developed in Appendix A, greater stationary inflow and longer expected duration are the two channels through which the transition regime sustains a larger implied stationary share of None. None therefore becomes both an accessible destination and a durable category once entered.

Its origin-specific memory nevertheless remains comparatively low. This does not mean that None-origin lineages frequently exit nonaffiliation, nor does high retention mechanically produce low memory. Rather, the relevant criterion is the distance between the entire origin-specific destination profile and the stationary distribution:

$$\mathrm{IM}_{i,1} = \frac{1}{2} \left\| P_{i,\cdot} - \pi_\infty \right\|_1 .$$

Origin-specific memory is not determined by diagonal retention alone; it is low when the full destination profile of an origin is close to the stationary distribution implied by the transition regime. In the case of None, rising retention and inflow are the transition features that support an increasingly None-heavy stationary distribution. At the same time, the None-origin row itself places most of its probability on None and, among affiliated destinations, primarily on Evangelical. The full None-origin destination profile is therefore comparatively close to the stationary composition implied by the regime. Its low memory reflects proximity to this shifting stationary benchmark, not a high propensity to leave None.

Catholicism exhibits the converse configuration. Declining Catholic retention directly reduces $\mathrm{MTE}_{\mathrm{Catholic}}$, while negligible replacement inflow and rising first-passage times

into Catholicism indicate that it becomes increasingly difficult to reach from other origins. Under stationary flow balance, greater exit and weak replenishment are the two features of the transition regime that support the contraction of $\pi_{\text{Catholic},\infty}$ toward a small long-run share. Yet declining retention does not by itself make the Catholic-origin destination profile resemble this low-Catholic stationary distribution. Even in the youngest cohort, approximately 60 percent of Catholic-origin respondents remain Catholic, whereas the stationary distribution implied by the same transition regime assigns Catholic only a small share. The Catholic-origin row therefore substantially overrepresents Catholic and correspondingly underrepresents the system's principal long-run destinations, especially None and Evangelical. Because

$$\text{IM}_{\text{Catholic},1} = \frac{1}{2} \sum_j \left| P_{\text{Catholic},j} - \pi_{j,\infty} \right|,$$

these excesses and deficits across the full row generate a sizeable distance from π_∞. Catholic memory can consequently remain comparatively high even while Catholic retention and MTE decline.

The Hunter representation expresses the same contrast dynamically:

$$P_{i,\cdot} - \pi_\infty = \left((P^2)_{i,\cdot} - P_{i,\cdot} \right) N.$$

The interpretive direction here runs from the right-hand side to the left-hand side. The signed vector

$$(P^2)_{i,\cdot} - P_{i,\cdot}$$

records how the origin-specific destination profile would change under one additional iteration of the transition regime, and N propagates that one-step adjustment through the system's passage-time structure. Their product determines the remaining signed disequilibrium $P_{i,\cdot} - \pi_\infty$, whose total-variation norm is individual memory.

For Catholic origin, the one-step adjustment and the passage-time structure jointly generate a comparatively large propagated residual. Weak replenishment into Catholic is consistent with continued reallocation away from a one-generation row that remains much more Catholic-heavy than the stationary distribution, while the compensating expansions of other destinations are propagated through the full matrix N. For None origin, the corresponding product has a smaller total-variation norm and therefore produces a smaller residual relative to stationarity. Equivalently, the None-origin row lies closer to the long-run composition implied by the regime. The latter statement describes the resulting left-hand-side distance; the Hunter representation identifies the signed row adjustment and passage-time propagation from which that distance is reconstructed.

The observed level of memory is therefore jointly determined by the magnitude and direction of the signed one-generation row adjustment, the full matrix of passage-time weights, and cancellation among positive and negative pathway contributions. It cannot generally be attributed to retention, inflow, or the size of any single element of N alone.

None can therefore be easy to reach, difficult to leave, and low in origin-specific memory, while Catholic origin can be easy to leave, difficult to reach, and still comparatively

high in origin-specific memory. These are not competing assessments of a single notion of persistence. They describe distinct but connected features of the same transition regime: how probability enters a tradition, how long it remains there before first exit, how much of the implied long-run system the tradition occupies, and how distinctive the full destination profile of that origin remains relative to stationarity.

6 Heterogeneity

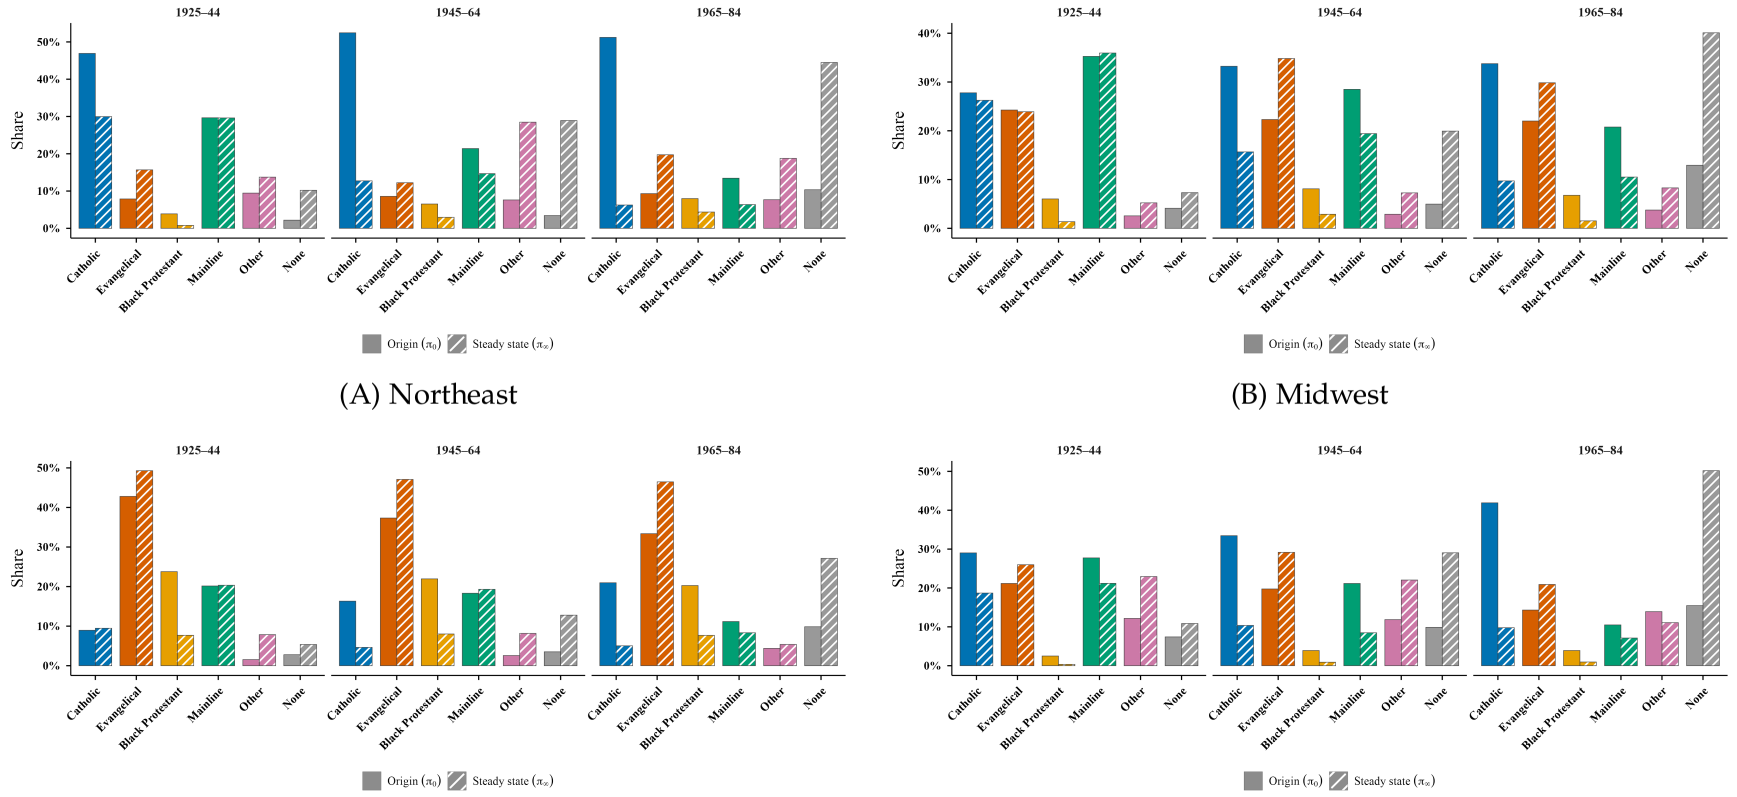
In this section, we examine heterogeneity in the main results along three dimensions: region of residence at age 16 (`REG16`, collapsed into the four Census regions, with respondents who lived abroad at age 16 excluded), current party identification (Democrat, Republican, or Other)[1], and sex. Because subgroup samples are smaller than the national sample, we use wider cohort bins to balance temporal resolution against estimability: 10-year bins for sex, 15-year bins for party identification, and 20-year bins for region. Sampling uncertainty is assessed using 2,500 respondent-level bootstrap resamples stratified by survey year; for the regional analyses, resampling is stratified jointly by survey year and age-16 region. Replicates that produce a non-ergodic transition matrix are treated as failed estimates and omitted. Cells for which more than 5% of replicates fail this criterion are not reported. Under this rule, three cells are omitted: West, 1925–44; Republican, 1970–84; and male, 1975–84. Among the reported cells, the failure rate does not exceed 2%.

6.1 Region

The regional contrasts are stable across every available comparison window (Figures 10 and 11). The South is the most persistent region — having the longest mixing times (t_{mix} from 26 to 16 generations), among the lowest overall mobility in every window (tied closely with the Northeast in the earliest cohort), and the only region whose dominant tradition is self-sustaining: Evangelical's π_∞ share stays at 0.46–0.49, even as its origin share falls from 0.42 to 0.33, representing the smallest relative gap between origin and steady state of any category in any region. Among the two cohort windows for which estimates are reported, the West is the least persistent and the most mobile (t_{mix} from 14 to 8 generations; OM_1 around 0.40), with None reaching half of its implied stationary distribution in the youngest window. The Northeast and Midwest fall in between, with the Northeast showing the starkest origin-to-steady-state collapse: Catholic origins are 0.51 of the youngest cohort but only 0.06 of its stationary distribution. The Midwest shows a comparable collapse running through Catholic (and Mainline): None's origin share rises only modestly, from about 4 to 13 percent, while its stationary share reaches roughly 40 percent by the youngest cohort.

[1]Because party identification is measured contemporaneously with adult religious affiliation, the party-specific transition matrices describe religious transitions among respondents currently identifying with each party. Their iteration represents a hypothetical regime conditional on current party identification, rather than an intergenerational trajectory of Democratic or Republican lineages.

Figure 10: Origin and stationary distributions by region at age 16.



(A) Northeast

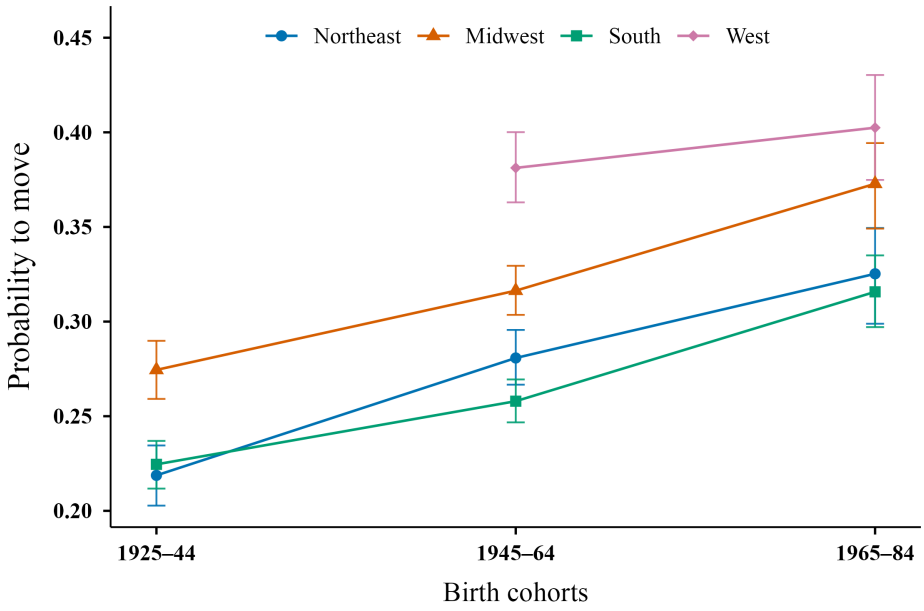
(B) Midwest

(C) South

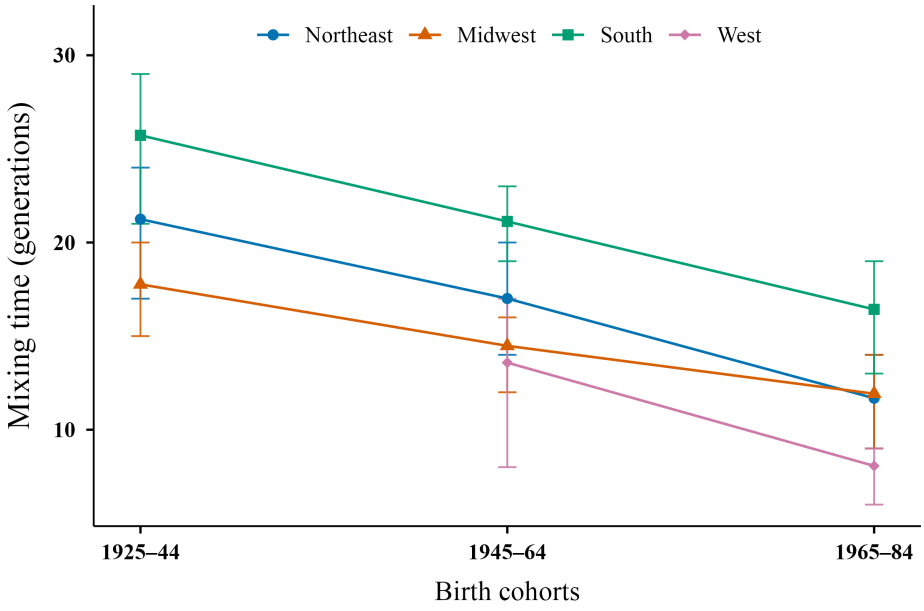
(D) West

Notes: Solid bars show the origin distribution (π_0); striped bars show the implied stationary distribution (π_∞), the left eigenvector of the cohort transition matrix associated with its leading eigenvalue, normalized to sum to one. π_∞ is computed directly from the baseline matrix, not from the bootstrap. 20-year birth cohorts.

Figure 11: Movement and memory summaries by region at age 16.



(A) Overall mobility (OM_1)



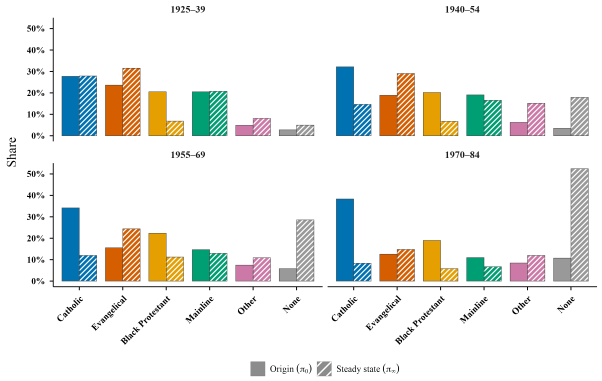
(B) Mixing time (t_{mix})

Notes: Panel A reports overall mobility at $t = 1$: the probability that a lineage's religion differs from its origin religion after one transmission step. Panel B reports the mixing time: the number of synthetic generations until the category distribution of every origin category converges to the implied stationary distribution. Bias-corrected bootstrap estimates with 95% confidence intervals (2,500 within-survey-year resamples; non-ergodic replicates excluded, cells with >5% rejections not reported — the West 1925–44 cell is not shown). 20-year birth cohorts.

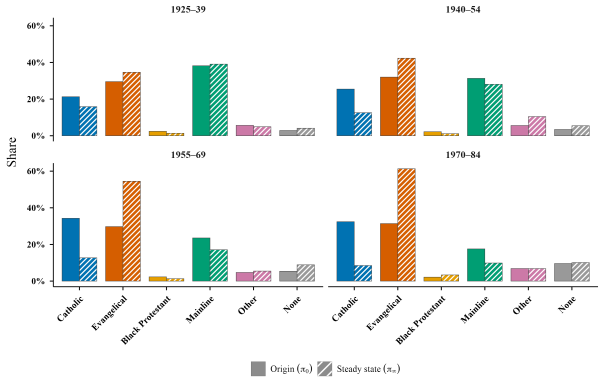
Taken together, these regional patterns show that the national reorganization toward None and Evangelical is not confined to any one part of the country, but proceeds at different speeds and through different channels. The South's persistence rests on Evangelical's unusually durable position; the West is furthest along in the shift toward None; and the Northeast and Midwest each show the sharpest origin-to-steady-state collapse in a different category, Catholic in the Northeast and None in the Midwest, rather than in the same category nationwide.

6.2 Political Views

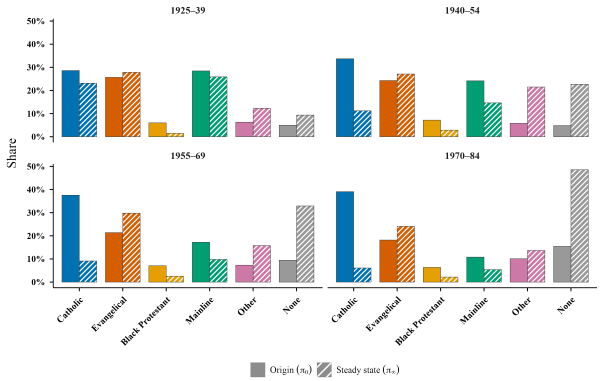
Figure 12: Origin and stationary distributions by party identification.



(A) Democrat



(B) Republican

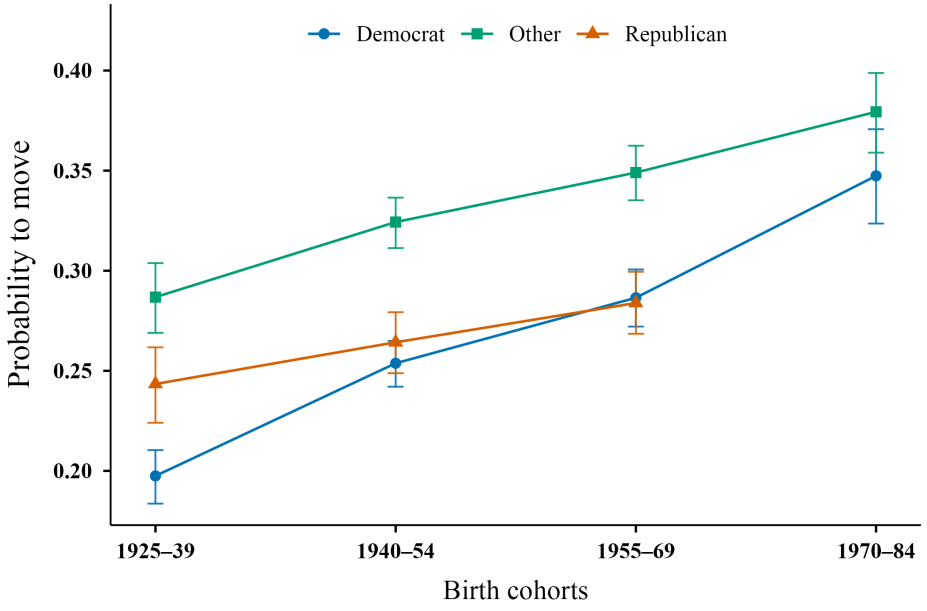


(C) Other

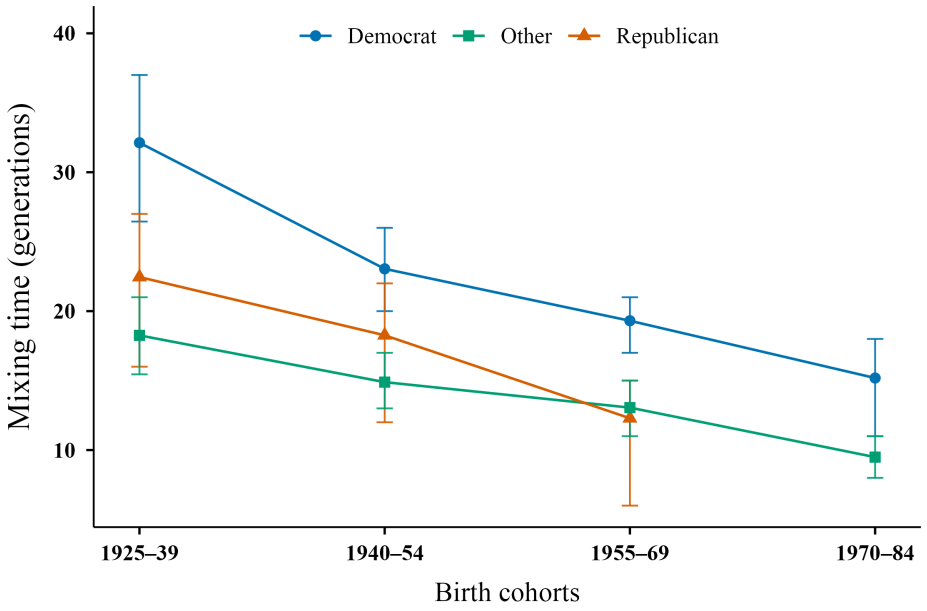
Notes: Solid bars show the origin distribution (π_0); striped bars show the implied stationary distribution (π_∞), the left eigenvector of the cohort transition matrix associated with its leading eigenvalue, normalized to sum to one. π_∞ is computed directly from the baseline matrix, not from the bootstrap. 15-year birth cohorts; party identification is defined as: strong/near Democrat, strong/near Republican, all others.

Partisan sorting is visible in both the levels and the destinations of persistence (Figures 12 and 13). Democrats are the most persistent group (t_{mix} from 32 to 15 generations) and Other identifiers the least (18 to 10), with Republicans in between, though Republican mixing time briefly dips just below Other's in the 1955–69 window (about 12 versus

Figure 13: Movement and memory summaries by party identification.



(A) Overall mobility (OM_1)



(B) Mixing time (t_{mix})

Notes: Panel A reports overall mobility at $t = 1$; Panel B reports the mixing time. Bias-corrected bootstrap estimates with 95% confidence intervals (2,500 within-survey-year resamples; non-ergodic replicates excluded, cells with >5% rejections not reported — the Republican 1970–84 cell is not shown). 15-year birth cohorts.

13 generations, well within overlapping confidence intervals). Overall mobility shows the mirror ordering, though Democratic mobility accelerates and nearly closes the gap with Republicans in the 1955–69 window (0.285 versus 0.283), the last cohort for which a Republican estimate is available.

The stationary distributions, however, reveal that the two parties are persistent toward opposite destinations: the Republican π_∞ concentrates on Evangelical, rising to 0.61 in the 1970–84 window (up from about 0.54 in 1955–69) while None stays near 0.10, whereas the Democratic π_∞ share of None climbs to 0.52 by 1970–84; Other identifiers show a similar but less extreme drift toward None, reaching about 0.48 in the same window.

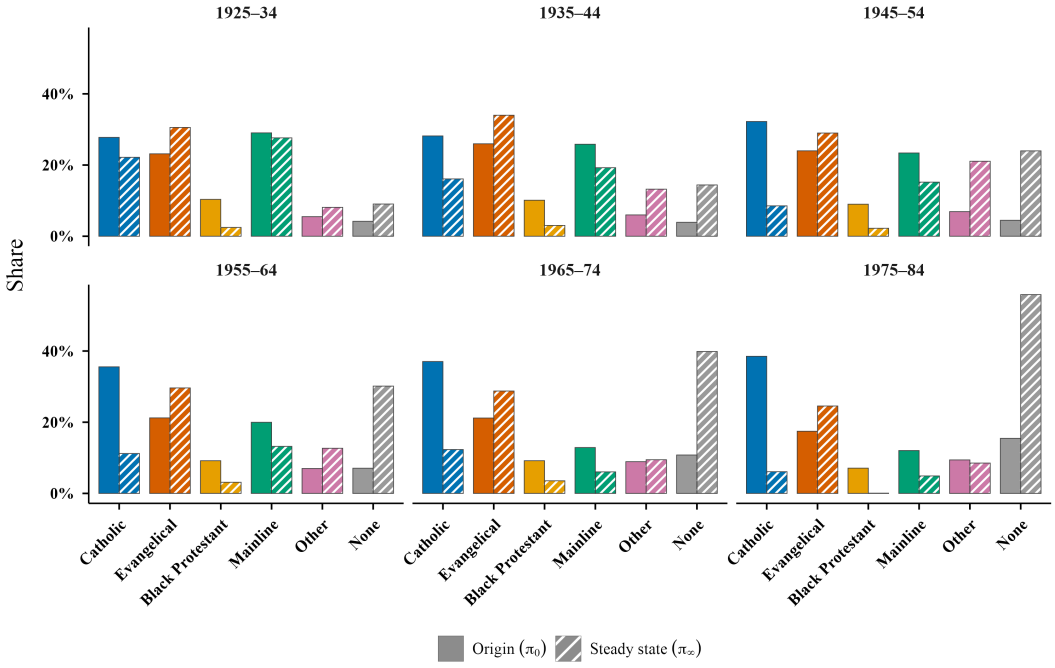
Overall, party identification sorts religious persistence into two distinct long-run destinations rather than simply into more or less persistent groups. Democrats and, to a lesser extent, Other identifiers are drifting toward None as their implied long-run identity, while Republicans are drifting toward Evangelical, with the parties' stationary distributions moving toward different poles.

6.3 Sex

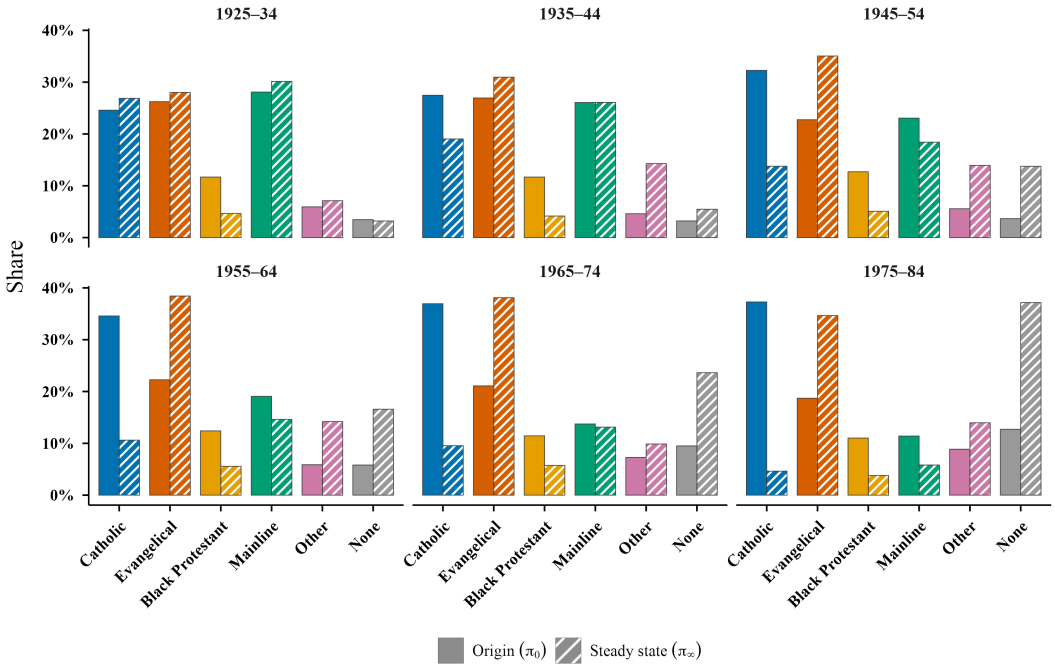
Women generally exhibit more persistence in religious affiliation: mixing times are longer for women in every early window (31 versus 22 generations in 1925–34), converging with men from the 1955–64 cohort onward, after which the two groups' confidence intervals overlap throughout the remaining windows for which both sexes are reported (Figure 15). Overall mobility rises for both sexes, with men consistently, if modestly, more mobile in every cohort for which both sexes are reported, from 1925–34 through 1965–74. The stationary distributions separate most clearly at the None margin (Figure 14): in the youngest windows None's π_∞ share reaches 0.56 for men against 0.37 for women, while women retain a larger Evangelical π_∞ share (0.35 against 0.25).

In sum, the sex-based patterns reveal that men are consistently more mobile at $t = 1$, yet women's mixing time is longer in the earliest cohorts before the two converge, so higher one-step mobility for men does not clearly translate into a faster overall approach to the stationary distribution. The two sexes also diverge in destination: men drift more toward None, women more toward Evangelical, distributing the same national two-pole reorganization differently across sexes.

Figure 14: Origin and stationary distributions by sex.



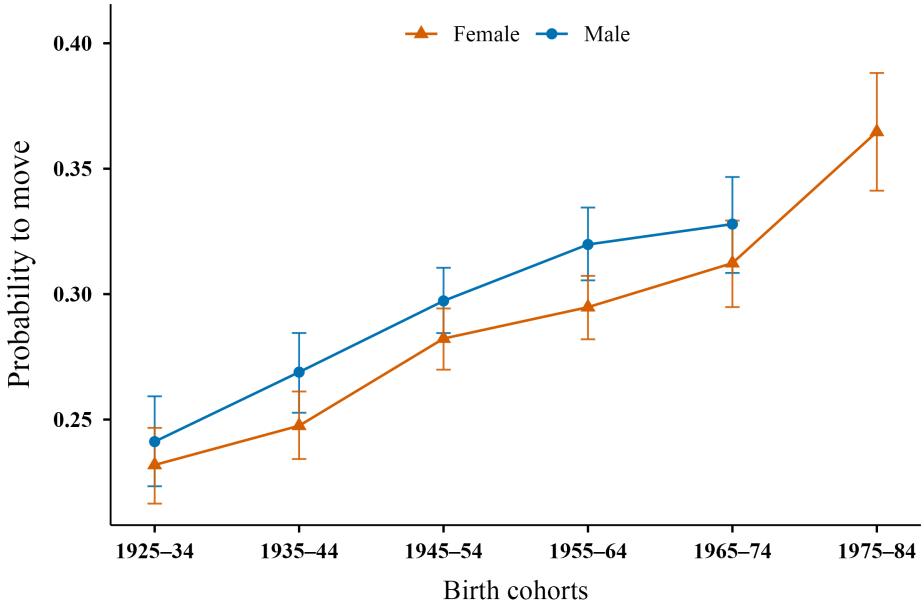
(A) Male



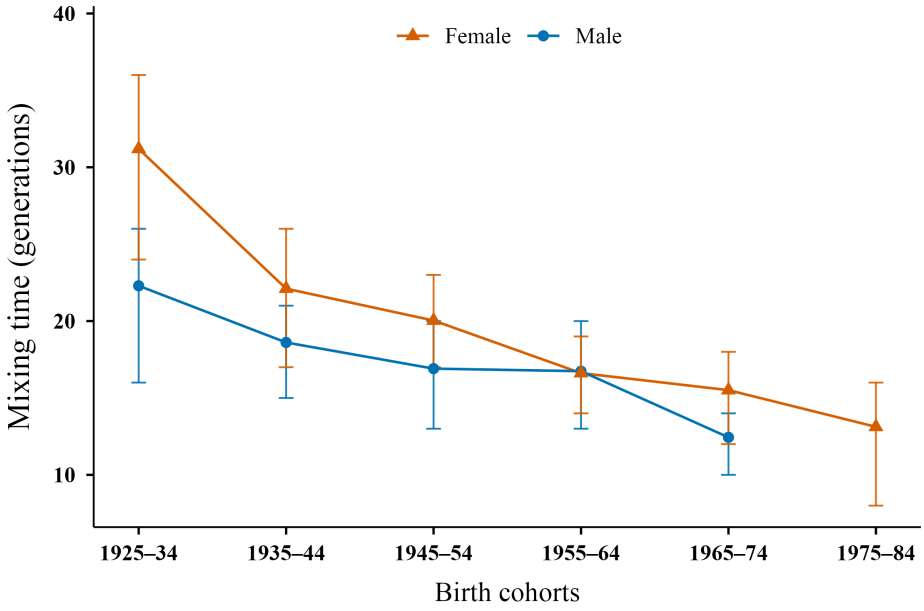
(B) Female

Notes: Solid bars show the origin distribution (π_0); striped bars show the implied stationary distribution (π_∞), the left eigenvector of the cohort transition matrix associated with its leading eigenvalue, normalized to sum to one. π_∞ is computed directly from the baseline matrix, not from the bootstrap. 10-year birth cohorts.

Figure 15: Movement and memory summaries by sex.



(A) Overall mobility (OM_1)



(B) Mixing time (t_{mix})

Notes: Panel A reports overall mobility at $t = 1$; Panel B reports the mixing time. Bias-corrected bootstrap estimates with 95% confidence intervals (2,500 within-survey-year resamples; non-ergodic replicates excluded, cells with >5% rejections not reported — the male 1975–84 cell is not shown). 10-year birth cohorts.

7 Conclusion

In this paper, we find that intergenerational religious persistence in the United States has weakened over the past six decades, but not uniformly. The resulting mobility patterns increasingly point toward two poles, with the population moving toward nonaffiliation and evangelical Protestantism while Catholicism and mainline Protestantism experience net losses with relatively little replacement. The estimated number of generations required for a cohort's transmission regime to approach its stationary distribution has fallen by nearly half, from about 25 generations for the earliest cohort to about 13 for the most recent. This pattern also differs increasingly along party lines: Democrats' long-run trajectories point more strongly toward nonaffiliation and Republicans' toward evangelical Protestantism, with the difference between the two widening over time. The same broader reorganization proceeds at different speeds and through different channels by region and sex, with the South's persistence associated with the particular durability of evangelical Protestantism and the West further along in the shift toward nonaffiliation, while men and women arrive at similar long-run patterns through different transition paths. None, once among the least persistent origins, has become one of the more durable states, while evangelical Protestantism remains the only affiliated tradition with positive net inflows in the cohorts we study.

These findings speak to, and extend, a literature that has established American religious decline as real and driven substantially by cohort replacement (Voas and Chaves, 2016), while leaving open the extent to which that decline is uniform across religious groups and populations (Schnabel and Bock, 2017; Voas and Chaves, 2018; Schnabel and Bock, 2018). Our results provide further evidence that the aggregate decline in persistence conceals substantial heterogeneity: it coexists with the emergence of two increasingly durable poles, evangelical Protestantism and nonaffiliation, each resilient in different ways. This pattern is consistent with accounts that treat secularization as staged and internally heterogeneous rather than uniform (Stolz et al., 2025), and with the argument that late-stage nonaffiliation may become positively institutionalized rather than remaining a residual or transitional category (Watts and Wilkins-Laflamme, 2026). Methodologically, our contribution is to provide a formal, Markov chain–based measurement framework that distinguishes *movement* from *memory*, thereby capturing dimensions of intergenerational change that are not directly summarized by the log-multiplicative models used in prior transition-based work on religious switching (Sherkat, 2001, 2014), and to apply this framework over a period that includes the recent rise in nonaffiliation.

Several limitations qualify these conclusions. First, we model each cohort's transition regime as fixed and time-homogeneous over a closed population. This specification does not incorporate migration, mortality, or fertility, each of which has historically shaped religious composition independently of intergenerational transmission and has buffered Catholic numbers even as U.S.-born Catholic persistence has fallen (Hout, 2016). We report a fertility-adjusted π_∞ in the appendix. A richer dataset would allow the framework to incorporate migration histories, fertility, and mortality by religious affiliation more rigorously. Second, race is a further dimension along which patterns of religious transmission may substantially differ. Our six-state scheme retains Black Protestant as a distinct category, but our sample does not support a fully race-stratified analysis at

a cohort resolution comparable to that reported here. Third, the None category also merits further study. A substantial share of respondents classified as unaffiliated move between affiliation and nonaffiliation over the life course, and our repeated cross-sections cannot distinguish stable nonaffiliation from this form of churn. Finally, our measures characterize the pace and direction of change in religious transmission. Identifying the mechanisms that generate a particular transition regime in a given cohort is a distinct question that the framework itself does not address.

References

Bengtson, V. L., N. M. Putney, and S. C. Harris (2013). *Families and Faith: How Religion Is Passed Down Across Generations*. New York: Oxford University Press.

Bhatia, R. (1997). *Matrix Analysis*. Graduate Texts in Mathematics. New York: Springer.

Blume, L., N. Cholli, S. Durlauf, and A. Lukina (2025). Immobility as Memory: Some New Approaches to Characterizing Intergenerational Persistence via Markov Chains. *Sociological Methods & Research 54*(4), 1396–1434.

Bok, J. (2025). Religious Exit Costs. *Journal for the Scientific Study of Religion 64*(2), 227–239.

Bottan, N. L. and R. Perez-Truglia (2015). Losing My Religion: The Effects of Religious Scandals on Religious Participation and Charitable Giving. *Journal of Public Economics 129*, 106–119.

Brauer, S. G. (2018). The Surprising Predictable Decline of Religion in the United States. *Journal for the Scientific Study of Religion 57*(4), 654–675.

Braunstein, R. (2022). A Theory of Political Backlash: Assessing the Religious Right's Effects on the Religious Field. *Sociology of Religion 83*(3), 293–323.

Butaeva, K., L. Chen, S. N. Durlauf, and A. F. Park (2026). A Tale of Two Transitions: Mobility Dynamics in China and Russia After Central Planning. *Journal of Applied Econometrics*.

Butaeva, K., S. N. Durlauf, and A. Shapoval (2026). Intergenerational Mobility in Late Qing Dynasty: Evidence from Northeast China. Working Paper 35313, National Bureau of Economic Research.

Campbell, D. E., G. C. Layman, and J. C. Green (2020). *Secular Surge: A New Fault Line in American Politics*. Cambridge Studies in Social Theory, Religion and Politics. Cambridge: Cambridge University Press.

Conlisk, J. and P. Sommers (1979). Eigenvector Status Proxies in Markov Chain Mobility Models. *Sociological Methods & Research 8*(2), 159–178.

Davern, M., R. Bautista, J. Freese, P. Herd, and S. L. Morgan (2025). General Social Survey 1972-2024 [Machine-readable data file]. Data accessed at https://gssdataexplorer.norc.org.

Dobrushin, R. L. (1956). Central Limit Theorem for Nonstationary Markov Chains. I. *Theory of Probability & Its Applications 1*(1), 65–80.

Erikson, R. and J. H. Goldthorpe (1992). *The Constant Flux: A Study of Class Mobility in Industrial Societies*. Oxford: Clarendon Press.

Featherman, D. L., F. L. Jones, and R. M. Hauser (1975). Assumptions of Social Mobility Research in the US: The Case of Occupational Status. *Social Science Research 4*(4), 329–360.

FitzGerald, F. (2017). *The Evangelicals: The Struggle to Shape America*. New York: Simon & Schuster.

Grinstead, C. M. and J. L. Snell (1997). *Introduction to Probability*. Providence, RI: American Mathematical Society.

Hayford, S. R. and S. P. Morgan (2008). Religiosity and Fertility in the United States: The Role of Fertility Intentions. *Social Forces 86*(3), 1163–1188.

Hollinger, D. A. (2022). *Christianity's American Fate: How Religion Became More Conservative and Society More Secular*. Princeton, NJ: Princeton University Press.

Hout, M. (2016). Saint Peter's Leaky Boat: Falling Intergenerational Persistence among U.S.-Born Catholics since 1974. *Sociology of Religion 77*(1), 1–17.

Hout, M. and C. S. Fischer (2002). Why More Americans Have No Religious Preference: Politics and Generations. *American Sociological Review 67*(2), 165–190.

Hout, M. and C. S. Fischer (2014). Explaining Why More Americans Have No Religious Preference: Political Backlash and Generational Succession, 1987–2012. *Sociological Science 1*, 423–447.

Hout, M., A. Greeley, and M. J. Wilde (2001). The Demographic Imperative in Religious Change in the United States. *American Journal of Sociology 107*(2), 468–500.

Hout, M. and O. P. Hastings (2016). Reliability of the core items in the general social survey: Estimates from the three-wave panels, 2006–2014. *Sociological Science 3*, 971–1002.

Hummer, R. A., R. G. Rogers, C. B. Nam, and C. G. Ellison (1999). Religious Involvement and U.S. Adult Mortality. *Demography 36*(2), 273–285.

Hungerman, D. M. (2013). Substitution and Stigma: Evidence on Religious Markets from the Catholic Sex Abuse Scandal. *American Economic Journal: Economic Policy 5*(3), 227–253.

Hunter, J. J. (2005). Stationary Distributions and Mean First Passage Times of Perturbed Markov Chains. *Linear Algebra and its Applications 410*, 217–243.

Kasselstrand, I., P. Zuckerman, and R. T. Cragun (2023). *Beyond Doubt: The Secularization of Society*. New York: New York University Press.

Kemeny, J. and J. Snell (1976). *Finite Markov Chains*, Volume 26 of *Undergraduate Texts in Mathematics*. New York: Springer-Verlag.

Kramer, S., C. Hackett, and M. Stonawski (2022). Modeling the Future of Religion in America. Technical report, Pew Research Center, Washington, DC. Scenarios, not forecasts. Concludes adult switching outweighs fertility and transmission for composition.

Lim, C., C. A. MacGregor, and R. D. Putnam (2010). Secular and Liminal: Discovering Heterogeneity Among Religious Nones. *Journal for the Scientific Study of Religion 49*(4), 596–618.

Margolis, M. F. (2018). How Politics Affects Religion: Partisanship, Socialization, and Religiosity in America. *The Journal of Politics 80*(1), 30–43.

Meyer, C. D. (2000). *Matrix Analysis and Applied Linear Algebra*. Philadelphia: Society for Industrial and Applied Mathematics (SIAM).

Pew Research Center (2025). Religious Switching. In *2023–24 U.S. Religious Landscape Study*. Washington, DC: Pew Research Center. Source for none-retention rising from 53 percent (2014) to about three-quarters, and the 6:1 Christian loss ratio.

Prais, S. J. (1955). Measuring Social Mobility. *Journal of the Royal Statistical Society. Series A (General) 118*(1), 56–66.

Puga-Gonzalez, I., D. Voas, L. Kiszkiel, R. J. Bacon, W. J. Wildman, K. Talmont-Kaminski, and F. L. Shults (2022). Modeling Fuzzy Fidelity: Using Microsimulation to Explore Age, Period, and Cohort Effects in Secularization. *Journal of Religion and Demography 9*, 111–137.

Schnabel, L. and S. Bock (2017). The Persistent and Exceptional Intensity of American Religion: A Response to Recent Research. *Sociological Science 4*, 686–700.

Schnabel, L. and S. Bock (2018). The Continuing Persistence of Intense Religion in the United States: Rejoinder. *Sociological Science 5*, 711–721.

Schnabel, L., S. Bock, and M. Hout (2024). Switch to Web-Based Surveys During COVID-19 Pandemic Left Out the Most Religious, Creating a False Impression of Rapid Religious Decline. *Sociology of Religion 85*(4), 429–453.

Schnabel, L., I. Horwitz, P. Hekmatpour, and C. Schleifer (2025). Breaking Free of the Iron Cage: The Individualization of American Religion. *Socius 11*, 1–27. NSYR 2003–2013 growth curves. Affiliation/belief crossover; socialization failure does not predict decline.

Sherkat, D. E. (2001). Tracking the Restructuring of American Religion: Religious Affiliation and Patterns of Religious Mobility, 1973–1998. *Social Forces 79*(4), 1459–1493.

Sherkat, D. E. (2014). *Changing Faith: The Dynamics and Consequences of Americans' Shifting Religious Identities*. New York: New York University Press.

Shorrocks, A. (1978). The Measurement of Mobility. *Econometrica: Journal of the Econometric Society 46*(5), 1013–1024.

Skirbekk, V., E. Kaufmann, and A. Goujon (2010). Secularism, Fundamentalism, or Catholicism? The Religious Composition of the United States to 2043. *Journal for the Scientific Study of Religion 49*(2), 293–310.

Smith, G. A., A. Cooperman, B. A. Alper, B. Mohamed, and M. Rotolo (2025). Decline of Christianity in the U.S. Has Slowed, May Have Leveled Off: 2023–24 Religious Landscape Study. Technical report, Pew Research Center, Washington, DC.

Sobel, M. E., M. Hout, and O. D. Duncan (1985). Exchange, Structure, and Symmetry in Occupational Mobility. *American Journal of Sociology 91*(2), 359–372.

Sommers, P. M. and J. Conlisk (1979). Eigenvalue Immobility Measures for Markov Chains. *Journal of Mathematical Sociology 6*(2), 253–276.

Sparkman, R., B. M. Wells, and B. Schapiro (2026, March). Addition of Religious Tradition in the GSS. General Social Survey Methodological Report 146, NORC at the University of Chicago.

Steensland, B., J. Z. Park, M. D. Regnerus, L. D. Robinson, W. B. Wilcox, and R. D. Woodberry (2000). The Measure of American Religion: Toward Improving the State of the Art. *Social Forces 79*(1), 291–318.

Stetzer, E. and R. P. Burge (2016). Reltrad Coding Problems and a New Repository. *Politics and Religion 9*(1), 187–190.

Stinespring, J. and R. T. Cragun (2015). Simple Markov Model for Estimating the Growth of Nonreligion in the United States. *Science, Religion and Culture 2*(3), 96–103.

Stolz, J., N. D. de Graaf, C. Hackett, and J.-P. Antonietti (2025). The Three Stages of Religious Decline around the World. *Nature Communications 16*, 7202.

Tocqueville, A. d. (2000). *Democracy in America*. University of Chicago Press.

Voas, D. (2008). The Continuing Secular Transition. In D. Pollack and D. V. A. Olson (Eds.), *The Role of Religion in Modern Societies*, pp. 25–48. London: Routledge.

Voas, D. (2009). The Rise and Fall of Fuzzy Fidelity in Europe. *European Sociological Review 25*(2), 155–168.

Voas, D. and M. Chaves (2016). Is the United States a Counterexample to the Secularization Thesis? *American Journal of Sociology 121*(5), 1517–1556.

Voas, D. and M. Chaves (2018). Even Intense Religiosity Is Declining in the United States: Comment. *Sociological Science 5*, 694–710.

Watts, G. and S. Wilkins-Laflamme (2026). The Secular Sacred Canopy: A Paradigm Revision in the Study of Secularization. *Sociological Theory*, 1–26.

Wodtke, G., W. Wang, K. Butaeva, and S. Durlauf (2026, February). Class Mobility in the Era of Rising Inequality: A Synthetic Dynasty Analysis. Working Paper 34800, National Bureau of Economic Research.

Woodell, B. and P. Schwadel (2020). Changes in Religiosity among Lesbian, Gay, and Bisexual Emerging Adults. *Journal for the Scientific Study of Religion 59*, 379–396.

Appendix

A Connecting Local Persistence, Stationary Prevalence, and Memory

Although mean time to exit, mean first-passage time, the stationary distribution, and individual memory are all functions of the same transition matrix P, they describe different features of the transition process. Mean time to exit measures the duration of an uninterrupted chain in a category; the stationary distribution describes the long-run prevalence implied by the entire transition regime; mean first-passage time measures the accessibility of one category from another; and individual memory measures the distance between an origin-specific destination distribution and the stationary distribution. Their relationship can be made explicit through stationary flow balance and the mean-first-passage-time representation developed by Hunter (2005).

Consider a target religious category j. The jth coordinate of the stationary condition

$$\pi_\infty = \pi_\infty P$$

is

$$\pi_{j,\infty} = \pi_{j,\infty} P_{jj} + \sum_{i \neq j} \pi_{i,\infty} P_{ij}. \tag{A.1}$$

Rearranging yields

$$\pi_{j,\infty}(1 - P_{jj}) = \sum_{i \neq j} \pi_{i,\infty} P_{ij}. \tag{A.2}$$

The left-hand side is stationary outflow from j: the long-run probability mass in j multiplied by its one-generation exit probability. The right-hand side is stationary inflow into j from all other categories. Thus, at the stationary distribution, total inflow and total outflow balance exactly.

Mean time to exit from j is

$$\mathrm{MTE}_j = (1 - P_{jj}) \sum_{t=1}^{\infty} t (P_{jj})^{t-1} = \frac{1}{1 - P_{jj}}. \tag{A.3}$$

Combining Equations (A.2) and (A.3) gives

$$\pi_{j,\infty} = \mathrm{MTE}_j \sum_{i \neq j} \pi_{i,\infty} P_{ij} \tag{A.4}$$

Equation (A.4) has a stock–flow interpretation:

$$\text{stationary prevalence of } j = \text{stationary inflow into } j \times \text{expected duration in } j.$$

A category can therefore acquire a large implied stationary share because it is difficult to leave, because it receives substantial inflow, or because both conditions hold. Conversely, a category that is easy to leave and receives little replacement inflow will have a small

implied stationary share. This is an accounting identity for the full transition system rather than a one-cell comparative-static result: the weights $\pi_{i,\infty}$ are themselves jointly determined by all elements of P. As elsewhere in the paper, π_∞ denotes the stationary distribution implied by indefinitely holding the cohort-specific transition regime fixed; it is not a forecast of future religious composition.

An important notational distinction is required here. The scalar power

$$(P_{jj})^t$$

in Equation (A.3) is the probability of remaining continuously in j for t transitions. Once the process leaves j, the chain has ended, even if it later returns. By contrast, the matrix-power entry

$$(P^r)_{jj} = \Pr(X_r = j \mid X_0 = j)$$

allows the process to leave and return before horizon r. For example,

$$(P^2)_{jj} = P_{jj}^2 + \sum_{k \neq j} P_{jk} P_{kj}.$$

Mean time to exit is therefore a local measure governed by the single diagonal element P_{jj}, whereas the matrix powers $(P^r)_{ij}$ reflect all direct and indirect pathways through the system.

Mean first-passage time provides a complementary measure. Let $m_{ij} = \text{MFPT}_{ij}$ denote the expected number of transition steps required to first reach j when starting from i. Following Hunter (2005), define the matrix $N = [n_{ij}]$ by

$$n_{ij} = (1 - \delta_{ij}) \pi_{j,\infty} m_{ij}, \tag{A.5}$$

so that $n_{jj} = 0$ and, for $i \neq j$,

$$n_{ij} = \sum_{r=0}^{\infty} \left[(P^r)_{jj} - (P^r)_{ij} \right] = \pi_{j,\infty} m_{ij}. \tag{A.6}$$

The first equality in Equation (A.6) compares two copies of the same transition process. One begins in target category j, while the other begins in category i. At horizon r, $(P^r)_{jj} - (P^r)_{ij}$ is the signed difference in their probabilities of occupying j. Although an individual horizon-specific difference may be negative, both probabilities converge to $\pi_{j,\infty}$ and their infinite sum is positive. The sum records the net cumulative j-occupancy head start obtained by beginning in j rather than in i.

The second equality converts that occupancy head start into passage time. The expected delay in reaching j is m_{ij} generations, and $\pi_{j,\infty}$ is the long-run amount of j-occupancy generated per generation. Their product is therefore the cumulative amount of j-occupancy displaced by beginning in i rather than in j. A falling m_{ij} means that j has become easier to reach from i, either through a larger direct transition probability P_{ij} or through faster indirect routes. A rising $\pi_{j,\infty}$ means that j occupies more of the implied

long-run system. These forces have opposing effects on n_{ij}:

$$n_{ij} = \pi_{j,\infty} m_{ij}.$$

For two cohort-specific regimes, denoted without and with primes, the finite change can be written exactly as

$$n'_{ij} - n_{ij} = (\pi'_{j,\infty} - \pi_{j,\infty}) m_{ij} + \pi'_{j,\infty}(m'_{ij} - m_{ij}). \tag{A.7}$$

Thus, if the stationary share of j rises while the time required to reach j falls, the direction of the change in n_{ij} cannot be inferred from either component alone.

The occupancy mechanism can also be seen recursively:

$$(P^{r+1})_{ij} = (P^r)_{ij} P_{jj} + \sum_{k \neq j} (P^r)_{ik} P_{kj}. \tag{A.8}$$

Next-period occupancy in j consists of mass already in j that is retained there plus mass entering from every other state. Stronger retention tends to preserve the occupancy advantage of a process that begins in j, whereas stronger direct or indirect inflow allows a process beginning in i to catch up more quickly. Equation (A.6) summarizes the net result of this competition over all horizons.

Hunter introduces N in the study of changes in stationary distributions under perturbations of a transition matrix. The same matrix identity also provides a representation of transient disequilibrium under one fixed transition law. Let

$$\Pi = \mathbf{1}\pi_\infty,$$

where every row of Π equals π_∞. Hunter's matrix satisfies

$$(I - P)N = \Pi - I.$$

Consequently, for any probability row vector q,

$$q - \pi_\infty = (qP - q)N. \tag{A.9}$$

The left-hand side is the signed stock of disequilibrium relative to the stationary distribution. The first term on the right,

$$qP - q,$$

is the signed one-generation redistribution generated by P, while N propagates that redistribution through the system's passage-time structure.

Applying Equation (A.9) to a synthetic dynasty founded in religious origin i, with current distribution $q = P^t_{i,\cdot}$, gives

$$(P^t)_{i,\cdot} - \pi_\infty = \left((P^{t+1})_{i,\cdot} - (P^t)_{i,\cdot}\right) N. \tag{A.10}$$

Therefore, the individual-memory measure in Equation (8) has the equivalent represen-

tation

$$\mathrm{IM}_{i,t} = \frac{1}{2} \left\| \left((P^{t+1})_{i,\cdot} - (P^{t})_{i,\cdot} \right) N \right\|_1 . \tag{A.11}$$

Componentwise,

$$(P^{t})_{ij} - \pi_{j,\infty} = \sum_{\ell} \left[(P^{t+1})_{i\ell} - (P^{t})_{i\ell} \right] n_{\ell j}. \tag{A.12}$$

Thus, the signed discrepancy in destination j is produced by the expansion or contraction of every intermediate category ℓ, weighted by the target-specific passage-time leverage $n_{\ell j}$.

The order of operations in Equation (A.11) is important. The pathway contributions in Equation (A.12) are first combined with their signs intact, and only then is the absolute value taken:

$$\mathrm{IM}_{i,t} = \frac{1}{2} \sum_{j} \left| \sum_{\ell} \left[(P^{t+1})_{i\ell} - (P^{t})_{i\ell} \right] n_{\ell j} \right| . \tag{A.13}$$

A relatively high value of $\mathrm{IM}_{i,t}$ may therefore reflect a large signed one-step redistribution, substantial passage-time leverage, limited cancellation among pathways, or some combination of these mechanisms. It does not imply that every relevant element of N is individually large.

These identities clarify the division of labor among the category-specific measures. MTE_j measures local chain duration and depends only on P_{jj}. The stationary share $\pi_{j,\infty}$ reflects the balance between systemwide inflow and outflow. The pairwise passage time m_{ij} measures how accessible j is from a particular origin i through all direct and indirect routes. Finally, $\mathrm{IM}_{i,t}$ measures how far the entire destination distribution of origin i remains from the full stationary distribution. At $t = 0$,

$$\mathrm{TV}(e_i, \pi_\infty) = 1 - \pi_{i,\infty},$$

so a larger stationary share mechanically places the pure origin category closer to the stationary distribution. At $t \geq 1$, however, the entire row $P^t_{i,\cdot}$ matters. A category can therefore be difficult to leave but low in memory, or easy to leave but high in memory, without contradiction.

B Alternative Memory Measures

The memory measures in the main text describe a representative origin (AIM) or each origin separately (IM). In this appendix, we report two further measures that have been used to estimate social mobility (Blume et al., 2025; Wodtke et al., 2026). As in the main text, distances are measured using total variation norms throughout. We begin with the 'worst-case' measure[1], which compares dynasties originating from different religious traditions directly,

$$d'(t) = \max_{i,j} \| (P^{t})_{i,\cdot} - (P^{t})_{j,\cdot} \|_{TV}. \tag{B.1}$$

[1] 'Worst-case', in this context, connotes no normative judgment.

The special case $d'(1)$ is known as Dobrushin's coefficient, $\delta(P)$ (Dobrushin, 1956), and measures the distance in affiliation between two individuals who were raised in the most dissimilar traditions. $\delta(P)$ has the advantage of not relying on π_∞, meaning that it bypasses concerns about inter-cohort comparability, given large differences in π_∞.

Next, we examine the second-largest eigenvalue of the transition matrix (Sommers and Conlisk, 1979; Conlisk and Sommers, 1979; Shorrocks, 1978).

To provide an intuitive explanation of the second-largest eigenvalue as a measure of memory, recall the standard spectral decomposition for diagonalizable matrices (Meyer, 2000, pp. 517–518; Bhatia, 1997, pp. 4–5):

$$P = V\Lambda V^{-1} = V\Lambda U = \sum_i \lambda_i \mathbf{v}_i \mathbf{u}_i, \tag{B.2}$$

where V is the matrix whose columns are the right eigenvectors $\mathbf{v}_i$, Λ is the diagonal matrix of eigenvalues λ_i, and $U = V^{-1}$ is the matrix whose rows are the left eigenvectors $\mathbf{u}_i$, normalized such that $\mathbf{u}_i \mathbf{v}_k = \delta_{ik}$. Recall that a vector $\mathbf{v}$ is a right eigenvector if $P\mathbf{v} = \lambda\mathbf{v}$, while a vector $\mathbf{u}$ is called a left eigenvector if $\mathbf{u}P = \lambda\mathbf{u}$.

Then, for the t-th generation in our setting, the expression can be written as:

$$P^t = \sum_i \lambda_i^t \mathbf{v}_i \mathbf{u}_i. \tag{B.3}$$

To derive the expression for the distribution of religions in generation t, we apply the matrix P^t to the initial row distribution μ_0:

$$\mu_t = \mu_0 P^t = \mu_0 \sum_i \lambda_i^t \mathbf{v}_i \mathbf{u}_i = \sum_i \lambda_i^t (\mu_0 \mathbf{v}_i) \mathbf{u}_i = \sum_i \lambda_i^t \alpha_i \mathbf{u}_i, \tag{B.4}$$

where $\alpha_i = (\mu_0 \mathbf{v}_i)$ denotes the scalar obtained by applying the initial distribution to the right eigenvector $\mathbf{v}_i$. Recall that the eigenvalues of P are ordered by modulus such that

$$|\lambda_1| = 1 > |\lambda_2| \geq \cdots \geq |\lambda_n|.$$

Because P is row-stochastic, the dominant eigenvalue λ_1 equals 1, with corresponding right eigenvector $\mathbf{v}_1 = \mathbf{1}$ and normalized left eigenvector $\mathbf{u}_1 = \mu_\infty$, the steady-state distribution (Sommers and Conlisk, 1979). Since $\mu_0 \mathbf{1} = 1$, we can rewrite Equation (B.4) in a form that is more intuitive for our analysis:

$$\mu_t = \mu_\infty + \lambda_2^t \alpha_2 \mathbf{u}_2 + \cdots + \lambda_n^t \alpha_n \mathbf{u}_n. \tag{B.5}$$

For the purpose of studying mobility, the transient left eigenvectors $\mathbf{u}_i$, for $i \geq 2$, can be interpreted as the principal directions along which the current distribution μ_t converges toward the steady-state distribution μ_∞. Each $\mathbf{u}_i$ identifies a "barrier" that may slow this evolution, and the elements of $\mathbf{u}_i$ correspond to category-specific coordinates along that direction. The scalar coefficients α_i measure the initial distribution's loading, or relative weight, on each barrier, while the eigenvalues λ_i indicate the persistence

of the corresponding barrier over generations t. Since $|\lambda_i| < 1$ for $i \geq 2$, as $t \to \infty$, all terms multiplied by λ_i^t converge to zero. Modes associated with smaller eigenvalue moduli decay more quickly, while the mode associated with $|\lambda_2|$ is asymptotically the slowest-decaying component, provided that its loading is nonzero. Consequently, the mode associated with $|\lambda_2|$ can be interpreted as the dominant, most persistent pattern constraining intergenerational mobility, and $|\lambda_2|$ governs the asymptotic rate of persistence in the system. Larger values of $|\lambda_2|$ correspond to slower convergence and, therefore, lower mobility conceptualized as memory. Formally, the convergence rate can be bounded as

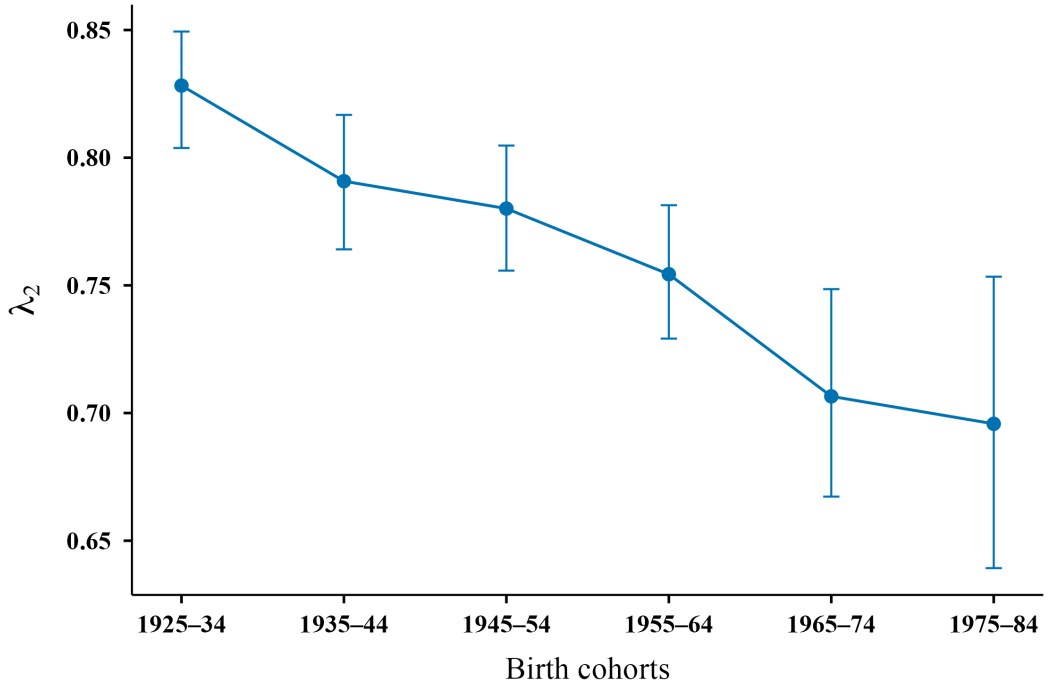
$$\|\mu_\infty - \mu_t\| \leq c\,|\lambda_2|^t, \quad \text{where} \quad c = \sum_{i=2}^{n} |\alpha_i|\,\|\mathbf{u}_i\|, \tag{B.6}$$

with c a constant depending on the initial distribution and the left eigenvectors.

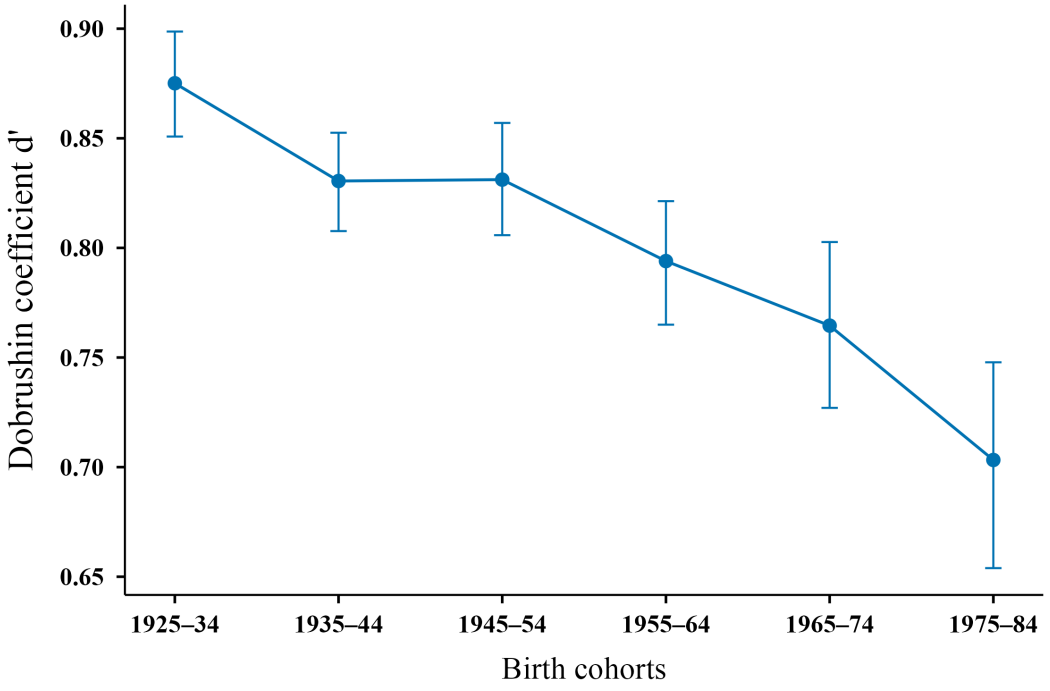
Figure B.1 presents the two headline scalars. λ_2 (Panel A) falls from about 0.83 in the 1925–34 cohort to about 0.70 in the 1975–84 cohort, but not at a constant rate: the decline is gradual through the first four cohorts, then steepens sharply between the 1955–64 and 1965–74 cohorts, after which it levels off, with confidence intervals widening substantially in the two youngest cohorts. As discussed in the main text, this bound also governs mixing time (t_{mix}, Figure 7B): because convergence to π_∞ proceeds at rate λ_2^t, a smaller λ_2 is generally associated with a shorter mixing time, though the constant c and the remaining eigenvalues also contribute.

Dobrushin's $d'(1) = \delta(P)$, by contrast, is a worst-case pairwise one-step measure that yields an upper bound on the tolerance-based mixing time when $d'(1) < 1$. Dobrushin's $d'(t = 1)$ (Panel B) follows a related but distinct path: it is roughly flat between the 1935–44 and 1945–54 cohorts (both around 0.83) before declining steadily, with the steepest single drop occurring in the final cohort. Both series confirm that the dependence of destination on origin has weakened over the period, and both show that this weakening is concentrated later in the period rather than proceeding at a constant pace throughout—a pattern that also appears in the AIM trend reported in the main text (Figure 7).

Figure B.1: Aggregate memory: second-largest eigenvalue and Dobrushin coefficient.



(A) λ_2



(B) Dobrushin coefficient $d'(t = 1)$

Notes: Panel A reports the modulus of the second-largest eigenvalue of P; lower values indicate faster mixing speed. Panel B reports the Dobrushin coefficient: the total-variation distance between the two most dissimilar origin rows; lower values indicate weaker dependence of destination on origin. Bias-corrected bootstrap estimates with 95% confidence intervals (2,000 within-survey-year resamples).

C Differential Fertility Rates

Different religious traditions have substantially different fertility rates which are not captured by the transition matrix, and these can be important in determining religious composition (Hout et al., 2001; Hayford and Morgan, 2008). To incorporate it into our steady-state calculations, let μ be the vector whose ith entry is the mean number of children reported by women over forty whose affiliation is tradition i, and consider the modified recursion

$$\pi_{t+1} = \frac{(\pi_t \odot \mu)\, P}{[(\pi_t \odot \mu)\, P]\mathbf{1}} \tag{C.1}$$

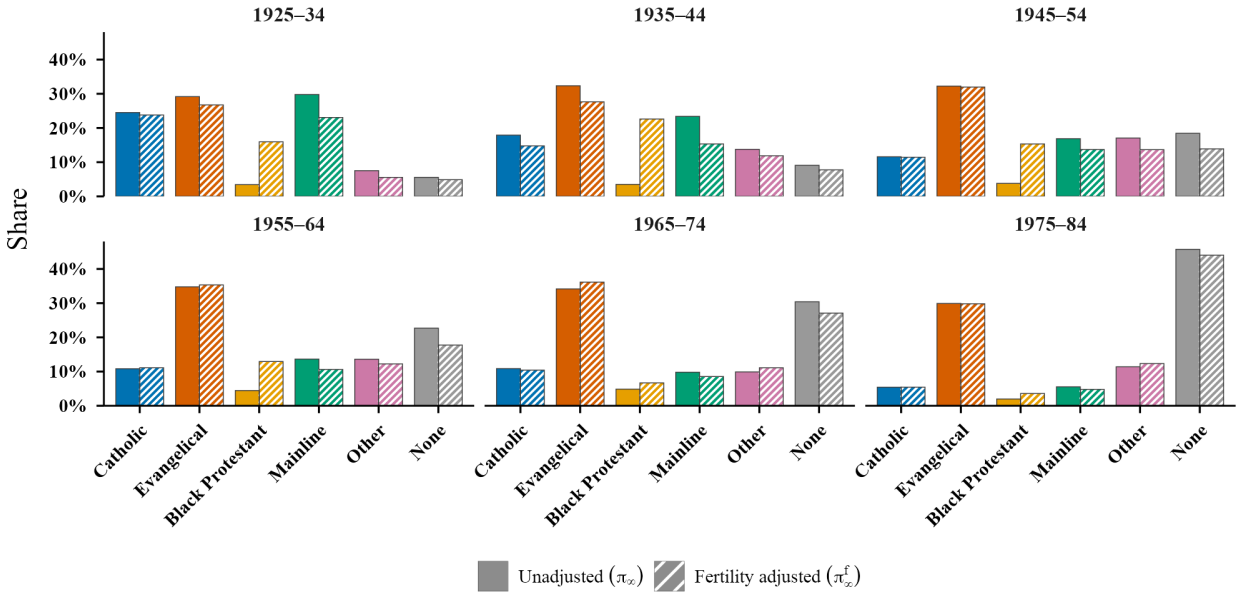
$$\pi_f^{\infty} = \frac{(\pi_f^{\infty} \odot \mu)\, P}{[(\pi_f^{\infty} \odot \mu)\, P]\mathbf{1}} \tag{C.2}$$

where $\odot$ denotes the Hadamard, or element-wise product: $(\pi \odot \mu)_i = \pi_i \mu_i$, such that each category's share is scaled by its own fertility before the transition matrix is applied. The resulting vector is renormalized to a distribution.

Figure C.1 compares π_∞ with π_∞^f for each cohort. For four of six categories, the adjustment changes almost nothing. Catholic, Evangelical, Other, and None shares move by no more than a few percentage points in each cohort, and their ordering remains unchanged throughout. Mainline decreases by 7–8 percentage points in the first two cohorts, with its rank decreasing in each. The second exceptional category is Black Protestant, whose steady-state share rises from roughly 4% to 23% in the 1935–44 cohort. Two features combine to produce this: first, the fact that this category's observed fertility is the highest in that cohort (3.74, compared to 3.22 for the next highest category, Catholic), and second, the fact that the unadjusted π_∞ is low due to extremely low inflow, meaning that the Black Protestant steady-state share rests almost entirely on retention, rather than recruitment. The fertility advantage is compounded within an essentially closed loop each generation, so a fertility differential of 0.5 produces a larger effect on the steady-state share than it does for more open traditions. Substantively, Black Protestantism is reproduced demographically rather than through conversion, and the unadjusted measure understates its position for that reason.

Two qualifications: first, because μ is estimated from women over forty, the 1975–84 cohort only contributes those observed at ages 40–49 in the 2015–2024 waves, so its fertility vector rests on far fewer observations. Second, the Black Protestant result is sensitive because it depends heavily on a single entry of μ estimated from a comparatively small subsample.

Figure C.1: Standard and fertility-adjusted stationary distributions



Notes: This figure compares the implied stationary distribution under the observed transition matrix (π_∞) with the fertility-adjusted stationary distribution (π^f_∞) by 10-year birth cohort. Fertility is measured as mean children reported by women over forty in each cohort in each tradition.

D Age-Standardized Estimates

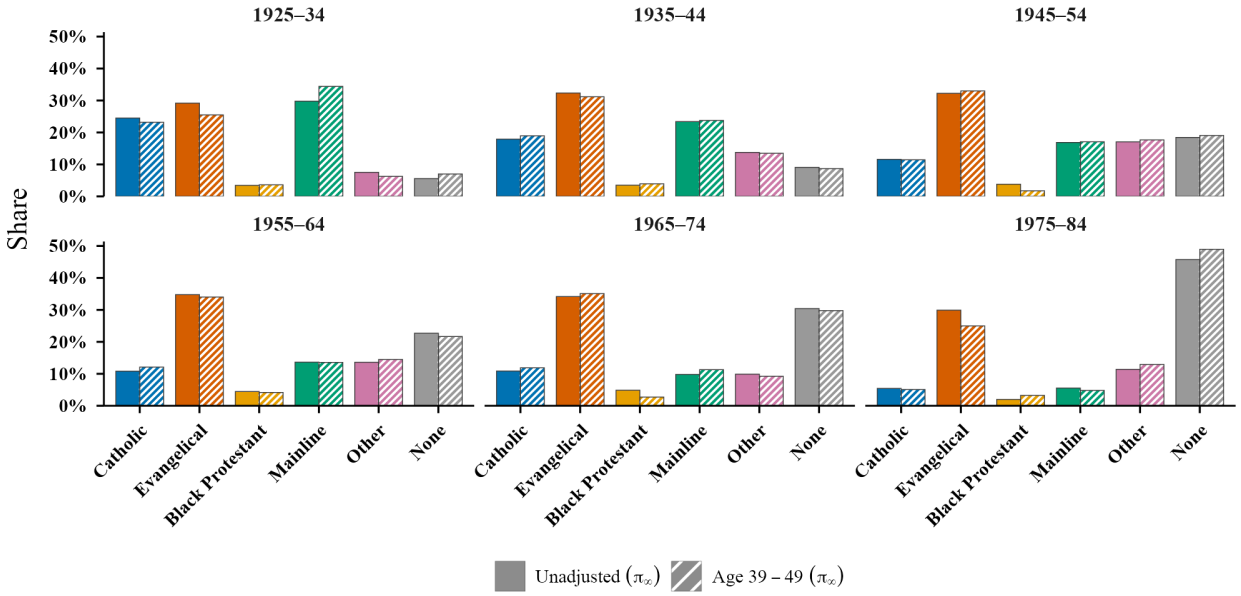
Birth cohorts in the sample are observed across different portions of the 30–75 age range, with the earlier cohorts observed disproportionately at older ages, and the more recent cohorts at younger ones. As mentioned above, the cohorts share a common support between ages 39–49. This appendix recomputes π_∞, OM_1, and t_{mix} for this subsample, thus holding age composition fixed.

Restricting the age range in this way discards a substantial share of observations, particularly in both the earliest and most recent cohorts, which have relatively smaller sample sizes to begin with.

Figure D.1 displays the steady-state shares for the full sample and the age-standardized sample. The ordering of the categories remains the same across each cohort, and the steady-state shares differ by no more than 3 or 4 percentage points. The trends remain the same: None's share rises steadily and the Catholic and Mainline shares contract at the same rate for both samples. The most significant discrepancy is the 4 percentage point difference between unadjusted and age-standardized Evangelical steady-state shares in the 1975–84 cohort. Given the aforementioned instabilities due to sample size, we do not over-interpret this difference; the lower standardized Evangelical share, if anything, would widen rather than narrow the gap between Evangelical and None shares in that cohort.

Figure D.2 contains the cohort series for overall mobility and t_{mix}. Overall mobility, in panel (A), rises monotonically in both samples, from roughly 0.24 in the earliest cohort

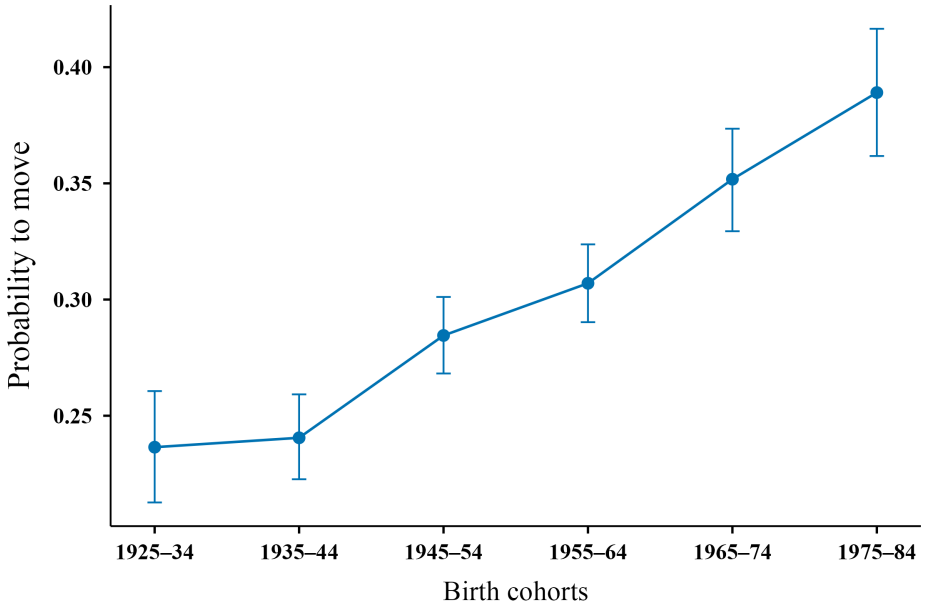
Figure D.1: Baseline and age-adjusted stationary distributions



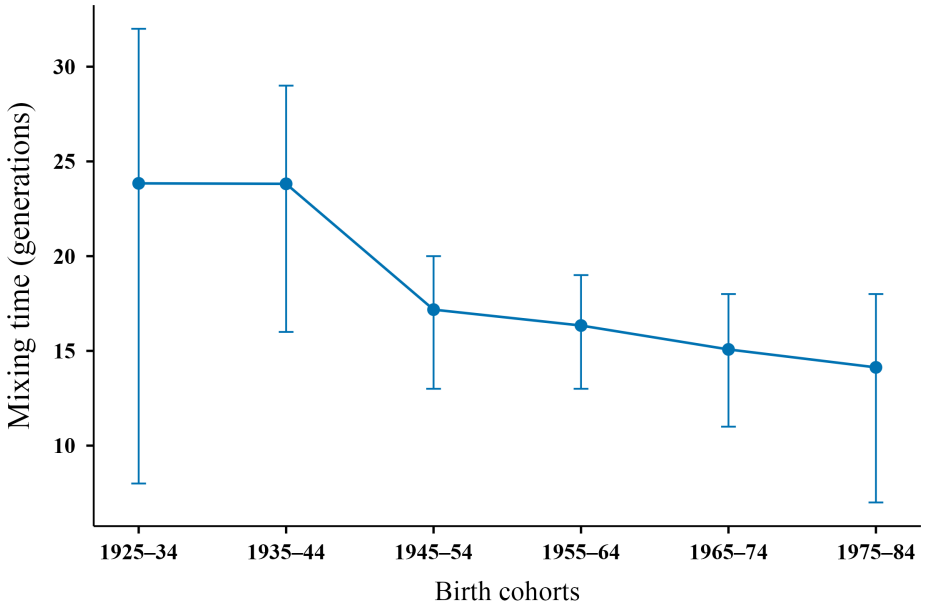
Notes: This figure compares the implied stationary distribution under the baseline sample (π_∞) to the implied stationary distribution under the common age restriction (39–49) by 10-year birth cohort.

to roughly 0.38 in the most recent, with the series basically identical for the 1945–54 and 1955–64 cohorts. If the rise in movement were an artifact of observing later cohorts younger, holding age composition fixed should attenuate the trend. Instead, if anything, it increases it. Mixing times likewise follow the same trend in both samples, declining substantially from around 25 to around 16 across the six cohorts. The age-standardized series is less smooth at both ends and its final point does not continue the decline; given the sample size considerations above, we do not read this as a reversal. The robust feature is the decline across the cohorts for which estimates are well supported.

Figure D.2: Movement and memory summaries with band-restriction



(A) Overall mobility (OM_1)



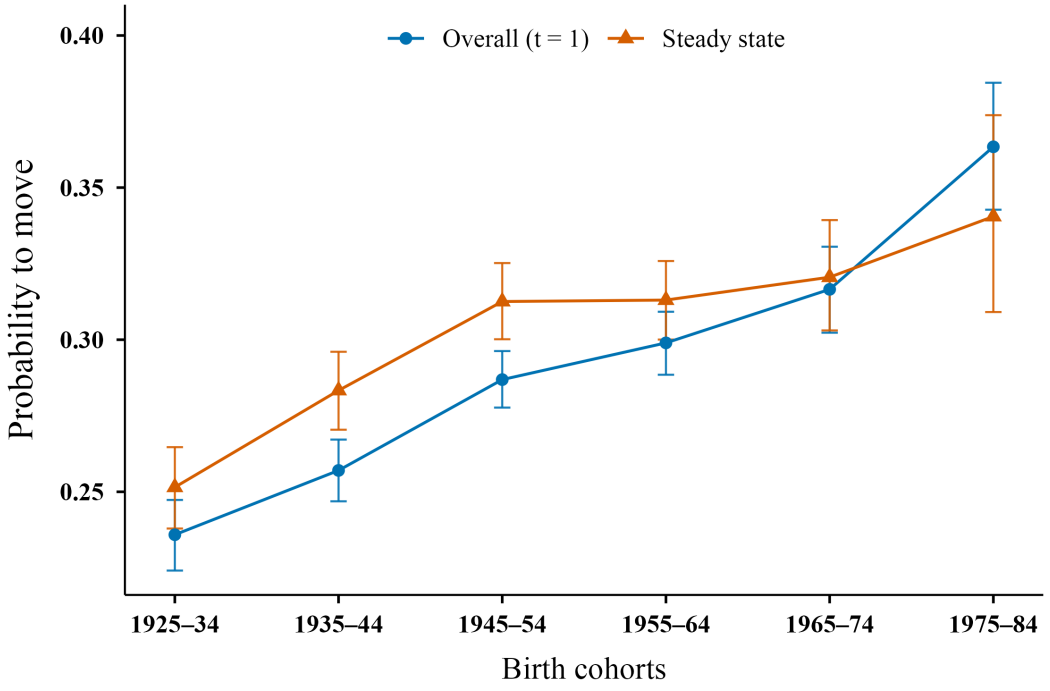
(B) Mixing time (t_{mix})

Notes: Panel A reports overall mobility at $t = 1$; Panel B reports the mixing time. Bias-corrected bootstrap estimates with 95% confidence intervals (2,500 within-survey-year resamples; non-ergodic replicates excluded, cells with >5% rejections not reported). 10-year birth cohorts.

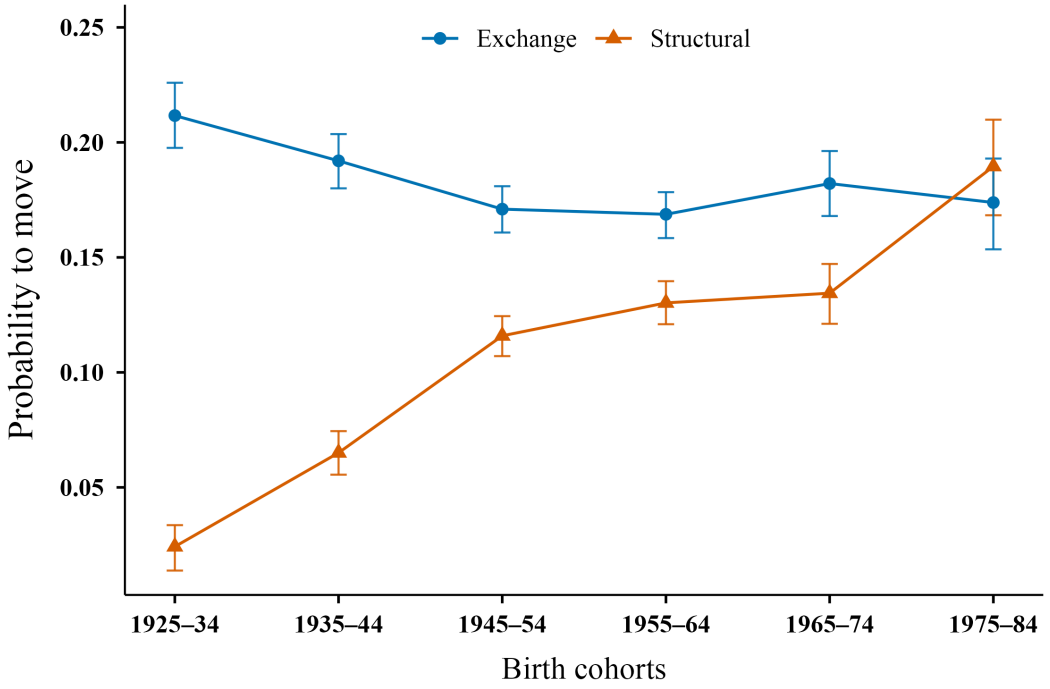
E Excluding Mixed-Mode Waves

The GSS moved away from exclusively in-person interviewing after 2018. In this section, we re-estimate all memory and movement measures after excluding the 2022 and 2024 waves, to check whether the results reported in the main analysis are sensitive to this shift. Every result is qualitatively unchanged: the OM/SSM crossover, the two-stage rise in structural mobility alongside flat exchange mobility, the MTE_i reversal between Catholic/Mainline and None, the AIM and mixing-time trajectories, and the category-specific memory ordering all closely reproduce the patterns reported in the main analysis (Figures E.1–E.4). The youngest cohort loses roughly a third of its observations and its confidence intervals widen accordingly, but the declining persistence and two-pole reorganization documented in the main analysis is not an artifact of the post-2018 shift to mixed-mode interviewing.

Figure E.1: Movement measures, excluding waves 2022 and 2024.



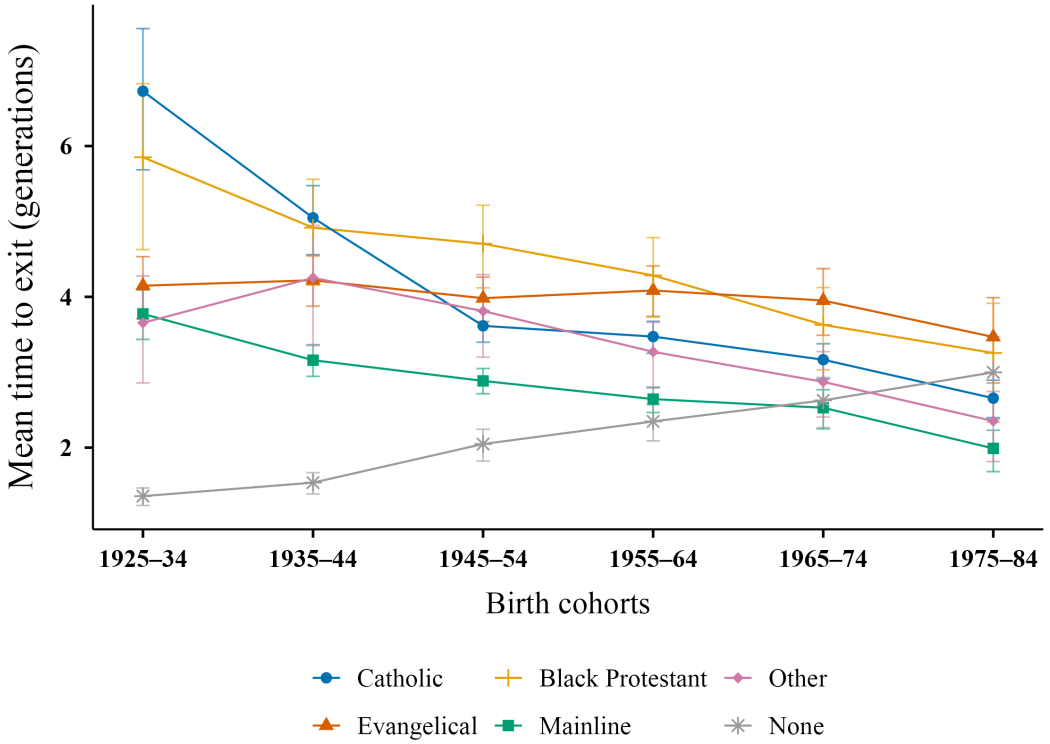
(A) Overall and steady-state mobility



(B) Structural and exchange mobility

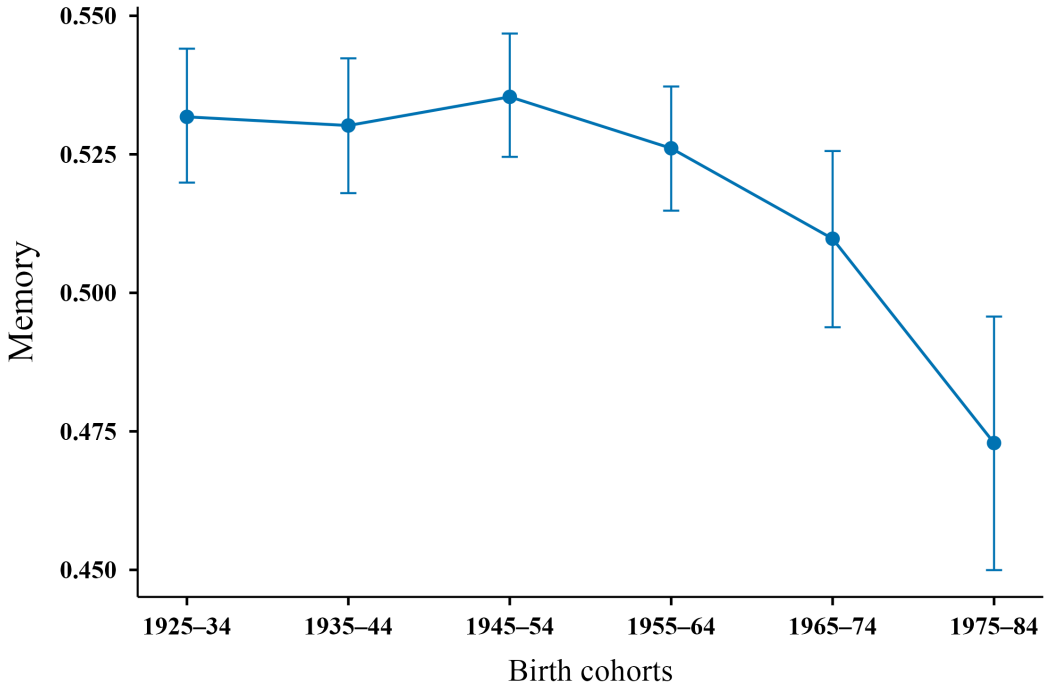
Notes: Panel A reports overall mobility at $t = 1$ (OM_1) and steady-state mobility (SSM); Panel B decomposes OM_1 into its structural (SM_1) and exchange (EM_1) components. Bias-corrected bootstrap estimates with 95% confidence intervals (2,000 within-survey-year resamples).

Figure E.2: Mean time to exit from each origin category (MTE_i), excluding waves 2022 and 2024.

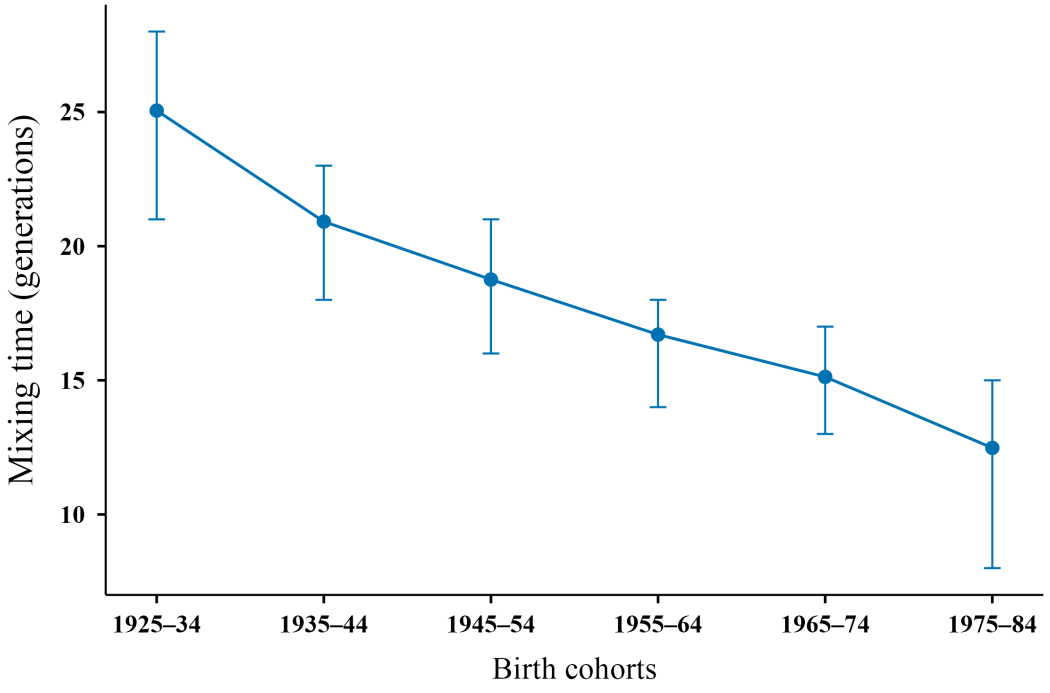


Notes: MTE_i is the expected number of transmission steps until a lineage's religion first differs from its origin state. Bias-corrected bootstrap estimates with 95% confidence intervals (2,000 within-survey-year resamples).

Figure E.3: Memory measures, excluding waves 2022 and 2024.



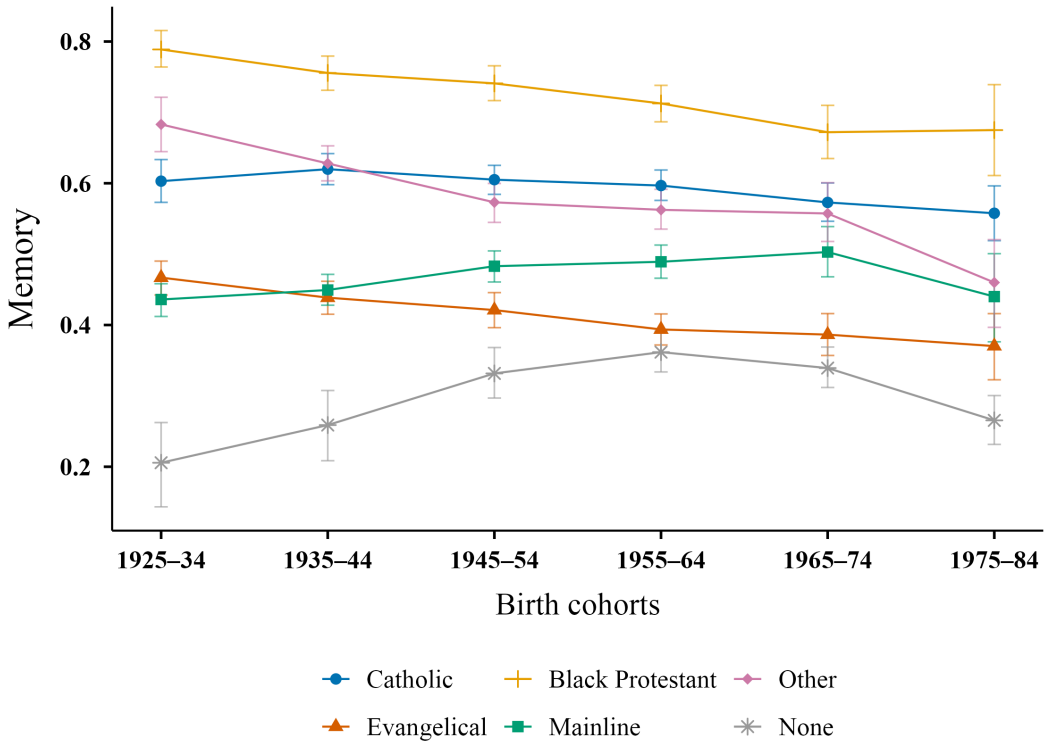
(A) Aggregate intergenerational memory (AIM, $t = 1$)



(B) Mixing time (t_{mix})

Notes: Panel A reports AIM, the π_0-weighted average of the category-specific individual memory measures at $t = 1$. Panel B reports the mixing time t_{mix}, the number of generations required for the transition matrix P^t to numerically converge to its steady state P^∞. Bias-corrected bootstrap estimates with 95% confidence intervals (2,000 within-survey-year resamples).

Figure E.4: Intergenerational memory at $t = 1$ by origin ($IM_{i,1}$), excluding waves 2022 and 2024.

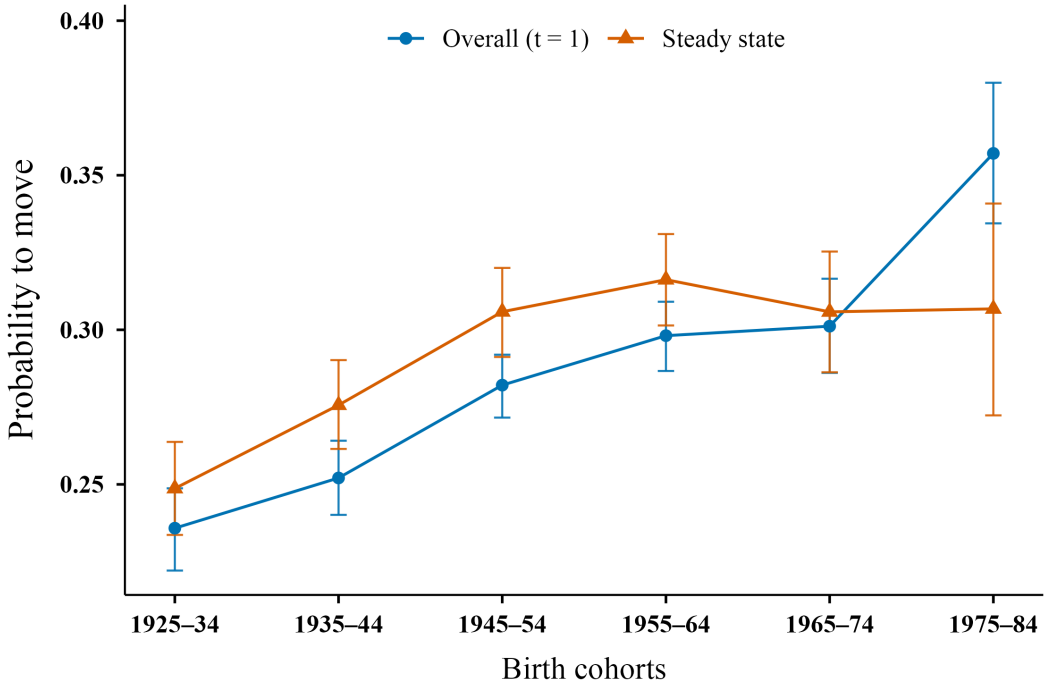


Notes: Bias-corrected bootstrap estimates with 95% confidence intervals (2,000 within-survey-year resamples).

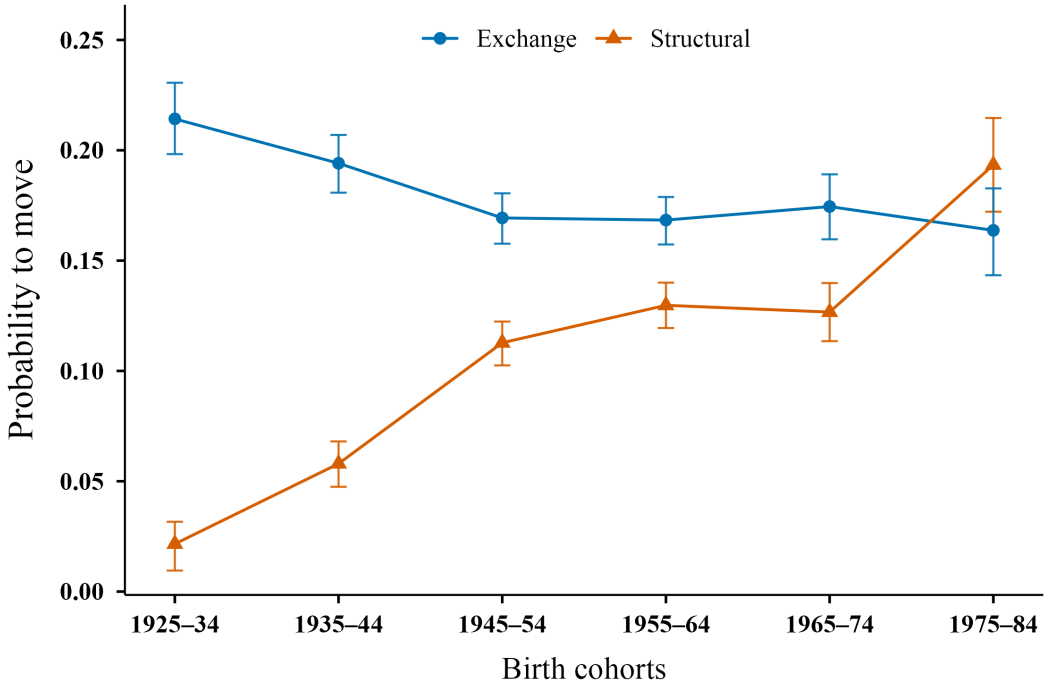
F Weighted Measures

Survey weights adjust for unequal probability of selection across respondents and sampling years. In this section, we re-estimate the main set of memory and movement measures with survey weights applied (wtssps $\times$ oversamp, renormalized within each survey year), under the six-state scheme with the 1982/1987 Black oversamples included, the same sample as the unweighted baseline, so that any difference isolates the effect of weighting. Every result is qualitatively unchanged: the OM/SSM crossover, the two-stage rise in structural mobility alongside flat exchange mobility, the MTE_i reversal between Catholic/Mainline and None, the AIM and mixing-time trajectories, and the category-specific memory ordering all closely reproduce the patterns reported in the main analysis (Figures F.1–F.4). The one visible difference is a widening of confidence intervals in the youngest cohorts, a design-effect cost of weighting rather than a change in the point estimates. The declining persistence and two-pole reorganization documented in the main analysis is therefore not an artifact of unweighted estimation.

Figure F.1: Movement measures, weighted.



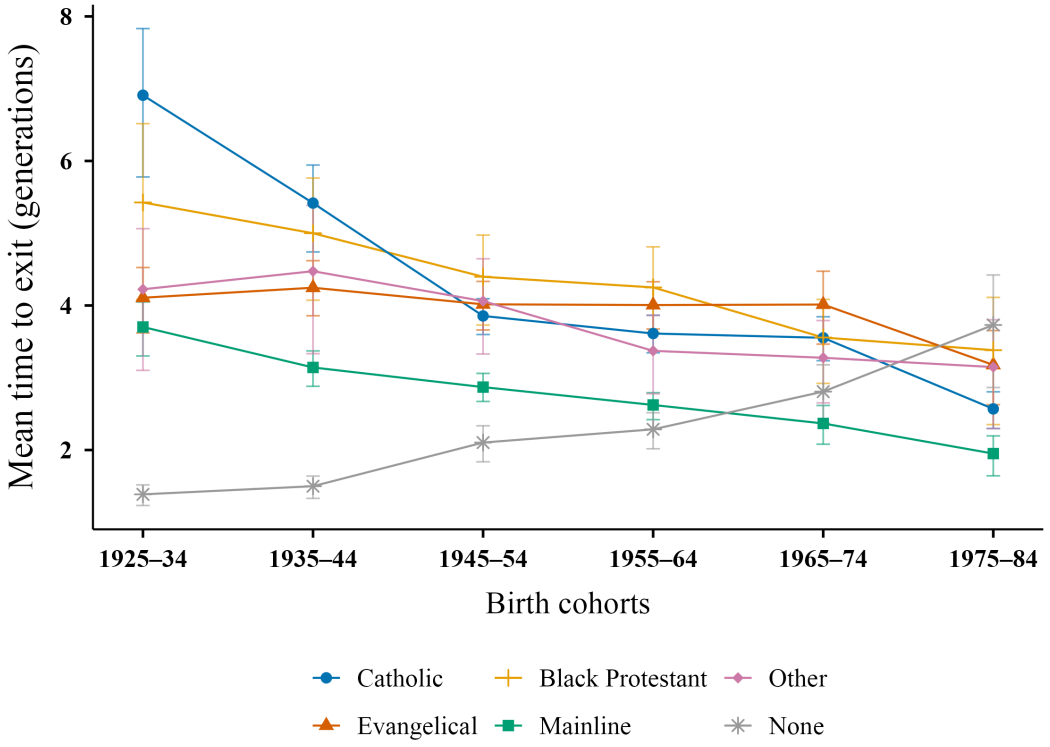
(A) Overall and steady-state mobility



(B) Structural and exchange mobility

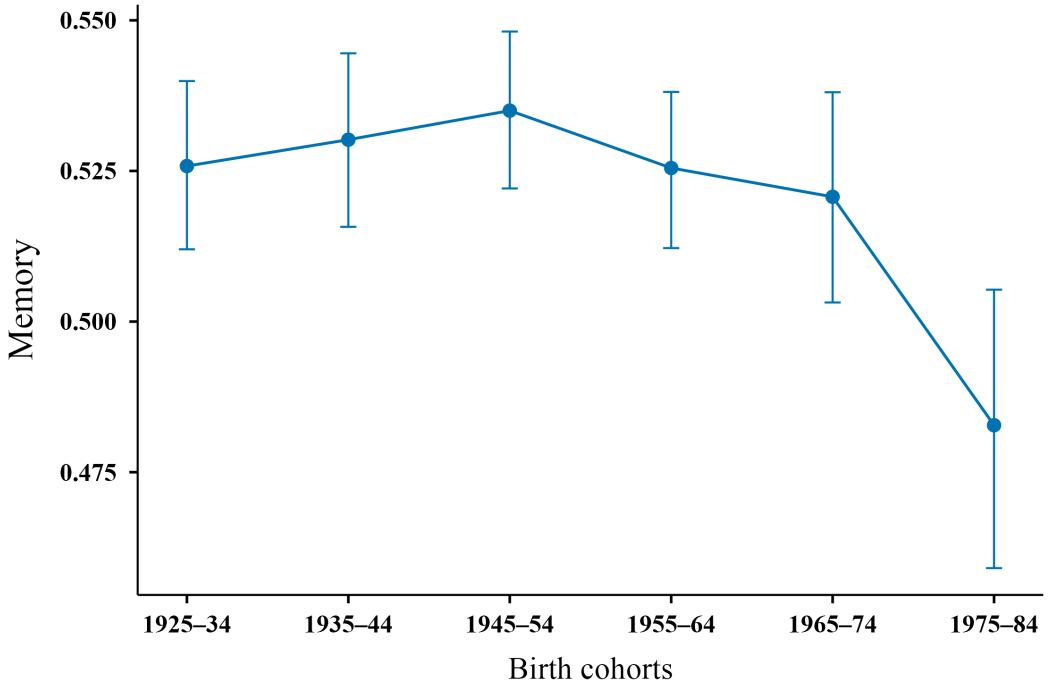
Notes: Panel A reports overall mobility at $t = 1$ (OM_1) and steady-state mobility (SSM); Panel B decomposes OM_1 into its structural (SM_1) and exchange (EM_1) components. Bias-corrected bootstrap estimates with 95% confidence intervals (2,000 within-survey-year resamples).

Figure F.2: Mean time to exit from each origin category (MTE_i), weighted.

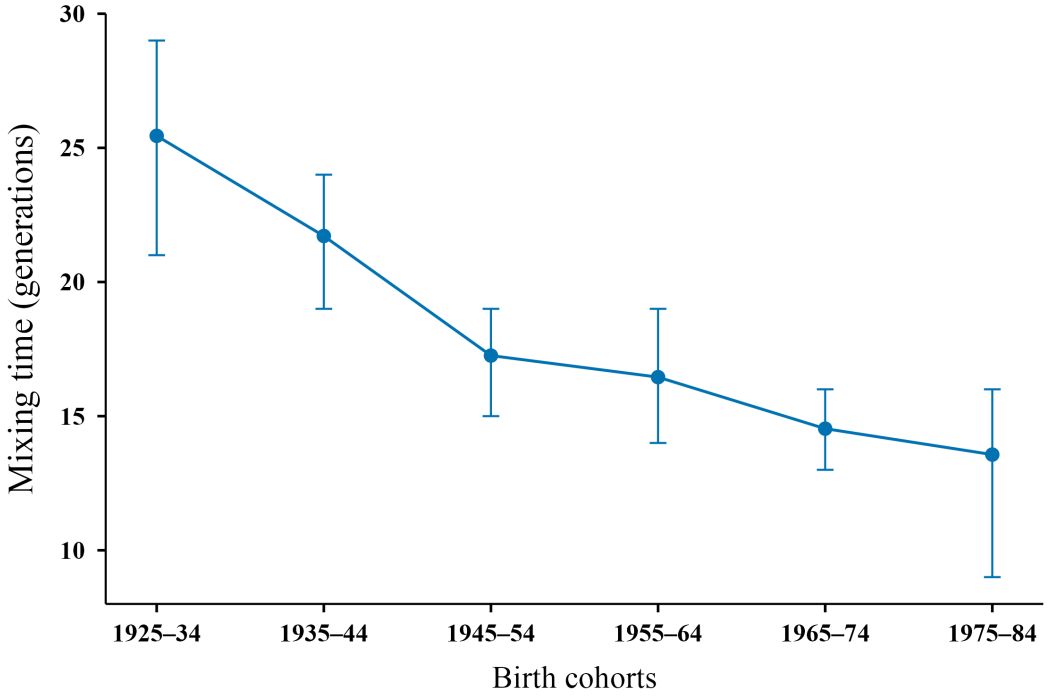


Notes: MTE_i is the expected number of transmission steps until a lineage's religion first differs from its origin state. Bias-corrected bootstrap estimates with 95% confidence intervals (2,000 within-survey-year resamples).

Figure F.3: Memory measures, weighted.



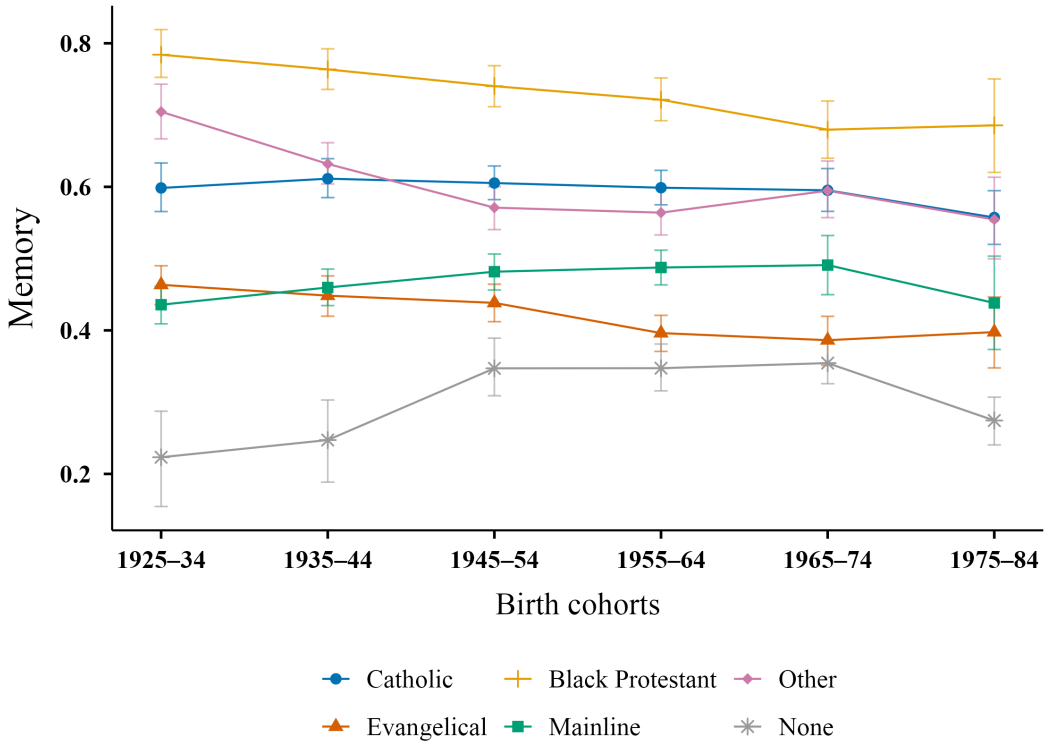
(A) Aggregate intergenerational memory (AIM, $t = 1$)



(B) Mixing time (t_{mix})

Notes: Panel A reports AIM, the π_0-weighted average of the category-specific individual memory measures at $t = 1$. Panel B reports the mixing time t_{mix}, the number of generations required for the transition matrix P^t to numerically converge to its steady state P^∞. Bias-corrected bootstrap estimates with 95% confidence intervals (2,000 within-survey-year resamples).

Figure F.4: Intergenerational memory at $t = 1$ by origin ($IM_{i,1}$), weighted.



Notes: Bias-corrected bootstrap estimates with 95% confidence intervals (2,000 within-survey-year resamples).