

The Rhythm of Aging: *Stability and Drift in the Individual Rate of Senescence*

Silvio C. Patricio*

*Interdisciplinary Center on Population Dynamics
University of Southern Denmark*

Abstract

Human aging is marked by a steady rise in the risk of dying with age—a process demographers call senescence. Over the past century, life expectancy has risen dramatically, but is this because we are aging slower, or simply starting it later? Vaupel hypothesizes that the pace at which individuals age may be constant, with gains in longevity coming from the delayed onset of senescence rather than its slowing down. We test this idea using a new framework that decomposes the pace of senescence into three components: a biological baseline, a long-term trend, and the cumulative impact of period shocks. Applying this to cohort mortality data from 12 countries, we find that once period shocks are accounted for, there is no statistical evidence of a long-term trend, consistent with Vaupel’s hypothesis. Rather than indicating a change in the process that drives senescence, these variations are consistent with echoes of shared historical events. These results suggest that while longevity has shifted, the rhythm of human aging may be conserved.

Keywords: Actuarial senescence, Gompertz law, Rate of aging, Cohort analysis, Period effects

1 Introduction

Aging is the gradual decline in physiological functioning—what we see as graying hair, slower steps, and growing vulnerability to illness and injury. Beneath these visible signs lies senescence, the biological process that drives aging. We focus on *actuarial senescence*—the age-related rise in mortality risk—which,

in most adult populations, shows an almost perfect exponential increase in mortality with age, well described by the Gompertz law (Gompertz, 1825).

In this representation, the Gompertz slope b measures how quickly risk accelerates as deterioration accumulates. Though not a direct biological measure, b is widely used as a proxy for the rate of aging (Olshansky and Carnes, 1997; Finch, 1990). A higher b means mortality rises more steeply with age; a lower b suggests a slower pace of senescence. Our concern is with the *individual rate of aging*, defined as the rate at which an individual’s mortality risk accelerates with age. For brevity, we will refer to this as the *rate of aging* or the *pace of senescence*.

Over the last century, more people have survived to older ages: life expectancy is higher; the age at which most deaths occur has moved to older ages; and later life is, for many, healthier (Oeppen and Vaupel, 2002; Vaupel et al., 2021). We live longer; the question is whether this reflects a *slower* or *later* aging.

Vaupel (2010) framed this as a testable hypothesis: *the rate at which the individual risk of dying increases with age for humans may be a basic biological constant that is very similar and perhaps invariant across individuals and over time*. From this perspective, gains in life expectancy would reflect delayed aging, not a change in the underlying process of senescence. But if the rate of aging is truly changing, it would mean that the biology of aging itself may be evolving—or being changed by the environment, behavior, or historical events (Finch and Crimmins, 2004; Crimmins and Beltrán-Sánchez, 2011; Kirkwood and Austad, 2000).

Empirical tests of this hypothesis have yielded mixed findings: Barbi et al. (2003) found that es-

*silviocabralpatricio@msn.com

estimates of b for Italian cohorts varied significantly depending on the statistical method used, raising the possibility that apparent changes could reflect model sensitivity rather than shifts in senescence. Similarly, Zarulli et al. (2012); Zarulli (2013) analyzed the aftermath of large mortality shocks—such as famine and wartime captivity—and found a flattening of the aging rate at the population level, likely due to selective survival rather than a biological response.

Other studies tested the constancy of b more directly. Salinari and De Santis (2014) rejected the hypothesis that b is constant across countries, sexes, and cohorts, though the differences they observed were modest. In a subsequent paper, Salinari and De Santis (2020) suggest that b might even vary with age, rising before leveling off. However, their parameter b is a cohort-level quantity shaped by selection, not the individual rate of aging itself. Besides this conceptual gap, their model does not separate cohort and age effects, which makes it hard to interpret whether the observed variation reflects fluctuations across cohorts or changes in the aging process.

This may suggest that the variations in b are historically driven. Period events—such as wars, pandemics, and economic crises—strike multiple cohorts at once, just at different ages (Vallin and Meslé, 2004), and their lasting consequences can subtly distort the mortality patterns within each exposed cohort through cumulative shifts (Horiuchi, 2003; Zarulli et al., 2012). As a consequence, when we estimate b cohort by cohort, we may be tracing not a pure signal of the aging process, but the lasting effects of these shared historical events. As these shocks accumulate over time, they can produce variations that mimic a change in the slope of mortality, even if the underlying biological rate is perfectly stable (Zarulli et al., 2012; Salinari and De Santis, 2014; Horiuchi and Wilmoth, 1998).

These kinds of latent effects accumulate gradually, move in one direction for a while, then turn (Nelson and Plosser, 1982). These resemble statistical processes, such as random walks, where movements accumulate over time. If period-driven shocks follow this pattern, they could mimic a changing b , even when the biology of aging holds steady (Yashin et al., 2000). As Alter and Riley (1989) noted, trajec-

ries of frailty and mortality are often shaped not just by individual biology, but by shared historical conditions.

We ask whether cohort-to-cohort variation in the Gompertz slope b reflects a shift in the pace of aging or the effects of period shocks. We decompose b into a baseline and a latent process. We model the latent process as a random walk with drift, the drift being the cohort trend.

A zero drift would suggest that cohort differences in b are consistent with period shocks whose effects persist but cancel on average; a nonzero drift points to a sustained cohort trend. This keeps b interpretable and provides a direct test of whether gains in longevity reflect *slower* or *later* aging.

2 Methodology

To estimate the pace of senescence, we model mortality at advanced ages, where deaths are more likely to reflect intrinsic aging processes than external causes. That is why analysis considers mortality after age 80, where the influence of non-senescent mortality diminishes, and the effects of selection and heterogeneity become defining features of the observed mortality pattern (Thatcher et al., 1998; Horiuchi and Wilmoth, 1998; Rau et al., 2008).

Vaupel’s hypothesis is mathematically grounded in the gamma-Gompertz model, where the exponential increase in the hazard of death due to aging is modulated by unobserved individual frailty (Vaupel et al., 1979; Vaupel and Missov, 2014). In this framework, even when populations become more heterogeneous, the rate of individual aging remains constant.

We use the gamma-Gompertz framework to represent individual aging dynamics within a heterogeneous population. Each individual is assumed to follow a Gompertz force of mortality:

$$\mu(x|Z = z) = z \cdot a e^{bx},$$

where a is the baseline mortality, b is the Gompertz slope, and z is an unobserved frailty term. The model assumes that Z follows a gamma distribution across individuals, with mean 1 and variance γ (e.g., Vaupel

et al., 1979), which leads to the cohort-level hazard:

$$\bar{\mu}(x) = \frac{ae^{bx}}{1 + \gamma \frac{a}{b} (e^{bx} - 1)}. \quad (1)$$

This expression describes how the aggregate force of mortality reflects both the exponential rise in individual risk with age and the dampening influence of selection: frailer individuals tend to die earlier, shifting the distribution of survivors with increasing age. The parameter b , which is assumed to be invariant across individuals, represents the underlying rate at which the individual risk increases with age. In this sense, b captures an individual process that is inferred from population mortality, and we interpret it as the individual rate of aging within the cohort. This assumption is the foundation of the model and allows for a direct test of whether the rate of aging is constant (Vaupel et al., 1979; Steinsaltz and Wachter, 2006; Vaupel and Yashin, 1985).

2.1 Decomposing Variation in b

When we estimate the rate of aging, b , separately for each birth cohort, we often get a series that fluctuates without a clear trend, as illustrated for Danish female cohorts in Figure 1 (left panel). While the series itself appears non-stationary, the series of its log-differences (right panel) shows a stationary process¹. This statistical pattern, where a series becomes stationary after differencing, is the classic signature of a random walk.

This statistical structure mirrors the demographic process itself. Historical events like wars, pandemics, or medical breakthroughs can be seen as "shocks" whose effects persist and accumulate, causing the estimated rate of aging to drift unevenly across cohorts. The average of the differences seen in Figure 1 (right panel) represents the long-term drift of this process. A drift of zero would suggest that historical shocks

¹Stationarity was formally confirmed using a suite of tests at the 5% significance level. The Augmented Dickey-Fuller (ADF) and Phillips-Perron (PP) tests both strongly reject the null hypothesis of a unit root ($p \leq 0.01$), while the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test fails to reject the null hypothesis of stationarity ($p \geq 0.10$). Together, these results provide robust evidence that the series of log-differences is stationary.

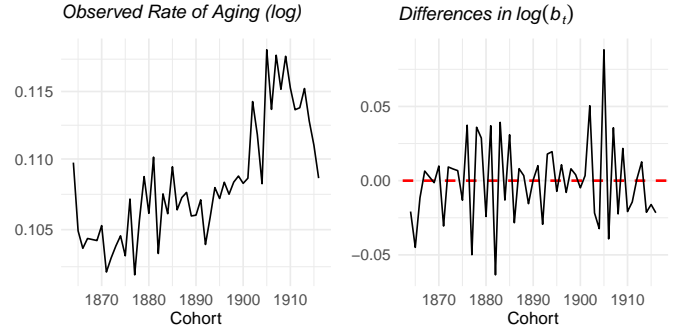


Figure 1: Estimated b across Danish female cohorts and cohort-to-cohort log-difference. The log-difference series is stationary, as confirmed by unit-root and stationarity tests (ADF and PP: $p \leq 0.01$; KPSS: $p \geq 0.10$)

cause short-term fluctuations but no lasting directional change in the pace of aging across cohorts. This insight motivates our choice to model these dynamics explicitly.

To formalize this idea, we model the log rate of aging for cohort t as:

$$\log b_t = \log b + X_t, \quad (2)$$

with

$$X_t = X_{t-1} + \beta + w_t, \quad (3)$$

where:

- b is the baseline rate of aging,
- X_t is a latent random walk,
- β is the drift term—capturing any long-run directional change,
- w_t is white noise with $\mathbb{E}(w_t) = 0$ and variance σ_{rw} .

Parameter interpretation

In this setup, X_t captures the accumulated impact of historical shocks—a stochastic process that can drift over time. The key implication is that $\mathbb{E}[\log b_t] = \log b + \beta t$: any change in b_t over time comes through the drift. This avoids problems with separating the trend from the cumulative effects and keeps the model simple and interpretable.

Our model is designed to test whether period shocks alone can explain the observed variation in the rate of aging. It is built under the assumption of a fixed biological baseline (b), meaning that any persistent directional change across cohorts—whether from cumulative historical effects or a genuine biological trend—is explicitly captured by the drift parameter, β . It measures whether the cumulative impact of period shocks—both negative (wars, pandemics, economic crises) and positive (medical progress, improved nutrition, sanitation)—pushes the observed rate of aging in a consistent direction over the long run.

The interpretation of β is therefore central:

- If we find that $\beta = 0$, it implies that period shocks, while causing cohort-to-cohort fluctuations, have no net directional effect over time. This result would provide strong evidence that there is no long-term trend, either historical or biological, in the rate of aging.
- If we were to find that $\beta \neq 0$, it would suggest a persistent trend in the rate of aging. For instance, if $\beta < 0$, it would suggest that the net effect of positive historical forces like medical advances consistently outweighs negative shocks, driving the observed rate of aging downward. If $\beta > 0$, it would suggest that the impact of negative shocks has been dominant.

This model structure allows us to formally test the parsimonious hypothesis that a single, stable biological rate of aging is sufficient to explain the data once stochastic, non-directional period effects are accounted for. A finding of $\beta \neq 0$ would reject this simple model, indicating a persistent directional force at play. However, it is crucial to note that the model cannot definitively distinguish a purely historical trend from a true change in the underlying biology of senescence. If human biology were evolving toward slower aging, this effect would also be captured in the β term. Therefore, while a non-zero β would confirm a change in the rate of aging, it could not, by itself, disentangle a historical cause from a biological one.

This structure lets us separate a stable biological rate ($\log b$) from the gradual buildup of shared period

effects (X_t), and to test whether that accumulation includes a long-term drift (β). By estimating these components, we can ask whether the changes we see in b_t reflect a true shift in the process that drives aging, or just the long memory of historical shocks.

2.2 Model Specification

To estimate the cohort-specific rate of senescence, we model observed death counts as Poisson-distributed events (Brillinger, 1986):

$$D_{x,t} \sim \text{Poisson}(\lambda(x,t) \cdot E_{x,t}),$$

where $D_{x,t}$ is the number of deaths at age x for cohort t , $E_{x,t}$ is the corresponding exposure, and $\lambda(x,t)$ is the force of mortality, modeled using the gamma-Gompertz hazard in Equation 1.

We use weakly informative priors on the model parameters:

$$\begin{aligned} a_t &\sim \text{half-Normal}(0, 1) \\ \gamma_t &\sim \text{Gamma}(1, 1/2) \\ X_t &\sim \text{Laplace}(X_{t-1} + \beta, \sigma_{\text{rw}}) \\ \log b &\sim \mathcal{N}(0, 10) \\ \beta &\sim \mathcal{N}(0, 10) \\ \sigma_{\text{rw}} &\sim \text{half-Normal}(0, 10) \end{aligned}$$

For the parameters a_t , and σ_{rw} , we use Half-Normal distributions. This follows modern recommendations for such parameters, as these priors offer robust regularization while avoiding the known problems of inverse-gamma distributions (Gelman, 2006). For the heterogeneity parameter γ_t , we instead use a Gamma prior, as recent evidence suggests this specification helps identify whether heterogeneity is supported by the data (Patricio and Missov, 2023).

The choice of a Laplace distribution for the innovations of the random walk $X_t - (X_{t-1} + \beta)$, is motivated by the nature of period shocks: its heavier tails can accommodate rare, large-magnitude events like wars or pandemics, while its sharp peak at zero reflects a belief that most cohort-to-cohort shifts are small.

The magnitude of the shocks is captured by the scale parameter of this distribution, σ_{rw} .

For the main parameters of interest, $\log b$ and β , we use Normal priors centered at zero, reflecting no initial assumption about the direction or magnitude of the baseline rate or its trend.

2.3 Estimation Process

We implemented the model in a Bayesian framework using `Stan`, accessed via its R interface `RStan` (Stan Development Team, 2025). Four independent Markov chains were run with 6,000 iterations each, discarding the first 4,000 as warm-up and retaining the remaining 2,000 per chain, yielding 8,000 posterior samples. Convergence was assessed using the \hat{R} diagnostic, which remained below 1.02 for all parameters (Vehtari et al., 2021). We report maximum a posteriori (MAP) estimates and summarize uncertainty using 95% Highest Posterior Density (HPD) to assess the credible interval.

2.4 Model Validation

To evaluate model adequacy, we conducted posterior predictive checks based on the alignment between observed and replicated death counts. Rather than relying on a single summary statistic, we assessed model fit using the full quantile structure of the data.

For each observed death count y_i , we computed its empirical quantile and compared it to the posterior predictive distribution of replicated counts y_i^{rep} . This quantile-wise comparison is used as a diagnostic tool that evaluates whether the posterior distribution reproduces the shape and spread of the observed data. When the model fits well, the posterior quantiles closely track the empirical quantiles, resulting in a QQ-plot that aligns with the identity line. Deviations from this line, such as systematic curvature, would indicate model misfit—for example, underdispersion or structural bias in the predictions.

Such distributional diagnostics are supported in the Bayesian literature as flexible and interpretable checks of model calibration (Gelman et al., 1996, 2000, 2013). They are valuable when the goal is not hypothesis testing but assessing whether the data

look plausible under the fitted model—an approach that has been advocated in broader statistical modeling contexts (Berkhof et al., 2000; Kruschke, 2021).

2.5 Data

We apply our decomposition model to cohort mortality data from the Human Mortality Database (HMD, 2025). The analysis includes male and female birth cohorts from 12 countries: the Nordic countries (Denmark, Finland, Norway, and Sweden), Western and Southern Europe (France, Italy, the Netherlands, and England & Wales), and selected non-European countries with extensive historical data (Australia, Canada, Japan, and the United States).

To ensure reliable estimation of late-life mortality patterns, we include only cohorts with at least 10 age groups observed above age 80. These countries were chosen for their consistent coverage of cohort mortality, allowing us to trace the rate of aging without the confounding effects of missing data or abrupt changes in registration systems. The cohort series generally begin in the mid- to late-19th century and extend through the end of the 20th century, offering a rich window into long-term demographic change. The specific cohorts included—by country, sex, and range—are listed in Table A1 in the Appendix.

3 Results

We applied the random-walk decomposition model to estimate three components of the Gompertz slope—the biological rate of aging (b), the drift term (β), and the variance of the period effect (σ_{rw}), for both males and females across the 12 countries. Figure 2 presents the posterior estimates, while the full numerical results with 95% credible interval are provided in Tables A2, A3, and A4.

Cross-Country Consistency in b

The estimates of b are relatively consistent across countries. For males, they typically center around 0.102; for females, around 0.107. These differences are small, and in most cases, the intervals overlap, suggesting no statistical difference in the estimated




Figure 2: Posterior estimates of the rate of aging (b , left), the drift (β , center), and the variance of the period effect (σ_{rw} , right) for males (blue) and females (red) across 12 countries. Horizontal bars represent 95% credible interval. The drift estimates are tightly centered around zero, and the period effect is larger in countries affected by major period shocks, such as France, Italy, and Japan.

rates of aging across sexes. In the U.S., Japan, and France, however, the male and female intervals do not overlap, suggesting a statistically higher rate of aging for females. Countries with larger populations and longer cohort series, such as Japan and the U.S., tend to produce narrower intervals, while smaller or data-limited settings, such as Finland, yield wider ones.

Across all countries and both sexes, the estimated drift terms are centered close to zero, with credible intervals that include zero—suggesting no statistically significant trend. This suggests that once period shocks are accounted for, there is no evidence of long-term change in the Gompertz slope. The width of the intervals for β varies across countries, reflecting differences in cohort coverage, population size, and the underlying volatility introduced by period shocks.

What Drives the Variation?

The variance of the latent period effect, σ_{rw} , differs across countries. Higher values imply bigger variability in b_t due to shocks in calendar time, and are most pronounced in countries like France, Italy, and Japan—populations that experienced major disruptions during the World Wars. These events can re-

shape the age-pattern of mortality through competing demographic mechanisms, intensifying processes such as selective survival (Zarulli, 2013; Steinsaltz and Wachter, 2006; Vaupel et al., 1979).

This selective survival mechanism directly impacts the Gompertz slope. A lower b means that for the surviving cohort, the force of mortality accelerates with age. This happens because the population-level death rate is an average over all survivors. When the most vulnerable individuals have been removed by a shock, the remaining ones comprise a more robust population that has a lower average risk of dying at each subsequent age, making the slope of the mortality curve less steep (Vaupel et al., 1979).

This shift in the apparent tempo of aging reflects a well-known demographic pattern: at very old ages, the continual loss of frailer individuals slows the rise in mortality for the survivors (Horiuchi and Wilmoth, 1998). The σ_{rw} parameter captures the overall variability introduced by these long-term, overlapping exposures to period shocks. In contrast, countries such as the Nordics show lower variance estimates, pointing to more stable cohort trajectories over time.

A Stable Rate Across Countries

Across the 12 countries, the results follow a consistent pattern: the rate of aging, b , stays within a narrow range, while the variance of the period effect, σ_{rw} , varies in ways that reflect historical and demographic differences. In every case, the drift term, β , remains around zero, providing no statistical evidence of a sustained directional shift in the rate of aging. This supports the most parsimonious explanation: a stable rate of aging shaped by the accumulation of non-trending period shocks.

The sex difference in b is consistent with the overall picture of stability. Males typically experience higher mortality at younger and middle ages, mainly due to biological and behavioral factors (Beltrán-Sánchez et al., 2015). This more intense early-life selection can act as a “winnowing” process, removing frailer individuals and leaving a more robust and homogeneous group of male survivors at high ages. As a result, the surviving group often has a lower b (Vaupel et al., 1979). By contrast, females experience lower

mortality throughout most of adulthood, may carry a wider range of underlying frailty into old age, resulting in a steeper observed rate of aging as selection continues to operate on a more diverse population.

This sex difference does not contradict our main finding of stability across cohorts. For males and females, the rates of aging show no statistical evidence of a trend, and the variation in b_t is small compared with the volatility introduced by the cumulative period effect. Together, these results suggest that variation in b_t reflects exposure to events in calendar time rather than changes in the aging process.

3.1 Model Diagnostics and Supplementary Estimates

Figure 3 shows the posterior predictive QQ-plots for males and females across all 12 countries. In both cases, the predicted quantiles align closely with the observed ones, falling along the identity line with only minor variation. This indicates that the model replicates the shape and spread of the empirical death count distribution well—not just on average, but across its full range.




Figure 3: Posterior predictive QQ-plots comparing observed and simulated quantiles of cohort death counts across 12 countries, separately for females (top) and males (bottom). Colored lines show posterior means by country; the black line represents the 45-degree identity line. The close alignment indicates that the model accurately captures the distribution of the observed data across its full range.

These results support the adequacy of the random-walk specification and the assumption that cohort-

level estimates of the Gompertz slope can be explained by baseline, drift, and period components. Country-specific decompositions of these components are provided in Appendix A2.

4 Discussion

This study set out to revisit Vaupel’s hypothesis that the rate of aging, b , is constant. We test whether fluctuations in b reflect changes in the rate of aging or echo shared historical shocks. Our findings are consistent with the latter. Once fluctuations driven by period shocks are accounted for, the estimates of b show no clear trend. This suggests that what appears to be a change in the pace of aging may not reflect a shift in this underlying process itself, but rather the cumulative impact of the historical environment experienced by different cohorts.

At the same time, the results do not point to a universal constant across all populations. The estimated values of b varies slightly between countries and sexes. These differences are small—close enough that they could reflect model specification or data limitations—but they are inconsistent with the idea of a single invariant value, as suggested by Vaupel’s hypothesis. What the evidence supports, then, is not strict universality but a narrow range of stability, in which modest variation coexists with a broadly conserved demographic rhythm of aging.

Support from Past Studies

Vaupel’s hypothesis gains new support from this decomposition. The stable component of b that our model identifies aligns with his original claim (Vaupel, 2010), while the observed variation is consistent with the accumulation of historical shocks. Where Barbi et al. (2003) saw model sensitivity, we interpret this sensitivity as a structured drift that reflects these historical shocks. And while Zarulli et al. (2012) and Salinari and De Santis (2014) interpreted cohort fluctuations as potential biological change, our results suggest those shifts may reflect shared exposures in calendar time. As Alter and Riley (1989) noted, mortality is embedded in historical context; what appears biological may instead reflect social history.

The idea that the rate of senescence is a stable and underlying biological parameter is supported by comparative research. A comprehensive study by Colchero et al. (2021), testing the “invariant rate of aging” hypothesis, found that the Siler model’s rate of aging parameter remains stable in primates, a pattern they link to species with slow life histories. This distinction between a stable rate of senescence and more variable, environmentally driven mortality is also observed in other species; for example, a study on Hazel Dormice found that while overall aging patterns responded to environmental conditions, the rate of senescence itself did not vary (Berg et al., 2025). On a broader scale, a survey across 101 wild mammal species found high variability in aging rates, with local conditions driving overall mortality patterns, yet the underlying pace of senescence remained stable (Lemaitre et al., 2020).

This interpretation is supported by recent theoretical models that explain why the Gompertz slope may be stable. Work by Nielsen et al. (2024), for example, shows that an exponential increase in mortality is a natural, emergent property of any complex organism with many interconnected parts (like the circulatory, nervous, and immune systems). Similarly, Flietner et al. (2024) uses network theory to show that the same mortality patterns can arise from random failures in networks structured like those found in biology.

Work by Hansen and Strulik (2025) reinforces this distinction. In their model, mortality rises as the compound result of two processes: the exponential accumulation of health deficits, and a power-law relationship between those deficits and death. The pace of aging remains fixed; it is the inputs that shift. In this sense, our findings parallel theirs: the onset of senescence can shift, even if its tempo remains constant.

The hallmarks of aging framework (López-Otín et al., 2013, 2023) highlights core processes—like telomere attrition and genomic instability—that operate consistently across species. Our finding of a stable demographic tempo of aging is consistent with the idea that these core biological mechanisms may also be highly conserved over time, even as overall mortality patterns shift. At the same time, con-

stancy holds within each country–sex subpopulation, but not uniformly across all subpopulations. There is little evidence that the Gompertz slope is identical everywhere; still, the observed variation is modest, and far smaller than would be expected if the underlying rate of aging were truly changing.

Onset vs. Tempo of Senescence

If the pace of aging does *not* change but its onset *does*, the timing of senescence may be influenced by early-life conditions, cumulative exposures, or broader social trajectories. Like Finch and Crimmins (2004) and Crimmins and Beltrán-Sánchez (2011), our findings suggest that aging’s timing may respond to environmental forces, even if its tempo does not, and the widening gaps in healthy life expectancy underscore the unequal ways in which this timing plays out (Crimmins and Saito, 2001).

Evidence from Patricio and Missov (2024) supports this view: the demographic onset of senescence—the age when mortality driven by senescence begins to dominate over extrinsic risks—has shifted to later ages. Although not a direct biological measure, this onset identifies the point at which age-related deterioration becomes the main force shaping population mortality.

The difference between when aging begins and how fast it proceeds may change how we interpret gains in lifespan, whether they extend healthspan, and whether morbidity is being compressed or redistributed. A stable rate of aging means the pace is remarkably consistent, but the clock starts at different moments.

Implications for Longevity Forecasting

Standard methods for forecasting mortality, such as the Lee-Carter model, implicitly allow the demographic rate of aging—the slope of the log-mortality curve—to drift over time. This occurs because the model’s time-varying index (k_t) interacts with the age-specific component (b_x), causing the entire age profile of mortality to “rotate” and change its steepness from year to year. Our results, however, suggest that this flexibility may be unnecessary. The stability of b observed in our study suggests that future

gains in life expectancy may continue to come from postponing the onset of senescence (lowering baseline mortality), rather than from slowing the rate of aging itself. Therefore, incorporating our findings into forecast models—either by constraining the Gompertz slope to be constant over the long term or by modeling its short-term fluctuations as a stochastic process (such as a random walk)—could potentially produce more plausible and robust long-term projections of human longevity.

Challenges and Directions from Geroscience

Still, our findings are not definitive, and new biological insights continue to challenge the notion of an immutable aging rate, highlighting the contrast between a stable demographic slope and emerging evidence of a malleable aging pace.

From a demographic perspective, a way to connect these views is to see the Gompertz slope b as a robust, system-level property of aging that is relatively constant. In this view, individual-level biological interventions might primarily influence the timing or heterogeneity of frailty accumulation, shifting the intercept of the Gompertz curve rather than its slope.

Recent studies using new biomarkers claim to quantify an individual’s pace of aging and show its susceptibility to lifestyle interventions (Belsky et al., 2020). While these tools measure a biological pace of aging, our results raise the question of whether they map directly onto the rate of mortality acceleration or if they primarily track the accumulation of damage that influences the onset of senescence.

Furthermore, the burgeoning field of senolytics, which involves selectively removing senescent cells (Chaib et al., 2022), offers a more direct biological intervention. A successful therapy that removes accumulated damage might be expected to lower the force of mortality, shifting the mortality schedule to later ages and thereby delaying the onset of senescence. An intervention that fundamentally alters core cellular repair mechanisms, however, would be required to change the Gompertz slope—a parameter our analysis suggests is highly conserved.

These findings invite a re-evaluation of the debate about the primary goals of anti-aging research.

Instead of focusing on altering a rate that appears highly conserved, perhaps the most fruitful interventions are those that postpone the onset of senescence or reduce the frailty individuals carry into old age.

Improving the Model

We assume that the variance of the period effect is constant across cohorts, which helps keep the model simple and interpretable. The tradeoff is between capturing more of the variation in the data and keeping a parsimonious model that robustly identifies the long-term trend in the rate of aging. The current approach prioritizes the stability and interpretability of the core Gompertz slope, recognizing that excessive flexibility could blur rather than clarify the underlying trend in mortality acceleration.

A model extension with time-varying volatility (e.g., a GARCH structure) could test whether the magnitude of historical shocks has changed over time, but our results show that even a simple random walk is sufficient to absorb the variation previously interpreted as a trend. A different line of future research could use causal inference frameworks to disentangle the impact of specific period events on cohort mortality trajectories, but would face the well-known identification problem in age-period-cohort models, which complicates the separation of these distinct effects.

Extending to Subgroups

Our results suggest that disparities in healthy longevity between socioeconomic groups may, at least in part, reflect different exposures to period effects rather than differences in intrinsic aging rates.

This perspective offers a possible explanation for the stark inequalities documented in recent studies. In Denmark, Strozza et al. (2025) report a widening gap in survival to retirement age between the highest and lowest socioeconomic groups, driven by stagnating mortality improvements among the most disadvantaged. In the U.S., Bergeron-Boucher et al. (2024) show that the intersection of sex, race, education, and marital status can produce lifespan differences of up to 18 years, with disadvantage accumulating across multiple social dimensions.

Our findings suggest a potential mechanism for this: while entire birth cohorts experience the same historical shocks—such as wars, pandemics, or economic recessions—the impact of these shocks is not uniform. Instead, it is heavily stratified by social position (Phelan et al., 2010). For example, during a war, men are more likely to face combat, while during a pandemic, frontline workers from disadvantaged backgrounds face greater exposure to infection (Hooper et al., 2020). This differential vulnerability to period effects, compounded by accumulated disadvantage, may explain a significant portion of health disparities (Elo, 2009), suggesting they stem less from intrinsic differences in the rate of aging and more from the unequal consequences of shared history.

Understanding aging as the interplay between a fixed rhythm and a variable historical context provides a framework for studying health inequality. This perspective also helps contextualize findings such as those by Graf et al. (2024), which show that higher education is linked to slower epigenetic aging. Whether this reflects a direct effect on the aging process or fewer negative period shocks shaping these measures remains an open question.

Prior work suggests that adverse exposures can accelerate epigenetic aging (Kim et al., 2023; Shen et al., 2023), but it is still unclear whether these differences are biological or the result of unequal exposure to shocks. Testing these dynamics could be a next step. A stable aging rate across socioeconomic groups would support the hypothesis of a shared underlying process, while finding a different rate could suggest that disparities are driven by unequal exposures and cumulative disadvantage.

Clarifying What Changes—and What Does Not

To our knowledge, this is the first study to decompose the variation in b into a latent period effect. In doing so, we provide a framework that aims to separate the stable component of aging from the noise of period events, using a model that is both parsimonious and grounded in demographic theory.

5 Conclusion

By focusing on late-life mortality and accounting for the accumulation of shared period effects, our analysis finds no statistically significant long-term trend in the Gompertz slope, b , across cohorts. This suggests that the underlying tempo of aging is stable, with the observed variation being consistent with the cumulative result of historical shocks.

While our findings are consistent with Vaupel’s hypothesis, they also highlight the importance of connecting it with emerging evidence of a malleable aging pace and the persistent disparities in longevity across socioeconomic groups. Our work suggests that the rhythm of aging may be highly conserved, but brings into focus the crucial questions of when, how, and for whom the clock starts.

In the end, our framework suggests that what often appears as a shift in aging may not be a change in the intrinsic rate of senescence, but rather the echo of history written in mortality patterns. By attempting to distinguish this underlying trend from historical noise, we can better understand what has—and has not—changed in the story of human longevity.

Acknowledgments

I thank James E. Oeppen, Marie-Pier Bergeron-Boucher, and Trifon I. Missov for their valuable comments on the manuscript. This research was supported by the AXA Research Fund through the “AXA Chair in Longevity Research”.

References

- Alter, G. and Riley, J. C. (1989). Frailty, sickness, and death: models of morbidity and mortality in historical populations. *Population Studies*, 43(1):25–45.
- Barbi, E. et al. (2003). Assessing the rate of ageing of the human population. *Max Planck Institute for Demographic Research Working Paper*, 8.
- Belsky, D. W., Caspi, A., Arseneault, L., Baccarelli, A., Corcoran, D. L., Gao, X., Hannon, E., Harrington, H. L., Rasmussen, L. J., Houts, R., et al. (2020). Quantification of the pace of biological aging in humans through a blood test, the dunedinpoam dna methylation algorithm. *elife*, 9:e54870.

- Beltrán-Sánchez, H., Finch, C. E., and Crimmins, E. M. (2015). Twentieth century surge of excess adult male mortality. *Proceedings of the National Academy of Sciences*, 112(29):8993–8998.
- Berg, T. B., Colchero, F., Jones, O. R., Sanderhoff, L., and Juškaitis, R. (2025). Variation in mortality and ageing rate in a fast-paced species: Insights from 24 years of hazel dormouse (*muscardinus avellanarius*) data. *Ecology and Evolution*, 15(6):e71440.
- Bergeron-Boucher, M.-P., Callaway, J., Strozza, C., and Oeppen, J. (2024). Inequalities in lifespan and mortality risk in the us, 2015–2019: a cross-sectional analysis of subpopulations by social determinants of health. *BMJ open*, 14(6):e079534.
- Berkhof, J., Van Mechelen, I., and Hoijsink, H. (2000). Posterior predictive checks: Principles and discussion. *Computational Statistics*, 15(3):337–354.
- Brillinger, D. R. (1986). A biometrics invited paper with discussion: the natural variability of vital rates and associated statistics. *Biometrics*, pages 693–734.
- Chaib, S., Tchkonja, T., and Kirkland, J. L. (2022). Cellular senescence and senolytics: the path to the clinic. *Nature medicine*, 28(8):1556–1568.
- Colchero, F., Aburto, J. M., Archie, E. A., Boesch, C., Breuer, T., Campos, F. A., Collins, A., Conde, D. A., Cords, M., Crockford, C., et al. (2021). The long lives of primates and the ‘invariant rate of ageing’ hypothesis. *Nature communications*, 12(1):3666.
- Crimmins, E. M. and Beltrán-Sánchez, H. (2011). Mortality and morbidity trends: is there compression of morbidity? *Journals of Gerontology Series B: Psychological Sciences and Social Sciences*, 66(1):75–86.
- Crimmins, E. M. and Saito, Y. (2001). Trends in healthy life expectancy in the united states, 1970–1990: gender, racial, and educational differences. *Social science & medicine*, 52(11):1629–1641.
- Elo, I. T. (2009). Social class differentials in health and mortality: Patterns and explanations in comparative perspective. *Annual review of sociology*, 35(1):553–572.
- Finch, C. E. (1990). *Longevity, senescence, and the genome*. University of Chicago Press.
- Finch, C. E. and Crimmins, E. M. (2004). Inflammatory exposure and historical changes in human life-spans. *Science*, 305(5691):1736–1739.
- Flietner, V., Heidergott, B., Hollander, F. d., Lindner, I., Parvaneh, A., and Strulik, H. (2024). A unifying theory of aging and mortality. *arXiv preprint arXiv:2412.12815*.
- Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models (comment on article by browne and draper). *Bayesian Analysis*.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2013). *Bayesian data analysis*. CRC press.
- Gelman, A., Goegebeur, Y., Tuerlinckx, F., and Van Mechelen, I. (2000). Diagnostic checks for discrete data regression models using posterior predictive simulations. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 49(2):247–268.
- Gelman, A., Meng, X.-L., and Stern, H. (1996). Posterior predictive assessment of model fitness via realized discrepancies. *Statistica sinica*, pages 733–760.
- Gompertz, B. (1825). On the nature of the function expressive of the law of human mortality, and on a new mode of determining the value of life contingencies: in a letter to francis baily, esq. frs &c. *Philosophical transactions of the Royal Society of London*, 115:513–585.
- Graf, G. H., Aiello, A. E., Caspi, A., Kothari, M., Liu, H., Moffitt, T. E., Muennig, P. A., Ryan, C. P., Sugden, K., and Belsky, D. W. (2024). Educational mobility, pace of aging, and lifespan among participants in the framingham heart study. *JAMA network open*, 7(3):e240655–e240655.
- Hansen, C. W. and Strulik, H. (2025). How do we age? a decomposition of gompertz law. *Journal of Health Economics*, page 102988.
- HMD (2025). The human mortality database. <http://www.mortality.org/>.
- Hooper, M. W., Nápoles, A. M., and Pérez-Stable, E. J. (2020). Covid-19 and racial/ethnic disparities. *Jama*, 323(24):2466–2467.
- Horiuchi, S. (2003). Interspecies differences in the life span distribution: Humans versus invertebrates. *Population and Development Review*, 29:127–151.
- Horiuchi, S. and Wilmoth, J. R. (1998). Deceleration in the age pattern of mortality at olderages. *Demography*, 35(4):391–412.
- Kim, K., Yaffe, K., Rehkopf, D. H., Zheng, Y., Nannini, D. R., Perak, A. M., Nagata, J. M., Miller, G. E., Zhang, K., Lloyd-Jones, D. M., et al. (2023). Association of adverse childhood experiences with accelerated epigenetic aging in midlife. *JAMA network open*, 6(6):e2317987–e2317987.
- Kirkwood, T. B. and Austad, S. N. (2000). Why do we age? *Nature*, 408(6809):233–238.
- Kruschke, J. K. (2021). Bayesian analysis reporting guidelines. *Nature human behaviour*, 5(10):1282–1291.

- Lemaître, J.-F., Ronget, V., Tidière, M., Allainé, D., Berger, V., Cohas, A., Colchero, F., Conde, D. A., Garratt, M., Liker, A., et al. (2020). Sex differences in adult lifespan and aging rates of mortality across wild mammals. *Proceedings of the National Academy of Sciences*, 117(15):8546–8553.
- López-Otín, C., Blasco, M. A., Partridge, L., Serrano, M., and Kroemer, G. (2013). The hallmarks of aging. *Cell*, 153(6):1194–1217.
- López-Otín, C., Blasco, M. A., Partridge, L., Serrano, M., and Kroemer, G. (2023). Hallmarks of aging: An expanding universe. *Cell*, 186(2):243–278.
- Nelson, C. R. and Plosser, C. R. (1982). Trends and random walks in macroeconomic time series: some evidence and implications. *Journal of monetary economics*, 10(2):139–162.
- Nielsen, P. Y., Jensen, M. K., Mitarai, N., and Bhatt, S. (2024). The gompertz law emerges naturally from the interdependencies between sub-components in complex organisms. *Scientific Reports*, 14(1):1196.
- Oeppen, J. and Vaupel, J. W. (2002). Broken limits to life expectancy.
- Olshansky, S. J. and Carnes, B. A. (1997). Ever since gompertz. *Demography*, 34(1):1–15.
- Patricio, S. C. and Missov, T. I. (2023). Using a penalized likelihood to detect mortality deceleration. *Plos one*, 18(11):e0294428.
- Patricio, S. C. and Missov, T. I. (2024). Makeham mortality models as mixtures. *Demographic Research*, 51:595–624.
- Phelan, J. C., Link, B. G., and Tehranifar, P. (2010). Social conditions as fundamental causes of health inequalities: theory, evidence, and policy implications. *Journal of health and social behavior*, 51(1_suppl):S28–S40.
- Rau, R., Soroko, E., Jasilionis, D., and Vaupel, J. W. (2008). Continued reductions in mortality at advanced ages. *Population and Development Review*, 34(4):747–768.
- Salinari, G. and De Santis, G. (2014). Comparing the rate of individual senescence across time and space. *Population*, 69(2):165–190.
- Salinari, G. and De Santis, G. (2020). One or more rates of ageing? the extended gamma-gompertz model (egg). *Statistical Methods & Applications*, 29(2):211–236.
- Shen, B., Mode, N. A., Hooten, N. N., Pacheco, N. L., Ezike, N., Zonderman, A. B., and Evans, M. K. (2023). Association of race and poverty status with dna methylation-based age. *JAMA Network Open*, 6(4):e236340–e236340.
- Stan Development Team (2025). *RStan: the R interface to Stan*. R package version 2.32.7.
- Steinsaltz, D. R. and Wachter, K. W. (2006). Understanding mortality rate deceleration and heterogeneity. *Mathematical Population Studies*, 13(1):19–37.
- Strozza, C., Vigezzi, S., Callaway, J., Aleksandrov, A., and Kashnitsky, I. (2025). Socioeconomic inequalities in survival to retirement age in denmark: a register-based analysis. *Genus*, 81(1):21.
- Thatcher, A. R., Kannisto, V., and Vaupel, J. W. (1998). *The Force of Mortality at Ages 80 to 120*, volume 5 of *Monographs on Population Aging*. Odense University Press.
- Vallin, J. and Meslé, F. (2004). Convergences and divergences in mortality: a new approach of health transition. *Demographic research*, 2:11–44.
- Vaupel, J. W. (2010). Biodemography of human ageing. *Nature*, 464(7288):536–542.
- Vaupel, J. W., Manton, K. G., and Stallard, E. (1979). The impact of heterogeneity in individual frailty on the dynamics of mortality. *Demography*, 16(3):439–454.
- Vaupel, J. W. and Missov, T. I. (2014). Unobserved population heterogeneity: A review of formal relationships. *Demographic Research*, 31:659–686.
- Vaupel, J. W., Villavicencio, F., and Bergeron-Boucher, M.-P. (2021). Demographic perspectives on the rise of longevity. *Proceedings of the National Academy of Sciences*, 118(9):e2019536118.
- Vaupel, J. W. and Yashin, A. I. (1985). Heterogeneity’s ruses: some surprising effects of selection on population dynamics. *The American Statistician*, 39(3):176–185.
- Vehtari, A., Gelman, A., Simpson, D., Carpenter, B., and Bürkner, P.-C. (2021). Rank-normalization, folding, and localization: An improved \hat{R} for assessing convergence of mcmc (with discussion). *Bayesian analysis*, 16(2):667–718.
- Yashin, A. I., Iachine, I. A., and Begun, A. S. (2000). Mortality modeling: A review. *Mathematical Population Studies*, 8(4):305–332.
- Zarulli, V. (2013). The effect of mortality shocks on the age-pattern of adult mortality. *Population*, 68(2):265–291.
- Zarulli, V. et al. (2012). Mortality shocks and the human rate of aging. *Max Planck Institute for Demographic Research—MPIDR Working Paper*.

Appendix

This appendix provides supplementary material to support the analyses presented in the main text. Section A1 lists the cohorts considered in the analysis and reports country-specific parameter estimates with credible intervals, and Section A2 contains the model diagnostics.

A1 Supplementary tables

This section reports the cohorts included in the analysis and the posterior estimates of the main model parameters, shown with 95% credible intervals by country and sex.

Table A1: Cohorts considered in the analysis. Countries, sexes, and ranges of cohorts included.

Country	Sex	First cohort	Last cohort	Length
Australia	Male	1845	1931	86
	Female	1845	1931	86
Canada	Male	1840	1932	92
	Female	1840	1932	92
Denmark	Male	1764	1934	170
	Female	1764	1934	170
Finland	Male	1807	1933	126
	Female	1807	1933	126
France	Male	1737	1932	195
	Female	1737	1932	195
Italy	Male	1794	1932	138
	Female	1794	1932	138
Japan	Male	1868	1933	65
	Female	1868	1933	65
Netherlands	Male	1777	1932	155
	Female	1777	1932	155
Norway	Male	1770	1933	163
	Female	1770	1933	163
Sweden	Male	1676	1933	257
	Female	1676	1933	257
United States	Male	1852	1933	81
	Female	1852	1933	81
England & Wales	Male	1764	1932	168
	Female	1764	1932	168

Table A2: Estimate and 95% credible interval for the individual rate of aging (b)

Country	Sex	Estimate	Lower C.I.	Upper C.I.
Australia	Male	0.1080	0.1005	0.1133
	Female	0.1095	0.1062	0.1135
Canada	Male	0.1026	0.0972	0.1073
	Female	0.1059	0.1006	0.1099
Denmark	Male	0.1095	0.1031	0.1142
	Female	0.1064	0.1037	0.1087
Finland	Male	0.0963	0.0902	0.1022
	Female	0.1072	0.1022	0.1138
France	Male	0.0999	0.0960	0.1039
	Female	0.1107	0.1066	0.1149
Italy	Male	0.1000	0.0961	0.1041
	Female	0.1105	0.1078	0.1132
Japan	Male	0.0931	0.0857	0.0979
	Female	0.1043	0.0984	0.1069
Netherlands	Male	0.1006	0.0971	0.1047
	Female	0.1010	0.0970	0.1053
Norway	Male	0.1064	0.1015	0.1115
	Female	0.1091	0.1050	0.1121
Sweden	Male	0.1078	0.1047	0.1114
	Female	0.1047	0.1017	0.1078
United States	Male	0.0955	0.0930	0.1016
	Female	0.1110	0.1089	0.1130
England & Wales	Male	0.0969	0.0929	0.1034
	Female	0.1006	0.0985	0.1040

Table A3: Estimate and 95% credible interval for drift (β)

Country	Sex	Estimate	Lower C.I.	Upper C.I.
Australia	Male	0.0033	-0.0100	0.0168
	Female	0.0037	-0.0081	0.0137
Canada	Male	0.0032	-0.0060	0.0139
	Female	0.0039	-0.0064	0.0137
Denmark	Male	0.0014	-0.0053	0.0100
	Female	0.0025	-0.0046	0.0104
Finland	Male	0.0020	-0.0087	0.0154
	Female	0.0038	-0.0057	0.0152
France	Male	0.0127	-0.0029	0.0307
	Female	0.0108	-0.0048	0.0241
Italy	Male	0.0086	-0.0086	0.0251
	Female	0.0047	-0.0135	0.0264
Japan	Male	0.0023	-0.0177	0.0231
	Female	-0.0001	-0.0166	0.0181
Netherlands	Male	0.0039	-0.0065	0.0142
	Female	0.0068	-0.0030	0.0173
Norway	Male	0.0018	-0.0062	0.0092
	Female	0.0027	-0.0050	0.0095
Sweden	Male	0.0012	-0.0052	0.0093
	Female	0.0023	-0.0052	0.0097
United States	Male	0.0038	-0.0090	0.0146
	Female	0.0023	-0.0094	0.0143
England & Wales	Male	0.0028	-0.0069	0.0149
	Female	0.0035	-0.0075	0.0139

Table A4: Estimate and 95% credible interval for magnitude of the period effects (σ_{rw})

Country	Sex	Estimate	Lower C.I.	Upper C.I.
Australia	Male	0.0501	0.0391	0.0681
	Female	0.0447	0.0346	0.0574
Canada	Male	0.0439	0.0330	0.0563
	Female	0.0422	0.0317	0.0549
Denmark	Male	0.0400	0.0315	0.0515
	Female	0.0400	0.0312	0.0516
Finland	Male	0.0488	0.0372	0.0661
	Female	0.0489	0.0368	0.0621
France	Male	0.1064	0.0887	0.1220
	Female	0.0975	0.0826	0.1159
Italy	Male	0.1037	0.0868	0.1275
	Female	0.1114	0.0930	0.1324
Japan	Male	0.0765	0.0601	0.1019
	Female	0.0679	0.0518	0.0894
Netherlands	Male	0.0537	0.0424	0.0664
	Female	0.0536	0.0437	0.0677
Norway	Male	0.0413	0.0327	0.0533
	Female	0.0372	0.0293	0.0471
Sweden	Male	0.0423	0.0342	0.0531
	Female	0.0454	0.0374	0.0548
United States	Male	0.0526	0.0388	0.0647
	Female	0.0527	0.0430	0.0665
England & Wales	Male	0.0646	0.0524	0.0801
	Female	0.0622	0.0512	0.0743

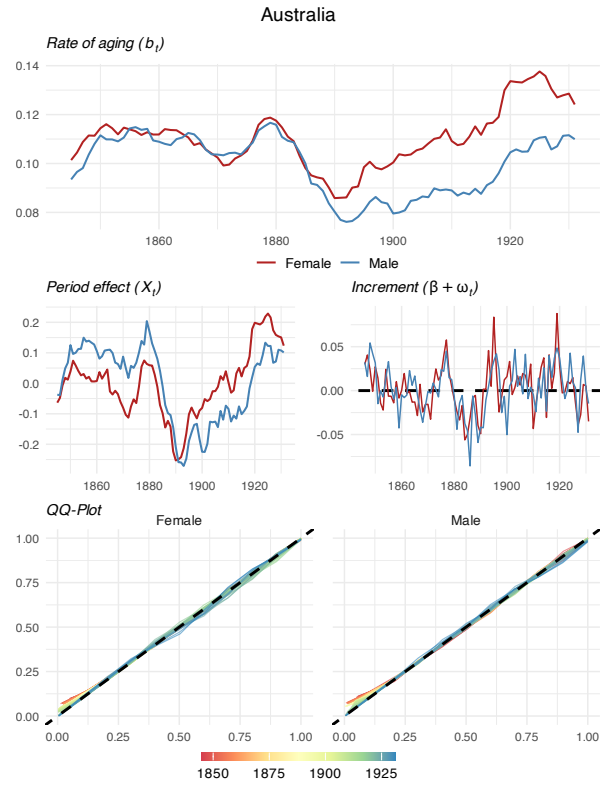


Figure A4: Random walk decomposition for Australia.

A2 Country-Specific Results and Model Diagnostics

This section provides the detailed country-by-country results of the random walk decomposition model. For each of the 12 countries, we present a four-panel figure presenting: the estimated series of the rate of aging, b_t ; the estimated latent period effect, X_t ; the estimated cohort increments, $\delta_t = b_t - b_{t-1}$; and a posterior predictive QQ-plot for model validation. We also provide trace plots for the main model parameters to visually confirm the MCMC convergence discussed in the main text.




Figure A5: Random walk decomposition for Canada.

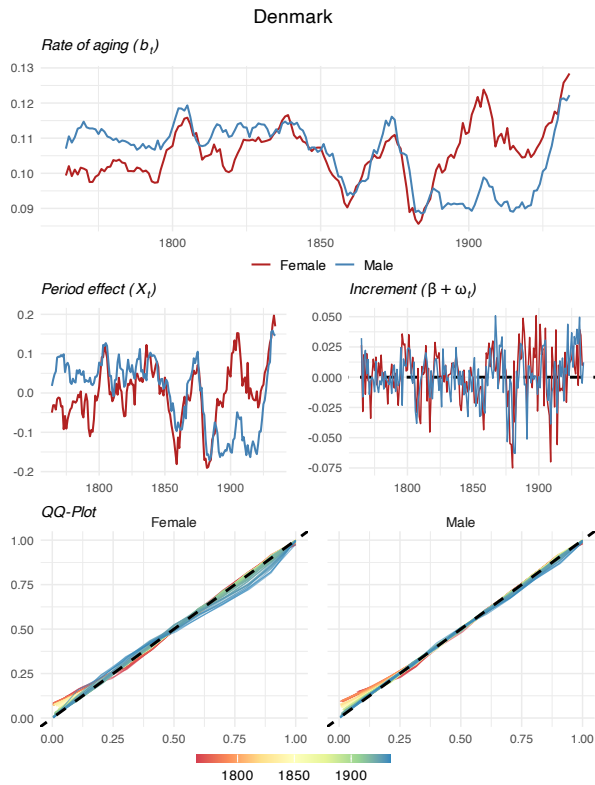


Figure A6: Random walk decomposition for Denmark.




Figure A8: Random walk decomposition for Finland.

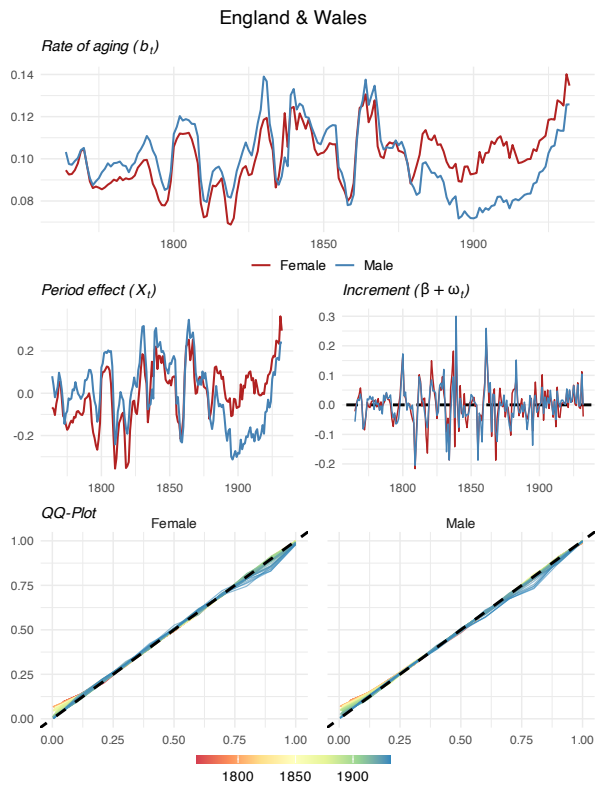


Figure A7: Random walk decomposition for England & Wales.

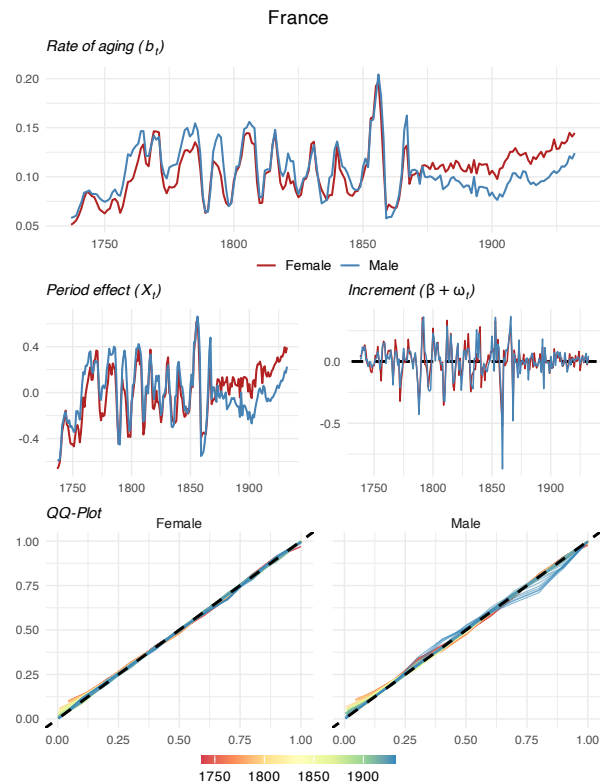


Figure A9: Random walk decomposition for France.

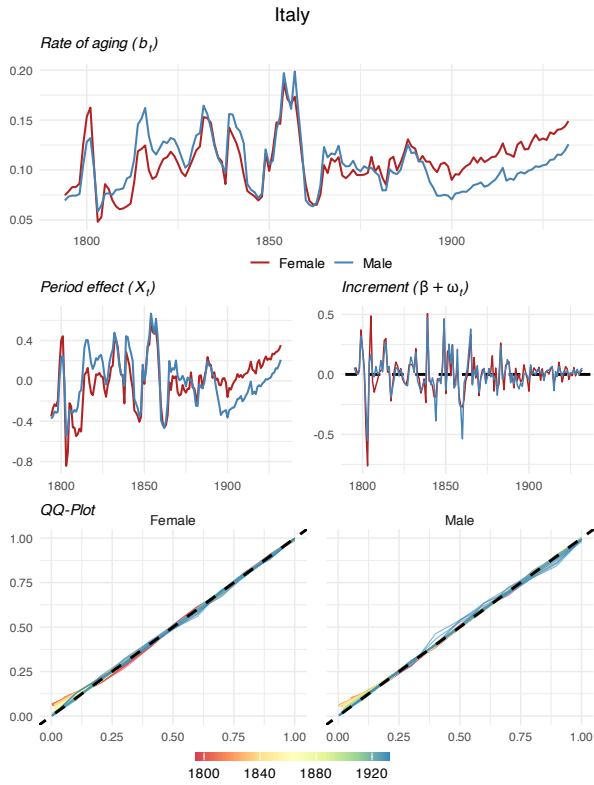


Figure A10: Random walk decomposition for Italy.

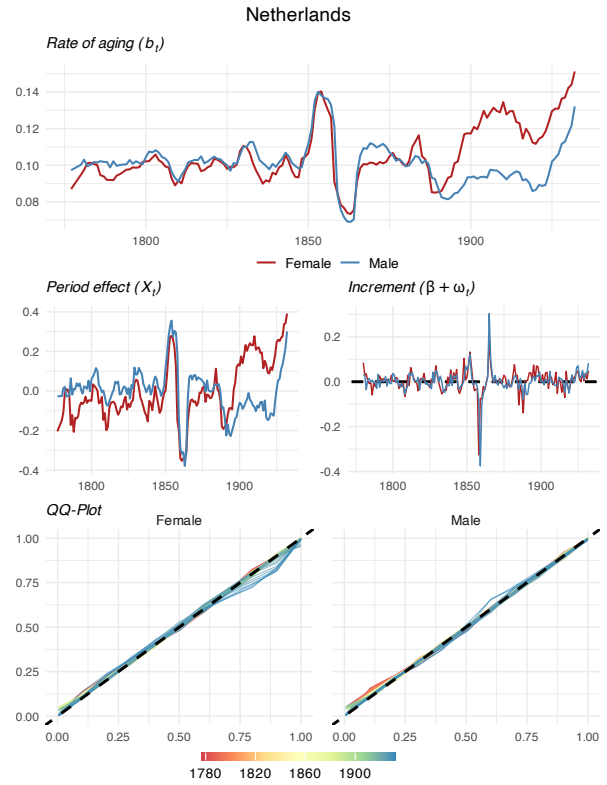


Figure A12: Random walk decomposition for Netherlands.

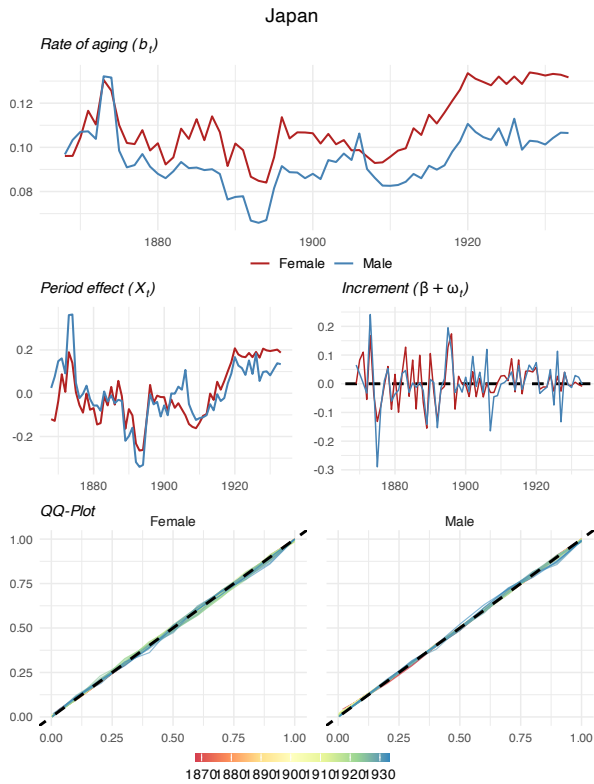


Figure A11: Random walk decomposition for Japan.

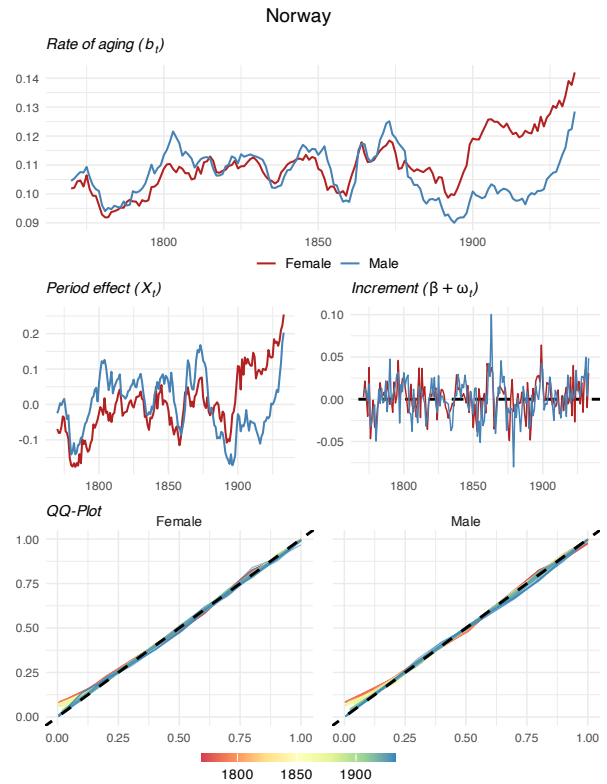


Figure A13: Random walk decomposition for Norway.

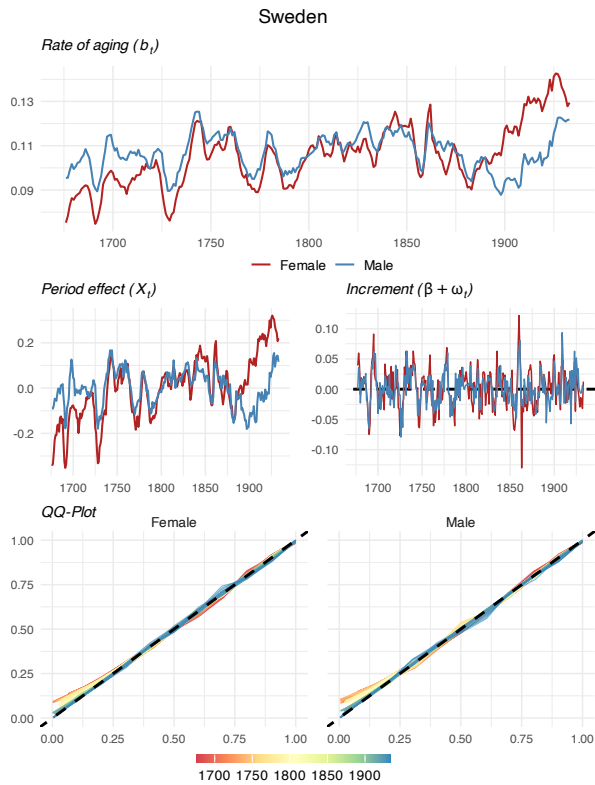


Figure A14: Random walk decomposition for Sweden.

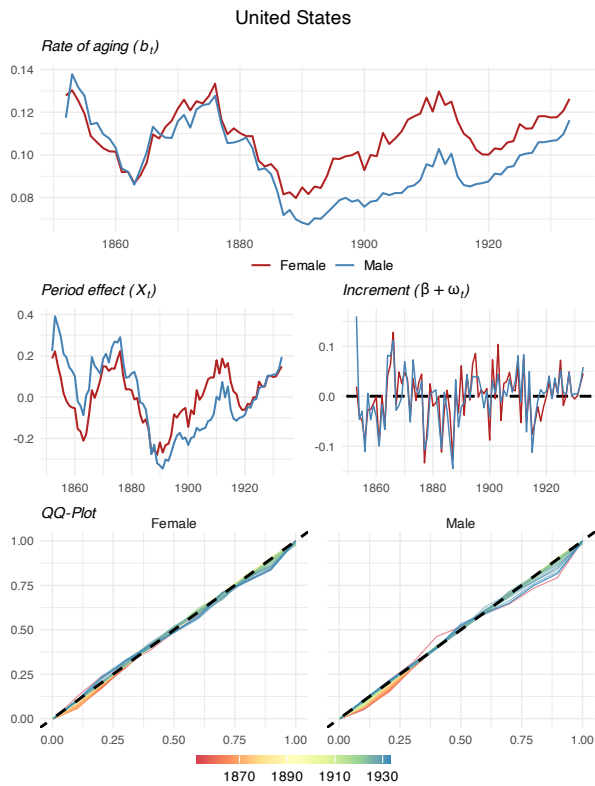


Figure A15: Random walk decomposition for United States.

A2.1 MCMC Convergence Diagnostics

To supplement the \hat{R} statistics reported in the main text, Figure A16 displays the MCMC trace plots for the main parameters for Italy (a high-volatility country) and Sweden (a low-volatility country). The chains are well-mixed and stationary, providing a visual check of convergence.

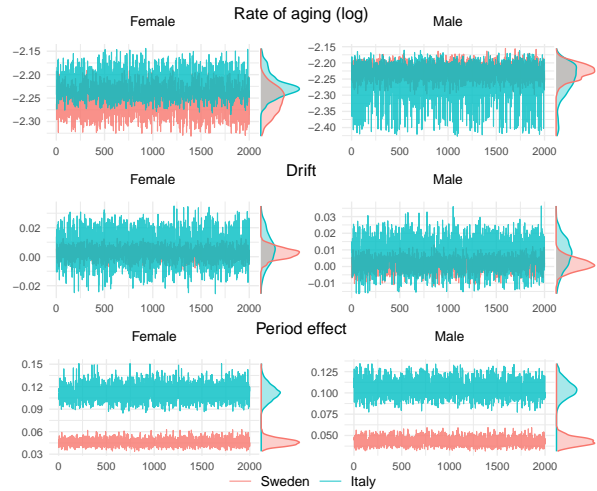


Figure A16: Trace plots for the baseline rate ($\log b$), drift (β), and period effect (σ_{rw}) for Italy and Sweden.