# Inequality of Opportunity in Biological Aging: Evidence from Epigenetic Clocks in the Health and Retirement Study

Daniel Graeber*

November 3, 2026

**Abstract**

How much of the inequality in how fast people age traces to circumstances they did not choose? We measure the inequality-of-opportunity in a third generation epigenetic measure of the rate biological aging, DunedinPACE, which is highly predictive for health and health-care usage. Circumstances fixed in childhood—parental education and occupation, the family's economic situation, race and ethnicity, sex, and region of birth—account for 12.0% of the inequality in the pace of aging; race and ethnicity make up 33.6% of this share and parental education 23.6%. The share peaks around the median, and its composition rotates across the distribution: race and ethnicity and region of birth dominate among the fastest agers, while sex gains weight among the slowest, where it matters as much as either. Circumstances imprint on the molecular level of aging.

**JEL:** I14, D63, J14. **Keywords:** inequality of opportunity, biological aging, epigenetic clocks, DunedinPACE, Shapley decomposition, HRS.

*DIW Berlin, IZA, and CEPA Potsdam. This version: November 3, 2026. Analyses use restricted-access DNA-methylation data from the Health and Retirement Study; only aggregate results are reported. Correspondence: dgraeber@diw.de. I thank the audience of a workshop at the German Institute for Economic Research (DIW Berlin) for helpful comments and suggestions. I used an AI assistant (Claude Code, developed by Anthropic) for editing the text and for writing and debugging the code. All remaining errors are my own.

# 1 Introduction

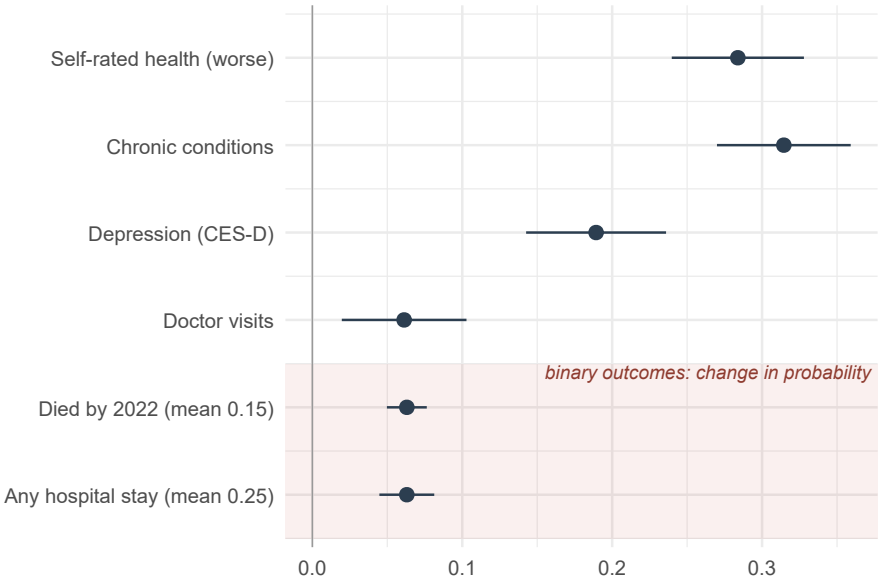
Chronological age is the same for everyone born in a given year, but people do not age at the same rate. Biological aging—the gradual, multi-system decline that raises the risk of disease and death—varies widely among people of the same calendar age. And importantly, much of that variation is present decades before the diseases of old age appear (Belsky et al., 2022). Epigenetic clocks measure this variation directly. They read DNA methylation, the chemical marks that regulate which genes are expressed, from a blood sample (Horvath, 2013; Levine et al., 2018; Lu et al., 2019), and faster biological aging on these clocks predicts morbidity, disability, and mortality over and above chronological age. If the rate of aging is unequally distributed, and if that inequality is tied to characteristics fixed early in life, then the fairness of the distribution of health—a core object of health economics—already manifests at the molecular level long before it causes costs.

This paper analyses how much of the inequality in the *pace* of biological aging reflects *ex-ante inequality of opportunity* (IOP), and which circumstances drive it. We measure the pace of aging with DunedinPACE (Belsky et al., 2022), a third generation DNA-methylation based estimate of the rate of biological aging in the Health and Retirement Study (HRS), and decompose its inequality with an exact Shapley–Owen decomposition. Six circumstance domains —parental education, parental occupation, the family's economic situation in childhood, race and ethnicity, sex, and region of birth—account for 12.0% of the mean log deviation of the pace of aging.

**Theory** We adopt the ex-ante approach to inequality of opportunity that goes back to Roemer (1998): outcome differences caused by circumstances beyond a person's control are unfair, differences caused by effort are not, and inequality of opportunity is measured as the share of outcome inequality that circumstances alone account for (Bourguignon et al., 2007; Ferreira and Gignoux, 2011; Ramos and Van de gaer, 2016). The pace of aging is a natural outcome for this framework. It corresponds to the net rate at which health capital depreciates in the models of Grossman (1972) and Galama and van Kippersluis (2019)—a flow that is central to these models but rarely observed—and theories of biological embedding and weathering predict that childhood socioeconomic position and race leave their mark on it (Hertzman, 1999; Geronimus et al., 2006). The pace of aging is also forward-looking. Figure 1 confirms this: in the HRS, a one-standard-deviation faster pace in 2016 predicts self-rated health worse by 0.28 standard deviations and 0.31 standard deviations more chronic conditions in 2018, and a 6.30-percentage-point higher probability of a hospital stay and a 6.30-percentage-point higher probability of death by 2022. The distribution of the pace of

aging therefore is tightly linked to outcomes of policy relevance.

**Figure 1:** The pace of aging predicts downstream health, healthcare use, and mortality




*Notes.* Each estimate is the survey-weighted coefficient estimate of a downstream outcome on standardized DunedinPACE controlling for age, age$^2$, and the estimated composition of the white blood cells, with 95% confidence intervals. The four continuous health outcomes enter as survey-weighted $z$-scores. The two outcomes in the shaded band are binary indicators. Their labels report the survey-weighted base rate. Downstream outcomes are measured in HRS 2018; mortality is cumulative through 2022.

**Empirical methods** The HRS is the reference panel study of aging in the United States: a nationally representative survey that has followed older Americans and their spouses since 1992, interviewing them every two years about health, cognition, work, income and wealth, and family. In 2016, nurses visited consenting respondents at home and drew blood. A subsequent analysis measured at several hundred thousand sites across the genome, how heavily each site was methylated—the molecular record that inform epigenetic clocks. That linkage delivers the circumstances alongside our measure of the pace of aging, and so makes an analysis of inequality of opportunity possible.

We follow the canonical two-step procedure of the ex-ante IOP literature. First, we estimate $\mathbb{E}[y \mid C]$ via OLS, the part of the pace of aging that a person's circumstances alone predict. Second, we apply an inequality index to this counterfactual distribution and divide by the index of the observed distribution. The primary index is the mean log deviation (MLD); the variance, whose ratio is the familiar $R^2$, and the Gini serve as robustness checks (Bourguignon et al., 2007; Ferreira and Gignoux, 2011). Additionally, an exact Shapley–Owen decomposition apportions the explained inequality across circumstance domains (Shorrocks, 2013). Because it averages each domain's marginal contribution over all orderings, it avoids

the path dependence of sequential decompositions and its shares sum exactly to the total. Because a linear OLS regression could be a to simplistic approximation of the data generating process, we also replicate the analysis with the debiased MLD estimator of Terschuur (2023), which builds on Escanciano and Terschuur (2023). We argue that both estimators agree.

Circumstances account for 12.0% of the mean log deviation of adjusted DunedinPACE. Race and ethnicity is the largest domain, at 33.6% of the explained inequality, followed by parental education (23.6%) and region of birth (22.9%); parental occupation (6.9%), the family's economic situation in childhood (6.6%), and sex (6.4%) account for the rest. The two childhood socioeconomic domains jointly add no statistically detectable explanatory power (Wald test, $p = 0.73$): much of what these represent are already approximated by parental education. A distributional analysis shows: Circumstances account for 9.2% of the inequality at the median pace of aging, against 3.0% at the fast-aging tail and 7.5% at the slow-aging tail, and the composition rotates across the distribution: the share of race and ethnicity falls from 30% among the fastest agers to 20% among the slowest, while the share of sex rises from 1% to 19%.

Two sensitivity analyses qualify these results. Across six epigenetic clocks, the two pace-of-aging measures give race and ethnicity the largest share, whereas three of the four biological-*age* clocks are dominated by sex, which several of them encode by construction. We critically discuss these results, which guide future inquiries into the socio-economic determinants of epigenetic aging. We argue that, whenever available, the most recent generation of epigenetic clocks should be used. We also investigate how sensitive our main result is to the choice of inequality index. The total is nearly identical under the variance (11.6%), and about 0.335 under the Gini, which is less sensitive to the tails. Consequently, if we apply the Gini, the parental education moves ahead of race and ethnicity. We read the results not as a contradiction of our main result, but as a mere reflection that both indices weight changes at different parts of the distribution differently.

The paper contributes to two literatures. The first brings epigenetics into economics, and it is young. Menta et al. (2025) aggregate several epigenetic clocks into a single Multi-EpiGenetic-Age measure and show that faster epigenetic aging in adolescence predicts worse schooling, mental health, and labor-market outcomes, and that early-life adversity—maltreatment, delayed school entry—accelerates it. A larger epidemiological literature documents socioeconomic gradients in epigenetic aging in Europe (Fiorito et al., 2017), in the UK (Hughes et al., 2018), and in the U.S. (Raffington et al., 2021; Schmitz et al., 2022). Crimmins et al. (2021) document the age, sex, race, and education gradients of thirteen HRS clocks, and Faul et al. (2023) their prognostic value in the same sample. Schmitz and Duque

(2022) show that in-utero exposure to the Great Depression is still visible in the epigenetic aging of HRS respondents eighty years later. This work establishes that epigenetic aging responds to the socioeconomic environment and predicts later outcomes. It does not, however, ask the distributional question within a normative framework: these studies estimate conditional means one socioeconomic characteristic at a time, without distinguishing circumstance from effort, and mostly for biological-*age* clocks—the stock of depreciated health capital—rather than for the pace of aging, the rate of depreciation.

The second literature measures inequality of opportunity. It has been applied to income (Bourguignon et al., 2007; Ferreira and Gignoux, 2011), wealth (Graeber et al.; Hufe et al., 2025), to self-reported and clinical health (Fleurbaey and Schokkaert, 2009; Rosa Dias, 2009; Trannoy et al., 2010), and, closest to our design, to biomarkers: Davillas and Jones (2020) decompose ex-ante IOP in a panel of UK blood-based biomarkers. This literature has not yet had a measure of the *rate* of aging to work with. We apply the methodology of IOP to a molecular, pre-clinical marker that is measured identically for everyone, is highly predictive for future outcomes, and show that its composition, not only its level, varies across the distribution.

Section 2 sets out the theory of inequality of opportunity and illustrate how it can inform economic theory. Section 3 introduces epigenetic clocks, the HRS methylation sample, and the raw circumstance gaps in the pace of aging. Section 4 sets out the measurement framework, the decomposition, the distributional analysis, inference, and the debiased estimator. Section 5 reports the total IOP, its composition, and its distributional profile. Section 6 investigates the sensitivity of our results to the clock and the inequality index. Section 7 concludes.

# 2 Inequality of opportunity in the pace of aging

**Circumstances, effort, and types.** Following Roemer (1998), write the outcome of person $i$ as $y_i = f(C_i, E_i, u_i)$, where $C_i$ comprises the *circumstances*, which lie outside the person's control, $E_i$ *effort*, for which the person may be held responsible, and $u_i$ luck. Equality of opportunity joins two principles (Roemer and Trannoy, 2016): *compensation*, under which outcome differences caused by circumstances are unfair, and *reward*, under which differences caused by effort are not. Individuals that share the same circumstances form a *type*.

**Ex-ante inequality of opportunity.** The *ex-ante* approach attributes to the opportunity set of each type the outcome its members can expect, $\mu(C) = \mathbb{E}[y \mid C]$; the *ex-post* approach instead compares people who exert the same effort (Fleurbaey and Peragine, 2013). Replac-

ing each person's outcome by the expected outcome of their type removes the inequality within types—the part due to effort and luck—and leaves the inequality between types. For an inequality index $I(\cdot)$, ex-ante inequality of opportunity is the share of total inequality that is attributable to this (Checchi and Peragine, 2010; Ferreira and Gignoux, 2011):

$$\text{IOP}_I \;=\; \frac{I\big(\mathbb{E}[y \mid C]\big)}{I(y)}. \tag{1}$$

It is zero when every type can expect the same outcome and one when circumstances determine the outcome completely.

**Why the ex-ante view.** Effort itself depends on circumstances, $E_i = h(C_i, v_i)$, so the outcome has the reduced form $y_i = g(C_i, v_i, u_i)$ (Bourguignon et al., 2007). The natural candidates for effort—smoking, diet, exercise, the use of medical care—are themselves shaped by circumstances, and Roemer (1998) measures effort by a person's rank within the effort distribution of their type, so that the part of behavior a type predicts counts as circumstance. Because $\mu(C)$ projects the reduced form on circumstances, it accounts for their direct effect and with the effect they exert through effort. It does so without requiring effort to be observed (Fleurbaey and Schokkaert, 2009; Rosa Dias, 2009). Two limits follow from the construction. Only *observed* circumstances enter $C$, so (1) is a lower bound on the role of circumstance (Ferreira and Gignoux, 2011); and because $g$ is a reduced form, the measure describes how much circumstances account for, not the causal effect of any one of them.

**The pace of aging as the outcome.** In the models of health capital of Grossman (1972) and Galama and van Kippersluis (2019), a person's stock of health evolves as $H_{t+1} = (1 - \delta_t)H_t + I_t$: it depreciates at rate $\delta_t$ and is replenished by investment $I_t$ in time, behavior, and medical care. A biological-*age* clock estimates the depreciation accumulated so far, a stock; the pace of aging estimates the net rate at which health capital is currently lost, $(H_t - H_{t+1})/H_t = \delta_t - I_t/H_t$, a flow (Section 3.1). Circumstances can act on both terms. Conditions in early life are *biologically embedded* (Hertzman, 1999): they are embodied in the physiological systems whose decline is aging (Krieger, 2005), and their consequences surface decades later in adult health (Almond and Currie, 2011; Almond et al., 2018). The cumulative stress of social disadvantage accelerates the decline over adulthood—the "weathering" documented for Black Americans (Geronimus et al., 2006)—while resources and health investment slow it (Dalgaard and Strulik, 2014). These accounts predict that the family's socioeconomic position and race shape the *rate* of aging in old age, not only its accumulated level. In the ex-ante framework, $\mathbb{E}[y \mid C]$ is then the pace at which a type can expect to age,

and (1) is the share of the inequality in the pace of aging that reflects unequal opportunities. The measure suits the framework for an additional reason: it is objective and cardinal, it avoids the reporting and justification biases of survey health measures, and it is observable before disease is diagnosed (Davillas and Jones, 2025), making it particularly relevant for policy.

# 3 Data

## 3.1 Background: what epigenetic clocks measure

As in Section 2, we think of health as a stock of capital that follows a law of motion: it determines how much time a person can allocate to (market) work and how they feel, and it is the result of past health capital, investments such as physical activity or diet, and depreciation (Grossman, 1972). Chronological age measures elapsed *time*. *Biological* age the accumulated degradation of biological functioning. The latter correspond to the cumulated net depreciation over time. The two differ because people of the same calendar age have aged by different amounts. The reason could be that different individuals are exposed to different environments.

*Epigenetics* is concerned about the study of molecules and mechanisms that can sustain alternative gene activity states in the context of the same DNA sequence (Cavalli and Heard, 2019). These mechanisms come in several forms—most prominently methylation of the DNA itself and modifications of the proteins the DNA is wrapped around—and *DNA methylation* is the most frequent and best-measured of them: small chemical tags (methyl groups) that attach to DNA at hundreds of thousands of specific sites and switch genes on or off. While the genome is fixed at conception, *methylation* is dynamic—it responds to time and to the environment.

An *epigenetic clock* is a prediction model estimated on these sites, typically a penalized (elastic-net) regression that selects a few dozen to about a thousand sites from the hundreds of thousands measured and fits these to a specific target that proxies (biological) aging.

Clocks differ in *what* target they are trained to predict, and this matters for interpretation. *First-generation* clocks (Horvath, 2013; Hannum et al., 2013) are trained to predict chronological age itself; their residual—how much older or younger than one's calendar age the molecular data imply—is termed *age acceleration*. *Second-generation* clocks are trained not on the calendar but on measurements of biological functioning: PhenoAge (Levine et al., 2018) on a composite of clinical biomarkers, and GrimAge (Lu et al., 2019) on mortality itself, so that a unit of GrimAge is a unit of predicted lifespan lost. These markers correspond

to estimates of the current *stock* of accumulated depreciation (a biological age, in years). This can also be thought of as the inverse of health capital.

Our primary outcome, the third generation clock DunedinPACE (Belsky et al., 2022), measures a *flow* rather than a stock. It was trained on the longitudinal *rate* of decline across organ systems in a birth cohort followed for two decades, and so estimates the current *pace* of aging—the derivative of biological age with respect to time—rather than its level. It is normed so that 1.0 is one biological year of aging per calendar year; a value of 1.05 means aging five percent faster than the calendar. Two features suit it to an inequality-of-opportunity analysis. First, as a rate, DunedinPACE provides more forward looking information than a stock of health (or deficits). Consequently, the knowledge about the rate of aging for interventions whereas a stock does not. Second, it maps directly into an underexplored fundamental of theory in health economics, *i.e.*, the net depreciation of health capital. This offers important insights into these fundamentals of economic theory.

## 3.2 Sample and variables

**The Health and Retirement Study.** The HRS (Sonnega et al., 2014) is the reference panel study of aging in the United States: a nationally representative, biennial survey of Americans over the age of fifty and their spouses, running since 1992 and refreshed with new birth cohorts as earlier ones age.[1] Roughly twenty thousand respondents are interviewed in each wave about health, cognition and functional limitation, work and retirement, income and wealth, family, and health-care use and expenditure, and the survey has been linked to administrative earnings and mortality records. The HRS is uniquely designed to investigate IOP in biological aging: it observes the same people repeatedly over decades and collects detailed retrospective information about childhood circumstances.

**Collection of methylation data.** In 2016 the HRS fielded a Venous Blood Study: trained phlebotomists visited consenting respondents at home and drew a venous blood sample. DNA was extracted from the white blood cells of a subsample of those respondents—selected by the HRS—and assayed for *methylation*—the chemical tags described in Section 3.1—on a commercial microarray that measures, at each of several hundred thousand individual sites in the genome, the fraction of a respondent's cells that carry a methyl group at that site. The HRS processed the resulting readings through the standard quality-control and normalization steps and then imputed the published clock algorithms.

[1]The HRS (Health and Retirement Study) is sponsored by the National Institute on Aging (grant number NIA U01AG009740) and is conducted by the University of Michigan.

**Sample and weighting.** We link epigenetic clocks to the respondents' survey records, which contains detailed information on circumstances. After restricting to respondents aged 50–95 at the time of the blood draw, the analysis sample is $N = 3{,}866$ individuals. Because the methylation subsample is only a subsample of a complex survey, all estimates are survey-weighted, using the weight the HRS constructs specifically for analyses of the epigenetic clocks: it adjusts for differential participation in the blood draw and post-stratifies to the 2016 HRS population born before 1960, by age, sex, and race and ethnicity.

**Outcome.** Our primary outcome is the reciprocal of DunedinPACE, which we call DunedinPACEr (Section 3.1). We invert the raw pace of aging so that the outcome points in the direction inequality indices assume: DunedinPACE is a "bad"—a higher value means faster aging—whereas an inequality index is conventionally written for a "good", something of which more is better. Thus, DunedinPACEr is increasing in *slower* aging, so that a distribution shifted upwards is unambiguously a better one.[2] Although the pace of aging is by construction only weakly related to chronological age (Appendix Figure A3), we residualize the outcome on a full set of single-year age indicators throughout, so that all reported inequality is strictly net of age.

**Blood-cell composition.** We also net out the composition of the blood sample itself. The DNA is extracted from white blood cells, and a blood sample is a mixture of cell types—granulocytes, several kinds of lymphocytes, and monocytes—each with its own methylation pattern. A person's methylation profile, and every clock computed from it, varies with the mix of cells in the blood sample as well as with the pace of aging. This mix also varies with recent infection, medication, and smoking, *i.e.*, idiosyncratic shocks. Our estimates of IOP would be biased in unknown directions if we would not account for this source of measurement error. We residualize the outcome on these shares together with the age indicators, similar to related work (Crimmins et al., 2021; Schmitz et al., 2022).

As sensitivity analysis, we also study five additional clocks that span both generations and both the stock and flow concepts: GrimAge (Lu et al., 2019), PhenoAge (Levine et al., 2018), the first-generation Horvath (Horvath, 2013) and Hannum (Hannum et al., 2013) clocks, and the mPOA pace measure. Each of them is inverted in the same way, since for a biological-*age* clock too a higher reading at a given calendar age is the worse outcome.

**Circumstances.** We use six circumstance domains, all fixed by early adulthood. They follow the IOP literature, which conditions on parental education and occupation, race or

[2] In Section 6.2, we find that the polarization of the outcome does not matter for the results.

ethnicity, and place of birth (Bourguignon et al., 2007; Ferreira and Gignoux, 2011), and its applications to health, which also condition on parents' social background (Trannoy et al., 2010), sex, and the family's economic situation in childhood (Rosa Dias, 2009; Davillas and Jones, 2020). These comprise (i) *parental education*—father's and mother's completed schooling, each binned into seven categories from "no schooling" to "college and above"; (ii) *race and ethnicity*—race (White, Black, other) and Hispanic ethnicity; (iii) *sex*; (iv) *region of birth*—the nine U.S. Census divisions, a residual U.S.-born category, and foreign-born; (v) *family socioeconomic status in childhood*—whether the family was pretty well off, about average, or poor financially while the respondent was growing up (from birth to age 16), whether financial difficulties ever forced the family to move, whether it received financial help from relatives, and whether the father was ever out of work for several months or more; and (vi) *parental occupation*—the father's occupation when the respondent was 16, in six broad groups, and how much of the time the mother worked outside the home. The last two domains are based on retrospective questions the HRS asks each respondent once. Every circumstance keeps a separate "unknown" category for respondents for whom it is not recorded—because they did not know or did not answer, or did not live with the parent in question—so no respondent is dropped for a missing circumstance. This matters mainly for parental education, which is unknown for 13% of respondents for the father and 7% for the mother (survey-weighted). Excluding these individuals would cause us to drop those who are less likely to provide this information. If non-response is also associated with the socio-economic status of these individuals, this exclusion could bias our results. Appendix Table A1 reports sample composition and mean DunedinPACE by circumstance category.

## 3.3 Descriptive properties of DunedinPACE

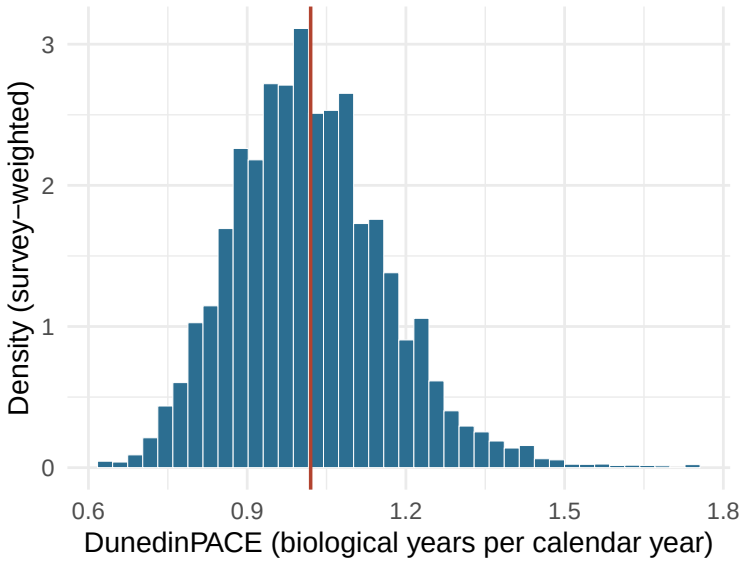
Figure 2 shows the survey-weighted distribution of DunedinPACE. The pace of aging is close to symmetric and unimodal, centered near one (weighted mean 1.01, standard deviation 0.14). That means that the average older American in the sample ages at roughly one biological year per calendar year. At the same time, the spread is wide—the 90th percentile ages about 40% faster than the 10th.

Figure A3 in Appendix A.4 plots the weighted mean of DunedinPACE within single-year age bins. The gradient is small—about +0.0024 per year of age, so chronological age accounts for a small share of the variance—which confirms that DunedinPACE is close to a pure rate rather than a restatement of calendar age, and that our age residualization removes only a small fraction of the signal.

Figure 3 compares the full survey-weighted distribution of DunedinPACE across circum-

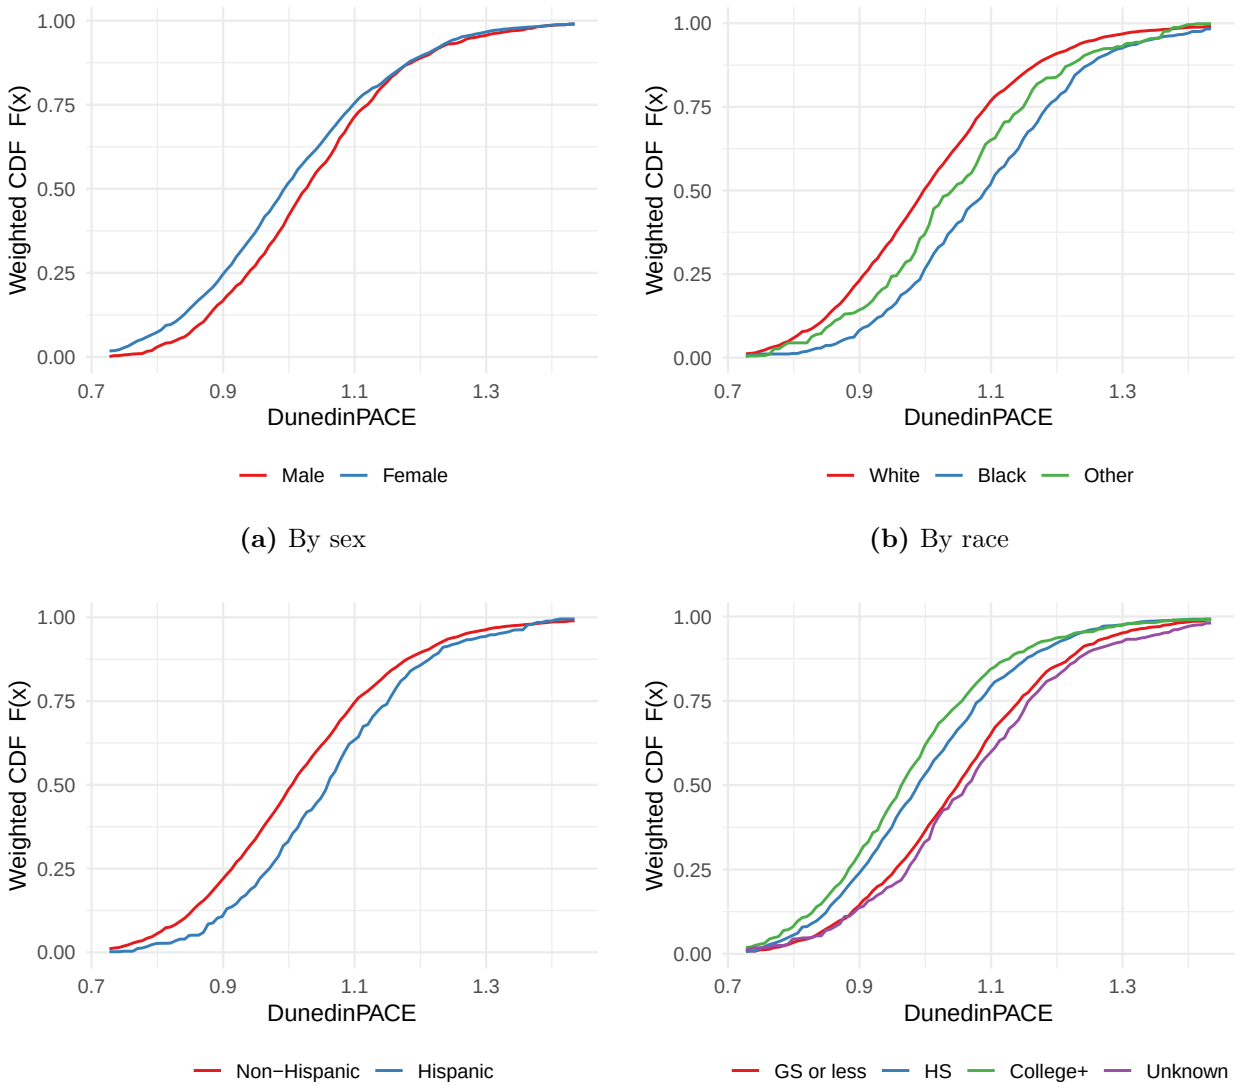
stance groups. A curve shifted to the *right* indicates faster aging. The distributions separate visibly and in the expected directions: the Black and Hispanic distributions are shifted right (faster aging) relative to the White and non-Hispanic distribution; women's distribution lies to the left of men's (slower aging); and the distribution shifts rightward as parental education declines. Noteworthy, the distribution of individuals who do not provide information about their father's education is similar or worse to the distribution of individuals whose father has the lowest educational degree. This indicates that we would discard observations from parts of the distributions that characterize fast agers.

**Figure 2:** Survey-weighted distribution of DunedinPACE




*Notes.* The red line marks the weighted mean. DunedinPACE is normed so that 1.0 is one biological year of aging per calendar year.

**Figure 3:** Distribution of DunedinPACE by circumstance group




**(a)** By sex

**(b)** By race

**(c)** By ethnicity

**(d)** By father's education

*Notes.* Survey-weighted cumulative distribution functions of DunedinPACE, evaluated on a common grid. A curve shifted to the right indicates faster aging. In panel (d), GS or less: no schooling or grade school (up to eight years); HS: some high school or high school; College+: some college or more; Unknown: father's education not recorded.

# 4 Empirical strategy

## 4.1 Measuring and decomposing inequality of opportunity

**Outcome and circumstances.** We estimate the ex-ante measure (1) of Section 2. In our application $y_i$ is DunedinPACEr adjusted for age and blood-cell composition—the residual

of a survey-weighted regression on single-year age indicators and the estimated cell shares, recentred to its mean[3]—and $C_i$ collects the twelve circumstance variables of the six domains of Section 3.2.

**The index.** Our primary index is the *mean log deviation* (MLD, the generalized-entropy index GE(0)), $\mathrm{MLD}(y) = \log \bar{y} - \overline{\log y}$. It is additively decomposable with population-share weights (Shorrocks, 1980), and it is the only such index that is also *path independent* (Foster and Shneyerov, 2000). That is, computing the between-type component first, by smoothing, or the within-type component first, by scaling every type to the common mean, gives the same split. The MLD-based IOP is therefore unambiguously the between-type share of inequality, and its complement the within-type share (Checchi and Peragine, 2010; Ferreira and Gignoux, 2011). We report the variance—whose ratio is the familiar $R^2$—and the Gini as robustness (Section 6.2).

**Estimation.** We estimate $\mathbb{E}[y \mid C]$ by survey-weighted least squares, with each circumstance entering as a full set of category indicators ($p = 50$ parameters besides the constant); Section 4.2 complements the results with a debiased machine-learning estimator. Besides the two limits of the measure itself (Section 2: it is a lower bound, and it describes a reduced form), estimation adds a third. The estimated share is biased *upward* in finite samples, because the fitted values inherit the sampling noise of the $p$ estimated coefficients, and the index reads that noise as inequality between types. The bias grows with every circumstance added, so a richer set of circumstances trades a smaller downward bias for a larger upward one (Brunori et al., 2019).

**Shapley–Owen decomposition.** The total $\mathrm{IOP}_I$ measures how much opportunity matters. But we also aim to understand *which* circumstances matter. Writing $\nu(S) = I\big(\mathbb{E}[y \mid C_S]\big)/I(y)$ for the share reproduced by the subset $S$ of circumstances, a naive sequential decomposition attributes to each domain the increment $\nu$ from adding it, but this depends on the order of entry when circumstances are correlated. The Shapley value resolves the ambiguity by averaging each domain's marginal contribution over all $K!$ orders of entry

[3]The residual has mean zero by construction, and the MLD needs a strictly positive outcome. We therefore add back the weighted mean of DunedinPACEr before purging, which leaves the mean of the outcome unchanged and the residualization otherwise untouched. The MLD is invariant to rescaling but not to adding a constant, so the level at which the residual is centred is a normalization that enters the index. It matters little here: for an outcome with mean $\bar{y}$ and small relative spread, the MLD is approximately $\mathrm{Var}(y)/(2\bar{y}^2)$, so the *share* explained by circumstances is, to first order, the variance share and does not depend on $\bar{y}$; the variance-based and the MLD-based shares in Section 6.2 are close (11.6% and 12.0%). Because the mean of DunedinPACEr is near one and its standard deviation near 0.14, the approximation is accurate.

(Shorrocks, 2013; Huettner and Sunder, 2012). For domain $k$,

$$\phi_k \;=\; \sum_{S \subseteq N \setminus \{k\}} \frac{|S|!\,(K-|S|-1)!}{K!} \Big[ \nu(S \cup \{k\}) - \nu(S) \Big], \tag{2}$$

where $N = \{1, \ldots, K\}$. The values are *efficient* ($\sum_k \phi_k = \mathrm{IOP}_I$), *symmetric*, and treat each domain as an indivisible group of dummies (the Owen extension to grouped regressors). We compute the values exactly, with the MLD as the value function. Because the Owen values within a domain sum to that domain's Shapley value, the domain decomposition needs only the $2^K = 64$ coalitions of the six domains. The variable-level decomposition of Appendix Table A2 uses all $2^{12}$ subsets of the twelve circumstance variables.

**The distributional profile.** The two steps above evaluate inequality at one point of the distribution. We also want to know whether IOP differs between the slow-aging and the fast-aging ends of the distribution of DunedinPACEr. For this we use the recentered influence function (RIF) of Firpo et al. (2009). Replacing the outcome by its RIF at quantile $\tau$,

$$\mathrm{RIF}(y_i; q_\tau) \;=\; q_\tau + \frac{\tau - \mathbf{1}\{y_i \le q_\tau\}}{f_y(q_\tau)}, \tag{3}$$

and running the same two steps on it recovers the effect of circumstances on the *unconditional* $\tau$-quantile of the pace-of-aging distribution. We keep the primary index throughout: at each $\tau$ we take the RIF of DunedinPACEr, purge it of age and cell composition with the same regressors, and compute the survey-weighted MLD share of the circumstances together with its Shapley–Owen split across the six domains. We report that split in two forms: as *absolute* Shapley values, the contribution of each domain to the MLD share at that quantile, which sum to the share; and as *relative* values, each domain's fraction of the share, which sum to one.

**Inference.** We obtain standard errors for the total share and for every domain share by a nonparametric bootstrap that resamples individuals ($B = 1{,}000$ replications) and recomputes the full weighted MLD decomposition within each replication. We report the bootstrap standard errors and, where an interval is shown, the normal-approximation 95% confidence interval: the point estimate plus or minus 1.96 bootstrap standard errors.

## 4.2 Debiased estimation of the total

The above estimate is based on a linear OLS regression, the standard practice in the *ex-ante* literature. Machine-learning methods, such as Lasso regressions or random forests, have become popular since they can partition the rich space of circumstances without inheriting the curse of dimensionality (Brunori et al., 2023). However, the shrinkage they impose, and the cross-validation used to tune it, introduce a bias in the first stage, which the final point estimate can inherit (Escanciano and Terschuur, 2023).

Escanciano and Terschuur (2023) address this. Their observation is that the inequality of $\mathbb{E}[y \mid C]$ is a *U-statistic* in the first-stage regression, and therefore admits a Neyman-orthogonal moment representation: a correction term can be added whose effect is to make the estimator's dependence on the first stage second-order. For the MLD, our main index, Terschuur (2023) states the analogous orthogonal moment: writing $\gamma_0(C) = \mathbb{E}[y \mid C]$, the MLD of the fitted values is $\ln \mathbb{E}[\gamma_0] - \mathbb{E}[\ln \gamma_0]$, and the correction term $(y - \gamma_0)/\gamma_0$ added inside the second expectation makes the moment's derivative in $\gamma$ vanish at $\gamma_0$, so the first-stage error again enters only at second order.[4]

Orthogonality removes the first-order effect of the first-stage error, but not the second-stage error. And for the MLD the remainder has a fixed sign. Write $r = (\hat{\gamma} - \gamma_0)/\gamma_0$ for the relative error of the cross-fitted first stage. The expected error of the debiased MLD is $-\mathbb{E}[f(r)]$ with $f(r) = \ln(1 + r) - r/(1 + r) \geq 0$, so the estimator is biased *downward*, by about $\frac{1}{2}\mathbb{E}[r^2]$, whatever the learner and in whichever direction it errs. Relative to the size of the target the bias is roughly $-\mathrm{MSE}(\hat{\gamma})/\mathrm{Var}(\gamma_0)$: the share of the variation in $\mathbb{E}[y \mid C]$ that the first stage fails to capture. It therefore shrinks as the first stage improves, but it is not negligible when, as here, circumstances explain only about a tenth of the variance of the outcome. Appendix A.3 derives these statements step by step and confirms them in a simulation. The two estimators err in opposite directions: the OLS plug-in upward through overfitting, the debiased estimator downward through the remainder. This way, for a fixed set of circumstances and sample size, we can bound IOP between the OLS estimand, which corresponds to an upper bound, and the debiased ML estimand, which constitutes a lower bound.

We implement the MLD version ourselves from Terschuur (2023), because the authors' `ineqopp` package for R computes only the Gini. We use two first stages, a Lasso and a random forest, tune their hyperparameters by five-fold cross-validation before estimating

[4]The construction is that of double/debiased machine learning (Chernozhukov et al., 2018). A moment is *Neyman-orthogonal* when its expectation does not move, to first order, as the first stage moves away from the truth; *cross-fitting*—fitting the first stage for each observation on folds that exclude it—keeps an observation's own noise out of its fitted value. Together they remove, to first order, both the overfitting that inflates a standard plug-in and the regularization bias of a learner.

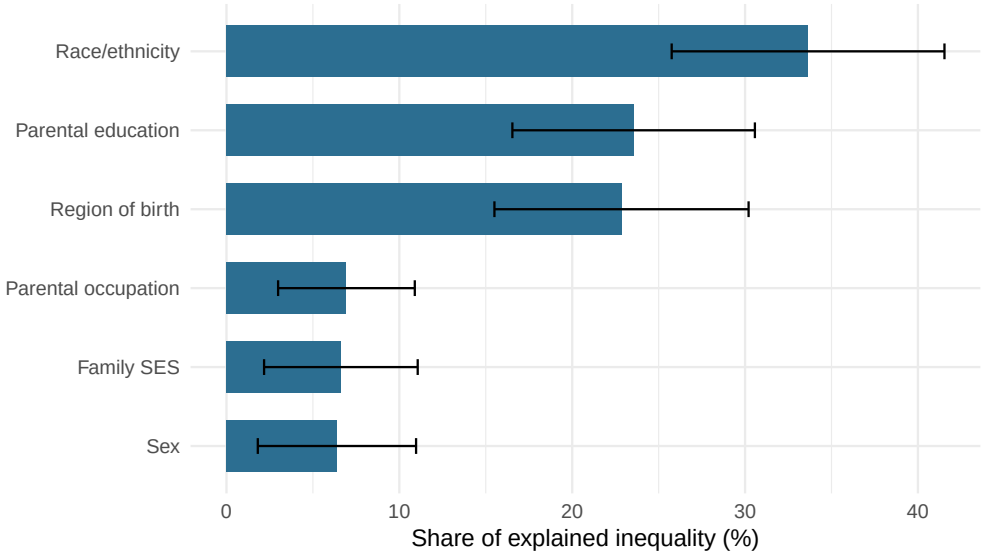
IOP, cross-fit the moment on five folds, and cluster the influence-function standard errors by household.

# 5 Results

## 5.1 Inequality of opportunity in biological aging

Circumstances account for 12.0% of the mean log deviation of adjusted DunedinPACEr (bootstrap standard error 1.2 percentage points, 95% CI [9.7%, 14.3%]; Table 5 and Figure 4). Race and ethnicity is the largest domain, with 33.6% of the explained inequality, followed by parental education at 23.6%. Together the two account for more than half. Region of birth (22.9%), sex (6.4%), parental occupation (6.9%), and the family's economic situation in childhood (6.6%) account for the remainder. Within these domains, the variable-level decomposition (Appendix Table A2) attributes the largest single share to race (30%), followed by region of birth (22%), father's education (17%), and mother's education (9%). Among the items characterizing the family's economic situation in childhood, none explains more than 3%.

**Figure 4:** Circumstance shares of explained inequality in the pace of aging




*Notes.* Shapley–Owen shares of the explained mean log deviation of DunedinPACEr adjusted for age and blood-cell composition, with 95% confidence intervals: the share plus or minus 1.96 bootstrap standard errors (1,000 replications).

**The debiased MLD.** Table 1 repeats the analysis, using the orthogonal moment and the ML first stage for the MLD (Terschuur, 2023), which we implement directly (Section 4.2).

This is the number comparable to our main estimate based on the linear OLS model as a first stage. We use a Lasso and a random forest as first stages to test how robust the estimate is to the functional form of the first stage. The debiased total is 0.088 $[0.068, 0.107]$ with a Lasso first stage, and the random forest gives the same value to the third decimal.

In percentage terms the debiased total is 8.8% of the inequality in the pace of aging, against 12.0% for our OLS estimate. The gap is not evidence of a wrong first stage: the two estimators err in opposite directions. The OLS plug-in is biased upward, because its fitted values inherit the sampling noise of 50 estimated coefficients (Section 4.1); the debiased estimator removes that first-order bias but carries a negative second-order remainder (Section 4.2). In expectation the two estimates therefore bracket the true share. The debiased estimator thus confirms the order of magnitude of the main decomposition, from a moment condition whose first-stage error enters only at second order.

Which circumstances the debiased estimator relies on is reported in Appendix A.2 as leave-one-out partial effects (Terschuur, 2023): the drop in the debiased total when a domain is removed and the moment is re-estimated, as a share of the total. Removing race and ethnicity lowers the total by 31%, region of birth by 10%, and sex by 8%, all significant at the 5% level, and parental education by 8%, just short of it. Removing the family's economic situation in childhood ($-1\%$) or parental occupation ($-1\%$) leaves it unchanged within sampling error. These effects are not directly comparable with the Shapley–Owen shares of Table 5. A leave-one-out effect measures what a domain adds *given all remaining circumstances*, so it depends on which other circumstances are in the set. When domains are correlated, what they explain jointly is attributed to neither of them, and the six effects sum to only 56% of the debiased total. The Shapley–Owen value instead averages a domain's marginal contribution over all subsets of the other domains, so it does not hinge on one particular set of remaining circumstances and its shares sum to the total. However, nevertheless, the analysis based on the debiased ML first stage largely confirms our main analysis.

**Table 1:** Debiased MLD inequality of opportunity in the pace of aging

| | Debiased | | Plug-in |
|---|---|---|---|
| First-stage learner | Relative IOp | 95% CI | Relative IOp |
| Lasso | 0.088 | [0.068, 0.107] | 0.084 |
| RF | 0.088 | [0.069, 0.106] | 0.075 |

*Notes.* Relative inequality of opportunity based on the *MLD*, $N$ = 3,866, survey-weighted. "Debiased" estimates the orthogonal moment of Terschuur (2023) with the first stage cross-fitted on five folds; the standard errors are influence-function based and clustered by household. "Plug-in" evaluates the MLD at the same first-stage fit without the correction. The first-stage hyperparameters are chosen by 5-fold cross-validation before estimation. The outcome is DunedinPACEr, adjusted for age and blood-cell composition, and the circumstances are the twelve variables of the six domains in Table 5. Debiased point estimates carry a negative second-order remainder (Section 4.2); no correction for it is applied to the numbers in this table.

## 5.2 Distributional analysis

We find that the conclusions about IOP in the rate of biological aging differ across the distribution. This is important to understand since individuals are assumed to be risk averse and the fast-aging tail is where the pace of aging translates most directly into disease and early death. We therefore show how IOP in the rate of biological aging, and its composition, vary across the distribution, using the RIF profile of Section 4.1.

Table 2 and Figure 5 report the distributional profile of IOP in biological aging. The *level* of opportunity is sharply increasing from the bottom to the middle of the distribution, after which it remains relatively stable: At $q_{10}$, IOP in biological aging is 3.0% and increases to 9.2% at $q_{50}$. Then, at $q_{75}$ and $q_{90}$, IOP corresponds to 8.3% and 7.5%. However, the 95% confidence interval suggests that we can not reject the absence of a trend beyond the median.

The *composition* of IOP rotates across the distribution. Because the outcome is DunedinPACEr, low quantiles are the *fastest* agers and high quantiles the *slowest*. Race and ethnicity and region of birth dominate the fast-aging tail—30% and 34% of the explained inequality at $q_{10}$—and race and ethnicity falls to 20% among the slowest agers. For sex, it is the other way around, rising from 1% at the fast-aging tail to 19% at the slow-aging tail, where it is as large as race and ethnicity (20%) and region of birth (21%). The family's economic situation in childhood also gains weight toward the slow-aging tail (from 5% to 12%) and

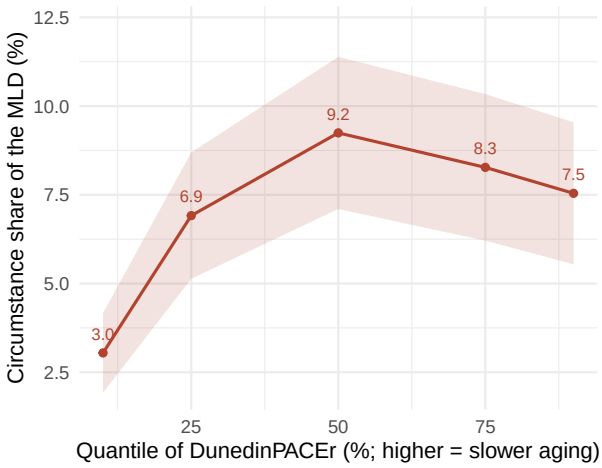
parental education peaks in the middle (27% at the median). The fast-aging tail reflects race, ethnicity, and region of birth. At the slow-aging tail the composition is more even, and sex matters as much as either. The underlying unconditional-quantile gradients are reported in Appendix A.1.

**Table 2:** Distributional IOP profile: the MLD share and its Shapley decomposition by quantile of DunedinPACEr

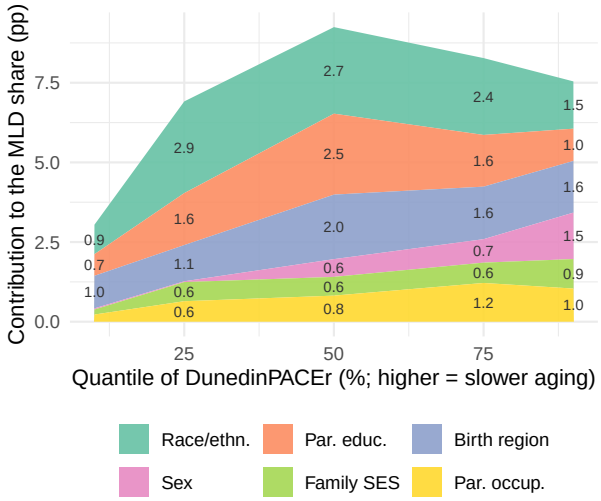
| | Circumstance | % of the MLD share (Shapley) | | | | | |
|---|---|---|---|---|---|---|---|
| Quantile | MLD share (SE) | Par. educ. | Race/eth. | Sex | Birthpl. | Fam. SES | Par. occ. |
| $q_{10}$ | 0.030 (0.006) | 22.3 | 30.4 | 1.2 | 33.6 | 5.2 | 7.4 |
| $q_{25}$ | 0.069 (0.009) | 23.6 | 41.6 | 0.4 | 16.4 | 8.7 | 9.3 |
| $q_{50}$ | 0.092 (0.011) | 27.5 | 29.4 | 6.0 | 21.9 | 6.4 | 8.9 |
| $q_{75}$ | 0.083 (0.011) | 19.6 | 29.2 | 8.9 | 19.9 | 7.7 | 14.7 |
| $q_{90}$ | 0.075 (0.010) | 13.4 | 19.7 | 19.3 | 21.5 | 12.3 | 13.8 |

*Notes.* At each quantile, the RIF of DunedinPACEr is adjusted age (single-year indicators) and blood-cell composition and recentred; "MLD share" is the survey-weighted share of its mean log deviation reproduced by the circumstances, and the last six columns Shapley–Owen-decompose that share across domains as relative values (shares sum to 100; the absolute values are in Figure 5, panel b). Low quantiles are the fastest agers, high quantiles the slowest. The MLD-share column is comparable *across quantiles* but not with the mean-based total of Table 5; see the text. Bootstrap standard errors in parentheses, from 200 replications that re-estimate the quantile, density, and decomposition.

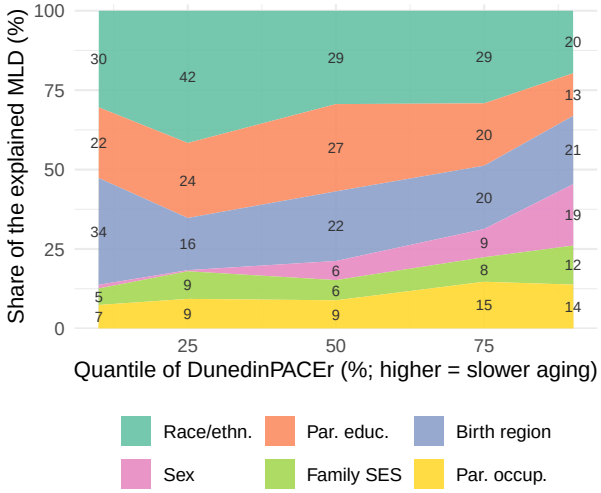
**Figure 5:** Distributional IOP profile




**(a)** How much: the circumstance share of the MLD




**(b)** Which circumstances: absolute Shapley values




**(c)** Which circumstances: relative Shapley values

*Notes.* Panel (a) plots the survey-weighted circumstance share of the mean log deviation of the adjusted RIF of DunedinPACEr by quantile, in percent, with a 95% confidence band (plus or minus 1.96 bootstrap standard errors). Panel (b) splits that share into the six domains' absolute Shapley values, stacked to the total, in percentage points of the share; each band is labelled with its contribution at that quantile. Panel (c) divides the same values by the share at each quantile, so that the bands stack to 100; labels in percent (shown from 3). Higher quantiles are slower agers. The vertical scale of panels (a) and (b) is comparable across quantiles but not with the mean-based total of Table 5.

# 6 Sensitivity analyses

## 6.1 Other epigenetic clocks

Table 3 and Figure 6 repeat the decomposition for our pace measure and five other epigenetic clocks, each inverted and residualized on age and blood-cell composition exactly like our main outcome DunedinPACE, so that in every row a higher value of the transformed marker is the better outcome. Two patterns emerge. First, the total amount of opportunity-linked inequality varies widely across assays, from 3.2% (PhenoAge) and 4.0% (Horvath) to 15.7% (GrimAge). Second, the *composition* splits along the stock–flow distinction of Section 3.1. For the two pace-of-aging measures, DunedinPACE and mPOA, race and ethnicity is the largest domain (34% and 44% of explained inequality), as in the main result. For three of the four biological-*age* clocks, *sex* is the largest domain—55% of GrimAge's explained inequality, 31% of Hannum's, and 26% of Horvath's—while race and ethnicity declines to between 12% and 22%. PhenoAge is the exception: once the cell mix is held fixed, its small circumstance share is carried by parental education (31%), the family's economic situation in childhood (24%), and region of birth, and sex accounts for only 2% of it. PhenoAge is trained on a clinical composite that includes white-cell counts (Appendix Table A4), so much of what distinguishes men and women in this clock runs through the composition of their blood. In the following, we describe additional potential sources of these differences. Overall, we read this as evidence that, if available, researchers should always rely on the most recent generation of clocks.

**Construction.** The differences in the level and decomposition of IOP could follow from how the markers were built (Appendix Table A4). A clock is a prediction model, and what it encodes is determined by its inputs and by the methylation sites it weights. GrimAge, for instance, takes the respondent's sex as an explicit input (Lu et al., 2019). This is not the case for DunedinPACE and mPOA were trained on the observed rate of physiological decline in a birth cohort followed over two decades (Belsky et al., 2020, 2022), so any sex gradient the decomposition finds in them is a gradient in measured aging. The rotation across clocks thus reflects what each algorithm is encoding, not a disagreement about the sources of inequality of opportunity.

**Reliability.** The *level* of the opportunity share depends on measurement reliability as much as on real inequality. Clocks differ markedly in test–retest reliability: the first-generation clocks and PhenoAge are among the noisiest, DunedinPACE and GrimAge among the most reliable (Higgins-Chen et al., 2022). Noise in the outcome inflates residual disper-

sion and so deflates the share that any fixed set of circumstances can explain. Part of the spread in the "MLD IOP" column of Table 3 is therefore a difference in measurement quality, and a larger share should not be read directly as more inequality of opportunity. Thus, more weight should be given to more reliable clocks.

**Stock versus flow.** Biological-*age* clocks estimate a level—an accumulated stock of depreciation—and are strongly correlated with chronological age, cohort, and hence differential survival. DunedinPACE estimates a rate and is nearly orthogonal to age (Appendix Figure A3). Age residualization therefore removes very different amounts of variance across markers, changing both the numerator and the denominator of the opportunity share, and the level clocks inherit circumstance gradients that operate through cumulative exposure and selective survival. A stock records depreciation that has already happened. A flow measures the rate at which health capital is being depreciated now, which is the object in Grossman-style models, the object a policy that slowed aging would change, and the reason the pace of aging forecasts the outcomes in Figure 1. For a question about opportunity, which concerns forward looking, the flow is therefore the preferred outcome.

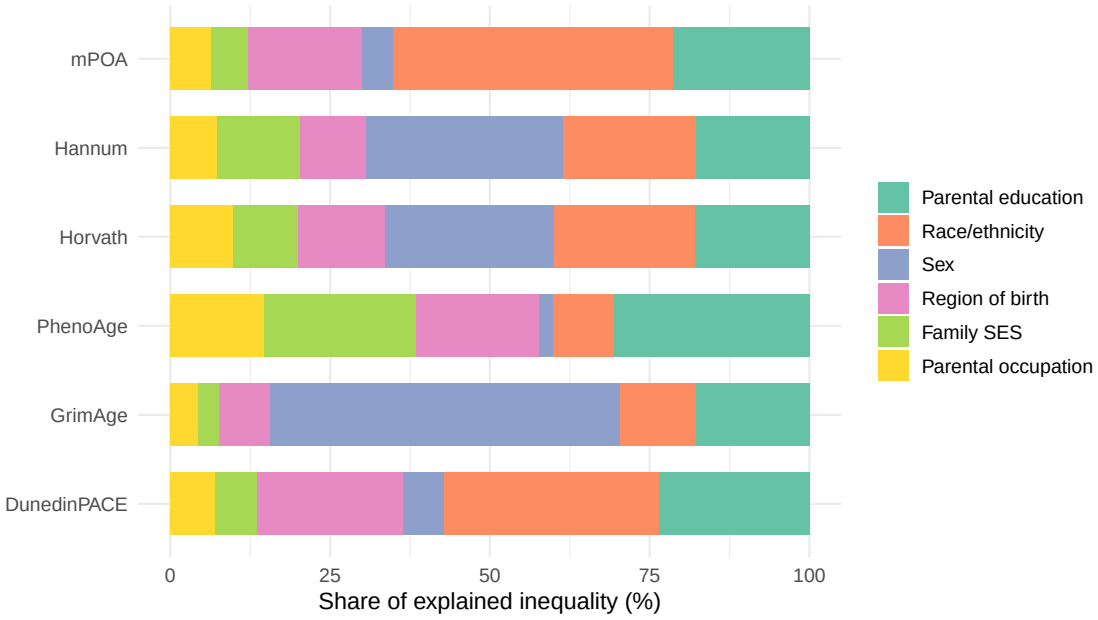
All three considerations—construction, reliability, and the stock–flow distinction—favor the pace of aging, so the cross-clock comparison strengthens rather than qualifies the main finding. This discussion can also help guiding the interpretation of clocks across contexts.

**Table 3:** Circumstance shares across epigenetic clocks

| Clock | MLD IOP | Par. educ. | Race/eth. | Sex | Birthpl. | Fam. SES | Par. occ. |
|---|---|---|---|---|---|---|---|
| DunedinPACE | 0.120 | 23.6 | 33.6 | 6.4 | 22.9 | 6.6 | 6.9 |
| GrimAge | 0.157 | 17.8 | 11.8 | 54.8 | 7.9 | 3.3 | 4.4 |
| PhenoAge | 0.032 | 30.5 | 9.7 | 2.2 | 19.3 | 23.8 | 14.6 |
| Horvath | 0.040 | 18.0 | 22.1 | 26.4 | 13.6 | 10.3 | 9.7 |
| Hannum | 0.047 | 17.8 | 20.8 | 30.8 | 10.4 | 12.9 | 7.3 |
| mPOA | 0.068 | 21.4 | 43.8 | 4.8 | 17.8 | 5.8 | 6.3 |

*Notes.* Each row is a separate survey-weighted Shapley–Owen decomposition on the MLD, like the main analysis. Every clock enters as its reciprocal and is then residualized on age and blood-cell composition and recentred, so that in each row a higher value is the better outcome: for the pace measures (DunedinPACE, mPOA) a lower rate is slower aging, and for the four biological-age clocks a lower level at a given calendar age is less accumulated depreciation. "MLD IOP" is the total MLD-based relative IOP; the remaining columns are each domain's % share of explained inequality.

**Figure 6:** Circumstance shares of explained inequality across six epigenetic clocks




*Notes.* Each clock is inverted and residualized on age and blood-cell composition as in Table 3, so that in every bar a higher value of the transformed marker is the better outcome. Bars are the six domains' Shapley–Owen shares of that clock's explained inequality, stacked to 100%.

## 6.2 The inequality index

Table 4 recomputes the ex-ante IOP and its Shapley decomposition under the variance (whose ratio is the familiar $R^2$) and the Gini. The total and the composition are nearly identical under the MLD and the variance (12.0% vs. 11.6%; race and ethnicity largest at 33.6% and 32.6%). Under the Gini, which weights the extreme tail less, the composition flattens: parental education (26.1%) moves ahead of race and ethnicity (21.9%), and the two childhood socioeconomic domains gain weight (family economic situation 10.1%, parental occupation 11.8%).

This shift reflects the distributional profile of Section 5.2. Because the race gradient is concentrated in the fast-aging tail, the tail-sensitive MLD and variance load more weight on it, while the flatter gradients of parental education and childhood economic circumstances gain weight under the Gini. The choice of index therefore matters for the *composition* of opportunity even when the total is stable, which is why we report all three in this sensitivity analysis. The orientation of the outcome matters little: the un-inverted DunedinPACE (lower panel of Table 4) implies the same conclusions.

**Table 4:** Ex-ante IOP: primary MLD index and robustness to the variance and Gini

| Inequality index | Total IOP | Par. educ. | Race/eth. | Sex | Birthpl. | Fam. SES | Par. occ. |
|---|---|---|---|---|---|---|---|
| *Primary outcome: DunedinPACEr (higher = slower aging)* | | | | | | | |
| MLD | 0.120 | 23.6 | 33.6 | 6.4 | 22.9 | 6.6 | 6.9 |
| Variance | 0.116 | 24.1 | 32.6 | 6.5 | 22.8 | 6.7 | 7.2 |
| Gini | 0.335 | 26.1 | 21.9 | 7.9 | 22.3 | 10.1 | 11.8 |
| *Comparison: DunedinPACE, un-inverted* | | | | | | | |
| MLD | 0.114 | 26.4 | 32.9 | 3.8 | 22.6 | 6.6 | 7.7 |
| Variance | 0.115 | 25.8 | 33.9 | 3.7 | 22.6 | 6.5 | 7.5 |
| Gini | 0.332 | 27.7 | 22.4 | 5.2 | 22.1 | 10.1 | 12.5 |

*Notes.* Each row re-estimates the ex-ante IOP and its Shapley–Owen decomposition using a different inequality index; the *MLD* (first row of each panel) is the paper's primary index, with the variance and Gini shown for robustness. The primary outcome is DunedinPACEr (upper panel); the lower panel repeats the exercise on the un-inverted DunedinPACE. Both are adjusted for age and blood-cell composition and recentred to their mean, and the variance total equals the baseline $R^2$. The MLD and variance are on a squared scale, the Gini total on the standard-deviation scale, so levels are comparable only within scale, whereas the domain shares (% columns) are comparable across indices. All regressions and all indices are survey-weighted.

# 7 Conclusion

Older Americans age at rates that differ systematically with circumstances they did not choose. Circumstances fixed in childhood—parental education and occupation, the family's economic situation, race and ethnicity, sex, and region of birth—account for 12.0% of the inequality in the pace of biological aging, and race and parental education account for more than half of that share. A distributional analysis shows: among the fastest agers, IOP in biological aging is mostly a matter of race and ethnicity and region of birth; among the slowest, sex matters as much as either.

Three implications follow. First, health-inequality measurement that stops at self-reported, diagnosed outcomes, and even common biomarkers understates the relevance of early-life circumstance, which is detectable in molecular aging before disease appears. Second, the weight of parental education is consistent with a developmental-origins channel in which the childhood socioeconomic environment shapes the trajectory of aging decades in advance. Third, because the composition of opportunity differs between the tails, the same average reduction in the pace of aging would have different distributional consequences depending on where in the distribution it is relevant: a reduction concentrated among the fastest agers would bear

mainly on the racial, ethnic, and regional components of IOP.

# Tables

**Table 5:** Shapley–Owen decomposition of DunedinPACEr inequality

| Circumstance domain | Shapley MLD | % | SE | 95% CI |
|---|---|---|---|---|
| Race/ethnicity | 0.040 | 33.65 | (4.03) | [25.76, 41.54] |
| Parental education | 0.028 | 23.55 | (3.58) | [16.54, 30.57] |
| Region of birth | 0.027 | 22.85 | (3.75) | [15.50, 30.20] |
| Parental occupation | 0.008 | 6.94 | (2.02) | [2.99, 10.89] |
| Family SES | 0.008 | 6.62 | (2.27) | [2.18, 11.06] |
| Sex | 0.008 | 6.39 | (2.34) | [1.81, 10.97] |
| **Total (relative IOP, MLD)** | 0.120 | | (0.012) | [0.097, 0.143] |

*Notes.* DunedinPACEr adjusted for age and blood-cell composition, $N = 3{,}866$. "Shapley MLD" is each domain's contribution to the total mean log deviation share, and the total row is the sum of that column; "%" rescales the same quantities to sum to 100. SE: bootstrap standard error of the share (1,000 nonparametric replications); 95% CI: the share plus or minus 1.96 SE.

# Appendix

**Table A1:** Sample composition and mean DunedinPACE by circumstance

| Circumstance | Category | Share (%) | Mean DunedinPACE |
|---|---|---|---|
| **All** | All | $N = 3{,}866$ | 1.020 |
| **Parental education** | | | |
| Father's education | No schooling | 2.42 | 1.092 |
| | Some grade school | 14.22 | 1.058 |
| | Grade school (8yr) | 13.65 | 1.038 |
| | Some high school | 9.49 | 1.023 |
| | High school | 26.75 | 0.992 |
| | Some college | 7.37 | 0.992 |
| | College+ | 13.10 | 0.968 |
| | Unknown | 13.00 | 1.070 |
| Mother's education | No schooling | 2.86 | 1.075 |
| | Some grade school | 12.27 | 1.069 |
| | Grade school (8yr) | 12.11 | 1.043 |
| | Some high school | 10.44 | 1.041 |
| | High school | 36.98 | 0.999 |
| | Some college | 8.69 | 0.994 |
| | College+ | 9.61 | 0.971 |
| | Unknown | 7.05 | 1.051 |
| **Race and ethnicity** | | | |
| Race | White | 83.31 | 1.008 |
| | Black | 10.40 | 1.094 |
| | Other | 6.16 | 1.051 |
| | Unknown | $n < 20$ | – |
| Ethnicity | Non-Hispanic | 91.00 | 1.016 |
| | Hispanic | 8.95 | 1.062 |
| | Unknown | $n < 20$ | – |
| **Sex** | | | |
| Sex | Male | 46.06 | 1.036 |
| | Female | 53.94 | 1.007 |
| **Region of birth** | | | |
| Region of birth | New England | 4.48 | 1.010 |
| | Mid Atlantic | 15.53 | 0.992 |
| | EN Central | 20.81 | 1.006 |
| | WN Central | 10.18 | 1.011 |
| | S Atlantic | 11.51 | 1.032 |
| | ES Central | 6.10 | 1.068 |
| | WS Central | 7.20 | 1.083 |

*Continued on next page*

**Table A1:** Sample composition and mean DunedinPACE by circumstance (continued)

| Circumstance | Category | Share (%) | Mean DunedinPACE |
|---|---|---|---|
| | Mountain | 3.75 | 1.011 |
| | Pacific | 7.81 | 0.978 |
| | US, division NA | 1.84 | 1.086 |
| | Foreign born | 10.79 | 1.038 |
| | Unknown | $n < 20$ | – |
| **Family socioeconomic status in childhood** | | | |
| Family finances, childhood | Pretty well off | 8.01 | 0.996 |
| | About average | 64.76 | 1.012 |
| | Poor | 25.88 | 1.048 |
| | Varied or unknown | 1.34 | 1.017 |
| Moved for financial reasons | Yes | 16.85 | 1.035 |
| | No | 82.81 | 1.016 |
| | Unknown | $n < 20$ | – |
| Help from relatives | Yes | 16.01 | 1.028 |
| | No | 82.62 | 1.019 |
| | Unknown | 1.37 | 0.996 |
| Father unemployed | Yes | 18.02 | 1.038 |
| | No | 74.32 | 1.011 |
| | Father never worked or absent | 6.69 | 1.074 |
| | Unknown | 0.97 | 1.039 |
| **Parental occupation** | | | |
| Father's occupation (age 16) | Managerial/professional | 20.07 | 0.983 |
| | Sales | 8.68 | 1.001 |
| | Clerical | 3.00 | 1.002 |
| | Service | 4.96 | 1.016 |
| | Manual/operators | 49.38 | 1.030 |
| | Armed forces | 1.81 | 1.011 |
| | Unknown or no father | 12.10 | 1.060 |
| Mother worked outside home | All of the time | 19.12 | 1.029 |
| | Some of the time | 40.59 | 1.015 |
| | Not at all | 38.96 | 1.019 |
| | No mother or unknown | 1.33 | 1.067 |

*Notes.* HRS 2016 Venous Blood Study, $N = 3{,}866$. DunedinPACE is the DNAm pace of aging (Belsky et al., 2022). Both columns are *survey-weighted*: shares of the weighted population and weighted means of the raw (unadjusted) outcome. The family socioeconomic and parental occupation items are retrospective reports on the respondent's childhood (birth to age 16; father's occupation at 16). "Unknown": the circumstance is not recorded (don't know, refused, not asked or missing); kept as a category rather than dropped. Categories with fewer than twenty respondents are suppressed ($n < 20$).

**Table A2:** Fine decomposition: individual circumstances

| Circumstance | Shapley MLD | % of explained | SE | 95% CI |
|---|---|---|---|---|
| Race | 0.036 | 29.79 | (3.98) | [21.99, 37.58] |
| Region of birth | 0.026 | 21.85 | (3.71) | [14.58, 29.13] |
| Father's education | 0.020 | 16.81 | (3.14) | [10.65, 22.96] |
| Mother's education | 0.011 | 9.44 | (2.27) | [4.99, 13.89] |
| Sex | 0.008 | 6.47 | (2.35) | [1.86, 11.09] |
| Father's occupation (age 16) | 0.006 | 5.33 | (1.77) | [1.87, 8.79] |
| Father unemployed | 0.003 | 2.67 | (1.26) | [0.20, 5.14] |
| Hispanic ethnicity | 0.003 | 2.66 | (1.15) | [0.40, 4.91] |
| Family finances, childhood | 0.002 | 1.99 | (1.04) | [-0.05, 4.02] |
| Help from relatives | 0.002 | 1.45 | (1.38) | [-1.27, 4.16] |
| Mother worked outside home | 0.001 | 0.99 | (0.84) | [-0.66, 2.63] |
| Moved for financial reasons | 0.001 | 0.56 | (0.59) | [-0.60, 1.72] |

*Notes.* DunedinPACEr adjusted for age and blood-cell composition, $N$ = 3,866. Shapley–Owen decomposition treating each circumstance variable as its own group. Shares of explained inequality sum to 100. SE: bootstrap standard error of the share (1,000 replications); 95% CI: the share plus or minus 1.96 SE.

## A.1 The gradients behind the distributional rotation

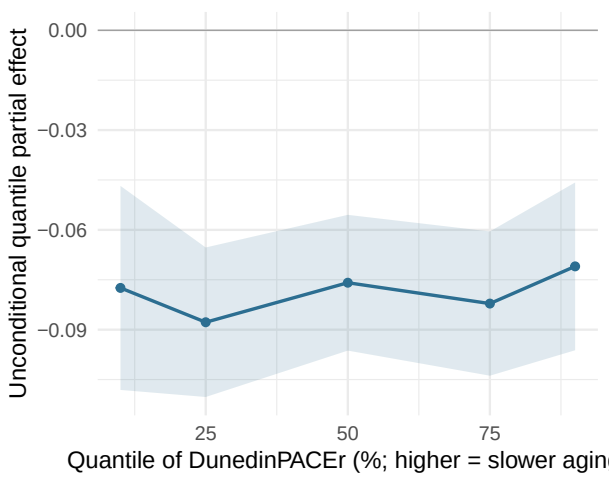
Section 5.2 decomposes the inequality of the pace of aging at each quantile and finds that the composition of opportunity rotates across the distribution. This appendix reports the circumstance-by-circumstance gradients that produce the rotation: RIF unconditional-quantile regressions of DunedinPACEr, which recover the effect of a circumstance on the *unconditional* quantiles of the outcome distribution rather than on the conditional quantiles a quantile regression would deliver. Because the outcome is DunedinPACEr, a positive coefficient means slower aging and higher quantiles are slower agers. Table A3 reports the coefficients at the tenth, fiftieth and ninetieth percentiles; Figure A1 traces them across the whole distribution.

**Table A3:** RIF unconditional-quantile regressions of DunedinPACEr

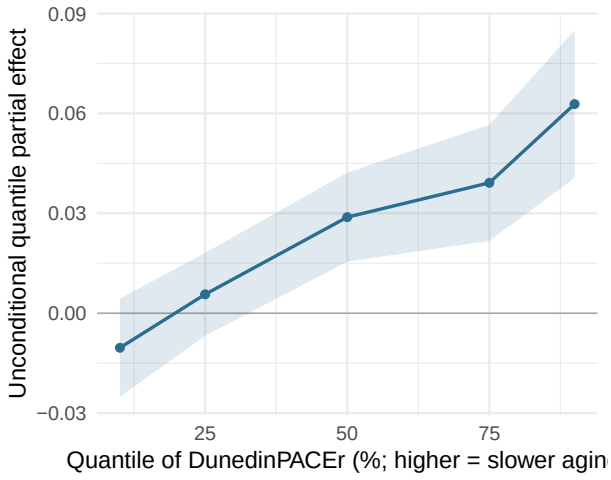
| Circumstance | $q_{10}$ | $q_{50}$ | $q_{90}$ |
|---|---|---|---|
| Black (vs. White) | -0.077*** | -0.076*** | -0.071*** |
| | (0.016) | (0.010) | (0.013) |
| Hispanic | -0.005 | -0.023 | -0.056*** |
| | (0.018) | (0.014) | (0.020) |
| Father college+ (vs. HS) | -0.010 | 0.017 | 0.038 |
| | (0.013) | (0.014) | (0.028) |
| Father no schooling (vs. HS) | -0.090*** | -0.035 | -0.013 |
| | (0.034) | (0.024) | (0.025) |
| Female (vs. male) | -0.010 | 0.029*** | 0.063*** |
| | (0.008) | (0.007) | (0.011) |
| Family poor (vs. average) | 0.012 | -0.005 | -0.017 |
| | (0.010) | (0.009) | (0.014) |
| Father managerial/prof. (vs. manual) | 0.013 | 0.017* | -0.016 |
| | (0.010) | (0.011) | (0.019) |

*Notes.* Survey-weighted RIF unconditional-quantile regressions at the 10th, 50th, and 90th percentiles; all six circumstance sets, single-year age indicators, and the five estimated blood-cell shares included. Heteroskedasticity-robust standard errors in parentheses. Reference groups: White, non-Hispanic, male, high-school-educated father. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

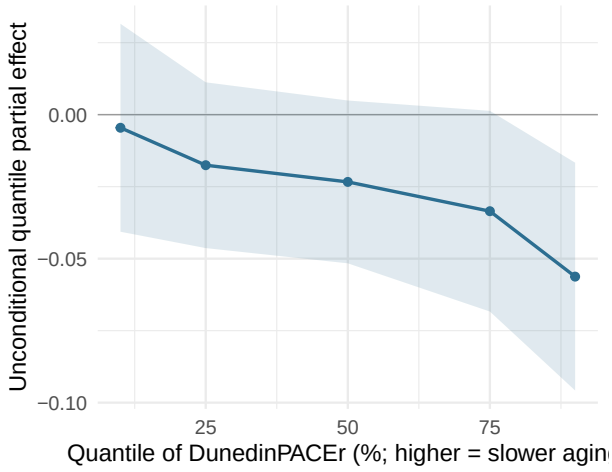
**Figure A1:** Unconditional quantile partial effects of selected circumstances on DunedinPACEr



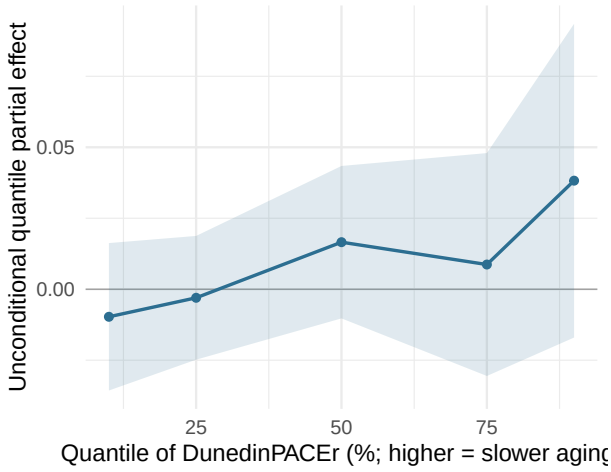

**(a)** Black (vs. White)




**(b)** Female (vs. male)




**(c)** Hispanic




**(d)** Father college+ (vs. high school)

*Notes.* Survey-weighted RIF regressions, with 95% confidence bands. A positive effect is slower aging; higher quantiles are slower agers.

**Table A4:** Epigenetic clocks used in the analysis: construction and interpretation

| Clock | Generation | Construction: tissue(s), CpG sites, training target and covariates | Scale and interpretation | Source |
|---|---|---|---|---|
| Horvath | First | Pan-tissue: whole blood plus ∼50 other tissues and cell types (∼8,000 samples). Elastic net on 353 CpGs. Target: *chronological age*. No biological or demographic covariates; sex-associated CpGs enter freely, so the predictor carries some sex signal. | DNAm age in *years*; "age acceleration" is the residual from calendar age. Level measure. | Horvath (2013) |
| Hannum | First | Whole blood only. Elastic net on 71 CpGs. Target: *chronological age*; the aging model was fit adjusting for sex and BMI, and sex-linked methylation is loaded into the predictor. | DNAm age in *years*; age acceleration is the residual from calendar age. Level measure. | Hannum et al. (2013) |
| PhenoAge | Second | Whole blood. Elastic net on 513 CpGs. Target: *Phenotypic Age*, a composite of nine clinical biomarkers (e.g. CRP, glucose, albumin, white-cell counts) plus chronological age. | Biological age in *years*; higher values imply a greater accumulated stock of depreciation. Level measure. | Levine et al. (2018) |
| GrimAge | Second | Whole blood. Built from DNAm surrogates for seven plasma proteins and DNAm smoking pack-years (∼1,030 CpGs), *plus chronological age and sex as explicit model inputs*. Target: *time-to-death* (mortality). | Biological age in *years*; a unit is a unit of predicted lifespan lost. Level measure; sex is built in by construction. | Lu et al. (2019) |
| DunedinPACE | Pace (2nd-gen) | Whole blood. Elastic net on 173 CpGs. Target: the longitudinal *rate* of decline across 19 organ-system biomarkers in a birth cohort followed from age 26 to 45. No demographic covariates. | A *rate*: 1.0 = one biological year of aging per calendar year; 1.05 = aging 5% faster. Flow measure; near-orthogonal to chronological age. | Belsky et al. (2022) |
| mPOA | Pace (2nd-gen) | Whole blood. Elastic net on 46 CpGs. Target: single-time-point *Pace of Aging* (19 biomarkers, cohort ages 26–38); the predecessor of DunedinPACE. No demographic covariates. | A *rate* of biological aging, interpreted as for DunedinPACE. Flow measure. | Belsky et al. (2020) |

*Notes.* "Generation" follows the standard taxonomy: first-generation clocks are trained on chronological age, second-generation clocks on health/mortality outcomes, and the pace measures on the longitudinal rate of aging. All clocks are penalized (elastic-net) predictors of DNA methylation. CpG counts are the number of methylation sites in each published predictor. The "covariates" item highlights why the circumstance decompositions differ across clocks (Table 3): GrimAge includes sex as an explicit model input and the first-generation clocks load sex-associated methylation, so their inequality is sex-dominated, whereas the pace measures (DunedinPACE, mPOA) contain no demographic covariates. DunedinPACE is the primary outcome of the paper.

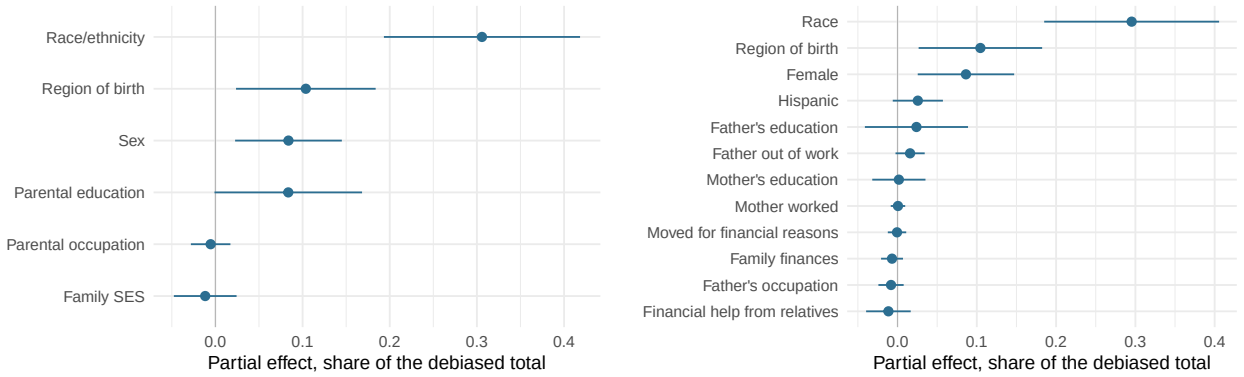
## A.2 Which circumstances does the debiased estimator use?

Section 5 decomposes the total with Shapley–Owen values, which divide the explained inequality additively among the domains. The debiased literature reports a different quantity: the *partial effect* of a circumstance, $\kappa_m = \text{IOP} - \text{IOP}_{(-m)}$, obtained by re-estimating the whole orthogonal moment with that circumstance (or that domain) removed (Escanciano and Terschuur, 2023). Figure A2 reports it for our MLD estimator, as a share of the total, with 95% confidence intervals from the same influence function.

The two decompositions answer different questions and should not be added to one another. A partial effect is a *leave-one-out* quantity: when circumstances are correlated—and here they are, since parental education, race and region of birth move together—the inequality that two of them explain jointly is attributed to neither alone, so the partial effects do not sum to the total (here the six domain effects sum to 56% of it) and each understates a correlated circumstance's contribution. The Shapley–Owen values in Table 5 average over orderings precisely to avoid that, and remain the paper's decomposition. What the partial effects add is a robustness check computed inside the debiased estimator itself: if a domain that carries a large Shapley share also moves the debiased total when it is removed, the ranking is not an artefact of the plug-in.

Race and ethnicity passes this check clearly: removing it lowers the debiased total by 31%. Removing region of birth lowers it by 10% and removing sex by 8%, both significant at the 5% level. Parental education (8%) falls just short of conventional significance, and the two childhood socioeconomic domains do not move the total by a detectable amount. For parental education this is the leave-one-out property at work rather than evidence that it does not matter: the circumstances that remain—the family's economic situation, the parents' occupations, race and region of birth—carry most of what parental education predicts, so the first stage loses little when it is withheld. The small negative point estimates for the two childhood socioeconomic domains are of the same kind, sampling noise around zero in a comparison of two cross-fitted first stages.

**Figure A2:** Debiased partial effects of the circumstances




**(a)** By domain **(b)** By circumstance

*Notes.* Partial effects $\kappa_m = \text{IOP} - \text{IOP}_{(-m)}$ of the debiased MLD estimator, expressed as a share of the estimated total, with 95% intervals from the influence function of the difference (exact for this same-sample comparison). Adjusted DunedinPACEr, $N = 3{,}866$, survey-weighted, first-stage learner as in Table 1. These are leave-one-out effects, not Shapley values: they do not sum to the total, and correlated circumstances share credit that is attributed to neither of them alone.

## A.3 Why the debiased MLD is biased downward

This appendix explains, step by step, the statement in Section 4.2 that the debiased estimator of the MLD errs downward by an amount that is proportional to how badly the first stage fits.

**Set-up.** Let $\gamma_0(C) = \mathbb{E}[y \mid C]$ be the type-mean function, which the first stage estimates by $\hat{\gamma}$. The target is the MLD of $\gamma_0(C)$,

$$\theta_0 = \ln \mathbb{E}[\gamma_0(C)] - \mathbb{E}[\ln \gamma_0(C)].$$

A plug-in estimator inserts $\hat{\gamma}$ for $\gamma_0$. Its error depends on the error of $\hat{\gamma}$ *to first order*, so any bias of the first stage passes into the result. The debiased estimator replaces $\ln \gamma_0(C)$ by the *score*

$$\psi(y, C; \gamma) = \ln \gamma(C) + \frac{y - \gamma(C)}{\gamma(C)}, \qquad \hat{\theta} = \ln \bar{y} - \frac{1}{n} \sum_i \psi(y_i, C_i; \hat{\gamma}),$$

where $\bar{y}$ estimates $\mathbb{E}[\gamma_0] = \mathbb{E}[y]$. The first stage $\hat{\gamma}$ is cross-fitted: the value used for observation $i$ is estimated on folds that do not contain $i$.

**Step 1: why the correction removes the first-order error.** Take the expectation of $\psi$ at a candidate function $\gamma$, using $\mathbb{E}[y \mid C] = \gamma_0$:

$$\mathbb{E}[\psi(y, C; \gamma) \mid C] = \ln \gamma(C) + \frac{\gamma_0(C) - \gamma(C)}{\gamma(C)}.$$

Move $\gamma$ away from $\gamma_0$ by a small amount $h$ and differentiate at $h = 0$: the first term contributes $h/\gamma_0$, the second $-h/\gamma_0$, so the two cancel. The expected score is therefore *flat* at the truth, a property called Neyman orthogonality, and a small first-stage error changes it only at second order. Without the correction term the flat spot would not exist and the first-stage error would enter linearly.

**Step 2: the exact size of the second-order term.** Write the first-stage error in relative terms, $\hat{\gamma} = \gamma_0(1 + r)$, with $r = (\hat{\gamma} - \gamma_0)/\gamma_0$. Inserting this in the expression above gives, for an observation whose first-stage value was estimated on other folds,

$$\mathbb{E}[\psi \mid C, \hat{\gamma}] = \ln \gamma_0 + \ln(1 + r) - \frac{r}{1 + r} = \ln \gamma_0 + f(r), \qquad f(r) = \ln(1 + r) - \frac{r}{1 + r}.$$

The average of $\ln \gamma_0(C_i)$ is exactly what the target subtracts, so

$$\mathbb{E}[\hat{\theta}] - \theta_0 = -\,\mathbb{E}\big[f(r)\big] + \underbrace{\big(\mathbb{E}\ln \bar{y} - \ln \mathbb{E}\, y\big)}_{\text{Jensen, } \le 0,\ \approx -\mathrm{MLD}(y)/n} .$$

The last term is the usual small-sample bias of the log of a mean. In our application it is below one per cent of the estimate and is ignored below.

**Step 3: the sign.** $f(0) = 0$ and $f'(r) = r/(1+r)^2$, which is negative for $r < 0$ and positive for $r > 0$. So $f$ has its minimum at zero and $f(r) \ge 0$ for every admissible $r$. The remainder can therefore only pull the estimate *down*, whether the first stage over- or under-predicts. A Taylor expansion gives $f(r) = r^2/2 - \frac{2}{3}r^3 + \ldots$, hence

$$\mathbb{E}[\hat{\theta}] - \theta_0 \;\approx\; -\tfrac{1}{2}\,\mathbb{E}[r^2] = -\tfrac{1}{2}\,\mathbb{E}\left[\frac{(\hat{\gamma} - \gamma_0)^2}{\gamma_0^2}\right].$$

The bias is *quadratic* in the first-stage error: halving the error cuts the bias to a quarter.

**Step 4: proportional to the fit of the first stage.** Two approximations, both exact to second order, express the bias in units of the target. First, a second-order expansion of the logarithm gives $\theta_0 = \ln \mathbb{E}\gamma_0 - \mathbb{E}\ln \gamma_0 \approx \mathrm{Var}(\gamma_0)/(2\mu^2)$ with $\mu = \mathbb{E}[y]$. Second, $\mathbb{E}[r^2] \approx \mathrm{MSE}(\hat{\gamma})/\mu^2$ with $\mathrm{MSE}(\hat{\gamma}) = \mathbb{E}[(\hat{\gamma} - \gamma_0)^2]$. Dividing,

$$\frac{\mathbb{E}[\hat{\theta}] - \theta_0}{\theta_0} \;\approx\; -\frac{\mathrm{MSE}(\hat{\gamma})}{\mathrm{Var}(\gamma_0)} \;=\; -\big(1 - R^2_{\hat{\gamma}\to\gamma_0}\big).$$

The relative bias equals minus the share of the variation in the type means that the first stage does not reproduce. A first stage that captures 90% of $\mathrm{Var}(\gamma_0)$ leaves a 10% downward bias; a poor first stage that captures half of it leaves 50%. Two things follow. (i) The denominator is small when circumstances explain little of the outcome, as here, so a modest absolute error is a sizable relative one. (ii) MSE has two parts, estimation variance, which shrinks with the sample size and is larger the more parameters are estimated, and approximation error, which does not: a learner that cannot represent interactions among circumstances carries the same misspecification error as OLS however large the sample. Orthogonality removes the first-order effect of estimation error; it cannot fix a misspecified limit.

**Simulation.** We check the algebra on simulated data in which the truth is known. There are twenty binary circumstances, ten of which matter with equal weight, and $\mathbb{E}[y \mid C] = 1 + s\sum_{j\le 10} C_j$ with outcome $y = \gamma_0(C)\,\varepsilon$ and log-normal noise of mean one, so the noise is

heteroskedastic. The scale $s$ is set so that circumstances explain a relative MLD of 0.12, the size of our main result, and the true $\theta_0 = 0.0245$. The sample has $n = 3{,}000$ observations, the estimator is the cross-fitted debiased MLD above with five folds, and each row is the average of 1,000 replications. We vary only the quality of the first stage, a cross-fitted OLS that uses a given set of the circumstances: too few relevant ones (approximation error), exactly the relevant ones, or the relevant ones plus irrelevant ones (estimation variance).

**Table A5:** The realized bias of the debiased MLD equals the remainder

| First stage uses | Realized bias (%) | Exact remainder (%) | Quadratic approx. (%) | Rule of thumb $-\text{MSE}/\text{Var}(\gamma_0)$ (%) |
|---|---|---|---|---|
| 2 of 10 relevant | −82.8 | −82.1 | −112.3 | −81.1 |
| 5 of 10 relevant | −54.9 | −54.4 | −65.9 | −52.5 |
| 8 of 10 relevant | −26.1 | −25.2 | −26.8 | −23.7 |
| all 10 relevant (correct) | −5.2 | −4.9 | −4.8 | −4.6 |
| 10 relevant + 5 irrelevant | −7.9 | −7.2 | −7.1 | −6.6 |
| 10 relevant + 10 irrelevant | −9.9 | −9.8 | −9.5 | −8.9 |

All entries are percent of the true $\theta_0$. Realized bias: mean of $\hat{\theta}$ over 1,000 replications minus $\theta_0$. Exact remainder: $-$mean $f(r)$ with the realized relative first-stage errors. Quadratic approximation: $-$mean $r^2/2$. Rule of thumb: $-\text{MSE}(\hat{\gamma})/\text{Var}(\gamma_0)$. $n = 3{,}000$, relative MLD 0.12. Script: `R/12_remainder_sim.R`.

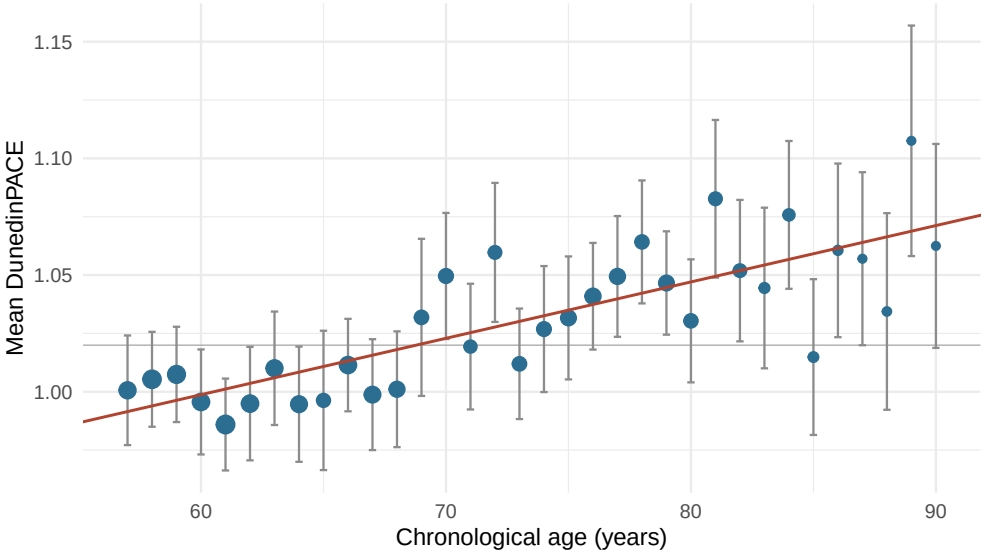
Table A5 shows three things. First, the realized bias is negative in every row, as Step 3 requires. Second, it matches the exact remainder to within about one percentage point, the gap being the Jensen term of Step 2; the algebra accounts for the whole bias, so the estimator is implemented correctly and the bias is a result of the method, not a bug. Third, the bias tracks the quality of the first stage as Step 4 says. It is −83% when the first stage uses two of the ten relevant circumstances, shrinks as it uses more, and is smallest (−5%) when the first stage contains exactly the relevant circumstances; adding irrelevant circumstances makes it larger again because the extra estimated coefficients raise the estimation variance. The quadratic approximation and the rule of thumb are accurate when the first stage is decent and overstate or understate the bias only when the errors are so large (the first rows) that the cubic term of $f$ matters.

**What this means for the paper.** The debiased estimate is a floor for the share of inequality that circumstances explain, and the shortfall is larger the worse the first stage predicts the type means. It is not evidence that the main OLS estimate is too high; the OLS estimator errs in the opposite direction because it fits the noise in the sample, and the

two estimates bracket the share of interest (Section 4.2). The analytic standard errors of the debiased estimate also ignore the sampling variation of the remainder, so they are a lower bound on the true spread.

## A.4 DunedinPACE and chronological age

**Figure A3:** DunedinPACE against chronological age




*Notes.* Survey-weighted mean of DunedinPACE within single-year age bins, with 95% intervals and a weighted linear fit. Bins with fewer than twenty observations are not plotted.

# References


Douglas Almond and Janet Currie. Killing me softly: The fetal origins hypothesis. *Journal of Economic Perspectives*, 25(3):153–172, 2011.

Douglas Almond, Janet Currie, and Valentina Duque. Childhood circumstances and adult outcomes: Act II. *Journal of Economic Literature*, 56(4):1360–1446, 2018.

Daniel W. Belsky, Avshalom Caspi, Louise Arseneault, Andrea Baccarelli, David L. Corcoran, Xu Gao, Eilis Hannon, Honalee L. Harrington, Line J. Rasmussen, Renate Houts, Kim Huffman, William E. Kraus, Dayoon Kwon, Jonathan Mill, Carl F. Pieper, Joseph A. Prinz, Richie Poulton, Joel Schwartz, Karen Sugden, Pantel Vokonas, Benjamin S. Williams, and Terrie E. Moffitt. Quantification of the pace of biological aging in humans through a blood test, the DunedinPoAm DNA methylation algorithm. *eLife*, 9: e54870, 2020.

Daniel W. Belsky, Avshalom Caspi, David L. Corcoran, Karen Sugden, Richie Poulton, Louise Arseneault, Andrea Baccarelli, and Terrie E. Moffitt. DunedinPACE, a DNA methylation biomarker of the pace of aging. *eLife*, 11:e73420, 2022.

François Bourguignon, Francisco H. G. Ferreira, and Marta Menéndez. Inequality of opportunity in Brazil. *Review of Income and Wealth*, 53(4):585–618, 2007.

Paolo Brunori, Vito Peragine, and Laura Serlenga. Upward and downward bias when measuring inequality of opportunity. *Social Choice and Welfare*, 52(4):635–661, 2019. doi: 10.1007/s00355-018-1165-x.

Paolo Brunori, Paul Hufe, and Daniel Gerszon Mahler. The roots of inequality: Estimating inequality of opportunity from regression trees and forests. *The Scandinavian Journal of Economics*, 125(4):900–932, 2023. doi: 10.1111/sjoe.12530.

Giacomo Cavalli and Edith Heard. Advances in epigenetics link genetics to the environment and disease. *Nature*, 571(7766):489–499, 2019. doi: 10.1038/s41586-019-1411-0.

Daniele Checchi and Vito Peragine. Inequality of opportunity in Italy. *Journal of Economic Inequality*, 8(4):429–450, 2010.

Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018. doi: 10.1111/ectj.12097.

Eileen M. Crimmins, Bharat Thyagarajan, Morgan E. Levine, David R. Weir, and Jessica Faul. Associations of age, sex, race/ethnicity, and education with 13 epigenetic clocks in a nationally representative u.s. sample: The Health and Retirement Study. *The Journals of Gerontology: Series A*, 76(6):1117–1123, 2021. doi: 10.1093/gerona/glab016.

Carl-Johan Dalgaard and Holger Strulik. Optimal aging and death: Understanding the Preston curve. *Journal of the European Economic Association*, 12(3):672–701, 2014. doi: 10.1111/jeea.12071.

Apostolos Davillas and Andrew M. Jones. Ex ante inequality of opportunity in health, decomposition and distributional analysis of biomarkers. *Journal of Health Economics*, 69:102251, 2020. doi: 10.1016/j.jhealeco.2019.102251.

Apostolos Davillas and Andrew M. Jones. Biological age and predicting future health care utilisation. *Journal of Health Economics*, 99:102956, 2025. ISSN 0167-6296. doi: https://doi.org/10.1016/j.jhealeco.2024.102956. URL `https://www.sciencedirect.com/science/article/pii/S0167629624001012`.

Juan Carlos Escanciano and Joël Robert Terschuur. Machine learning inference on inequality of opportunity. Preprint arXiv:2206.05235, arXiv, 2023. URL `https://arxiv.org/abs/2206.05235`. Later versions circulated as "Debiased Machine Learning U-Statistics"; the estimator is implemented in the authors' `ineqopp` package for R.

Jessica D. Faul, Jung Ki Kim, Morgan E. Levine, Bharat Thyagarajan, David R. Weir, and Eileen M. Crimmins. Epigenetic-based age acceleration in a representative sample of older Americans: Associations with aging-related morbidity and mortality. *Proceedings of the National Academy of Sciences*, 120(9):e2215840120, 2023. doi: 10.1073/pnas.2215840120.

Francisco H. G. Ferreira and Jérémie Gignoux. The measurement of inequality of opportunity: Theory and an application to Latin America. *Review of Income and Wealth*, 57(4): 622–657, 2011.

Giovanni Fiorito, Silvia Polidoro, Pierre-Antoine Dugué, Mika Kivimaki, et al. Social adversity and epigenetic aging: A multi-cohort study on socioeconomic differences in peripheral blood DNA methylation. *Scientific Reports*, 7:16266, 2017. doi: 10.1038/s41598-017-16391-5.

Sergio Firpo, Nicole M. Fortin, and Thomas Lemieux. Unconditional quantile regressions. *Econometrica*, 77(3):953–973, 2009.

Marc Fleurbaey and Vito Peragine. Ex ante versus ex post equality of opportunity. *Economica*, 80(317):118–130, 2013. doi: 10.1111/j.1468-0335.2012.00941.x.

Marc Fleurbaey and Erik Schokkaert. Unfair inequalities in health and health care. *Journal of Health Economics*, 28(1):73–90, 2009.

James E. Foster and Artyom A. Shneyerov. Path independent inequality measures. *Journal of Economic Theory*, 91(2):199–222, 2000. doi: 10.1006/jeth.1999.2565.

Titus J Galama and Hans van Kippersluis. A theory of socio-economic disparities in health over the life cycle. *The Economic Journal*, 129(617):338–374, 01 2019. ISSN 0013-0133. doi: 10.1111/ecoj.12577. URL `https://doi.org/10.1111/ecoj.12577`.

Arline T. Geronimus, Margaret Hicken, Danya Keene, and John Bound. "weathering" and age patterns of allostatic load scores among blacks and whites in the United States. *American Journal of Public Health*, 96(5):826–833, 2006. doi: 10.2105/AJPH.2004.060749.

Daniel Graeber, Johannes König, and Viola Hilbert. Inequality of opportunity in wealth: Levels, trends and drivers in germany. *Oxford Bulletin of Economics and Statistics*, n/a

(n/a). doi: https://doi.org/10.1111/obes.70112. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/obes.70112.

Michael Grossman. On the concept of health capital and the demand for health. *Journal of Political Economy*, 80(2):223–255, 1972.

Gregory Hannum, Justin Guinney, Ling Zhao, Li Zhang, Guy Hughes, SriniVas Sadda, Brandy Klotzle, Marina Bibikova, Jian-Bing Fan, Yuan Gao, Rob Deconde, Menzies Chen, Indika Rajapakse, Stephen Friend, Trey Ideker, and Kang Zhang. Genome-wide methylation profiles reveal quantitative views of human aging rates. *Molecular Cell*, 49(2):359–367, Jan 2013. doi: 10.1016/j.molcel.2012.10.016. URL https://pubmed.ncbi.nlm.nih.gov/23177740/.

Clyde Hertzman. The biological embedding of early experience and its effects on health in adulthood. *Annals of the New York Academy of Sciences*, 896(1):85–95, 1999. doi: 10.1111/j.1749-6632.1999.tb08107.x.

Albert T. Higgins-Chen, Kyra L. Thrush, Yunzhang Wang, Christopher J. Minteer, Pei-Lun Kuo, Meng Wang, Peter Niimi, Gabriel Sturm, Jue Lin, Ann Zenobia Moore, Stefania Bandinelli, Christiaan H. Vinkers, Eric Vermetten, Bart P. F. Rutten, Elbert Geuze, Cynthia Okhuijsen-Pfeifer, Marte Z. van der Horst, Stefan Schreiber, Vadim N. Gladyshev, Luigi Ferrucci, Robert L. Grossman, Martin Picard, and Morgan E. Levine. A computational solution for bolstering reliability of epigenetic clocks: Implications for clinical trials and longitudinal tracking. *Nature Aging*, 2(7):644–661, 2022.

Steve Horvath. DNA methylation age of human tissues and cell types. *Genome Biology*, 14 (10):R115, 2013.

Frank Huettner and Marco Sunder. Axiomatic arguments for decomposing goodness of fit according to Shapley and Owen values. *Electronic Journal of Statistics*, 6:1239–1250, 2012.

Paul Hufe, Martyna Kobus, Andreas Peichl, and Paul Schüle. Multidimensional equality of opportunity in the united states. *European Economic Review*, 180:105162, 2025. ISSN 0014-2921. doi: https://doi.org/10.1016/j.euroecorev.2025.105162. URL https://www.sciencedirect.com/science/article/pii/S0014292125002120.

Amanda Hughes, Melissa Smart, Tyler Gorrie-Stone, Eilis Hannon, et al. Socioeconomic position and DNA methylation age acceleration across the life course. *American Journal of Epidemiology*, 187(11):2346–2354, 2018. doi: 10.1093/aje/kwy155.

Nancy Krieger. Embodiment: A conceptual glossary for epidemiology. *Journal of Epidemiology and Community Health*, 59(5):350–355, 2005. doi: 10.1136/jech.2004.024562.

Morgan E. Levine, Ake T. Lu, Austin Quach, Brian H. Chen, Themistocles L. Assimes, Stefania Bandinelli, Lifang Hou, Andrea A. Baccarelli, James D. Stewart, Yun Li, Eric A. Whitsel, James G. Wilson, Alexander P. Reiner, Abraham Aviv, Kurt Lohman, Yongmei Liu, Luigi Ferrucci, and Steve Horvath. An epigenetic biomarker of aging for lifespan and healthspan. *Aging (Albany NY)*, 10(4):573–591, April 2018. doi: 10.18632/aging.101414. URL `https://doi.org/10.18632/aging.101414`.

Ake Lu, Austin Quach, James G. Wilson, Alexander P. Reiner, Abraham Aviv, Kumar Raj, Lifang Hou, Andrea A. Baccarelli, Yun Li, Eric A. Whitsel, and Steve Horvath. Dna methylation grimage strongly predicts lifespan and healthspan. *Aging (Albany NY)*, 11(2): 303–327, 2019. doi: 10.18632/aging.101684. URL `https://www.aging-us.com/article/101684`.

Giorgia Menta, Pietro Biroli, Divya Mehta, Conchita D'Ambrosio, and Deborah A. Cobb-Clark. Aggregating epigenetic clocks to study human capital formation. IZA Discussion Paper 18114, Institute of Labor Economics (IZA), 2025.

Laurel Raffington, Daniel W. Belsky, Meeraj Kothari, Margherita Malanchini, et al. Socioeconomic disadvantage and the pace of biological aging in children. *Pediatrics*, 147(6): e2020024406, 2021. doi: 10.1542/peds.2020-024406.

Xavier Ramos and Dirk Van de gaer. Approaches to inequality of opportunity: Principles, measures and evidence. *Journal of Economic Surveys*, 30(5):855–883, 2016.

John E. Roemer. *Equality of Opportunity*. Harvard University Press, Cambridge, MA, 1998.

John E. Roemer and Alain Trannoy. Equality of opportunity: Theory and measurement. *Journal of Economic Literature*, 54(4):1288–1332, 2016. doi: 10.1257/jel.20151206.

Pedro Rosa Dias. Inequality of opportunity in health: Evidence from a UK cohort study. *Health Economics*, 18(9):1057–1074, 2009. doi: 10.1002/hec.1535.

Lauren L. Schmitz and Valentina Duque. In utero exposure to the Great Depression is reflected in late-life epigenetic aging signatures. *Proceedings of the National Academy of Sciences*, 119(46):e2208530119, 2022. doi: 10.1073/pnas.2208530119.

Lauren L. Schmitz, Wei Zhao, Scott M. Ratliff, Julia Goodwin, et al. The socioeconomic gradient in epigenetic ageing clocks: Evidence from the Multi-Ethnic Study of Atherosclerosis and the Health and Retirement Study. *Epigenetics*, 17(6):589–611, 2022. doi: 10.1080/15592294.2021.1939479.

Anthony F. Shorrocks. The class of additively decomposable inequality measures. *Econometrica*, 48(3):613–625, 1980. doi: 10.2307/1913126.

Anthony F. Shorrocks. Decomposition procedures for distributional analysis: A unified framework based on the Shapley value. *Journal of Economic Inequality*, 11(1):99–126, 2013.

Amanda Sonnega, Jessica D. Faul, Mary Beth Ofstedal, Kenneth M. Langa, John W.R. Phillips, and David R. Weir. Cohort profile: the health and retirement study (HRS). *International Journal of Epidemiology*, 43(2):576–585, 2014. doi: 10.1093/ije/dyu067.

Joël Robert Terschuur. Educational inequality of opportunity and mobility in Europe. Preprint arXiv:2212.02407, arXiv, 2023. URL `https://arxiv.org/abs/2212.02407`.

Alain Trannoy, Sandy Tubeuf, Florence Jusot, and Marion Devaux. Inequality of opportunities in health in France: A first pass. *Health Economics*, 19(8):921–938, 2010.