UNIVERSITÀ DEGLI STUDI DI PADOVA

DIPARTIMENTO DI SCIENZE ECONOMICHE ED AZIENDALI "M. FANNO"

CORSO DI LAUREA MAGISTRALE IN APPLIED ECONOMICS

TESI DI LAUREA

ACCESS TO HEALTH INSURANCE AND BIOLOGICAL AGE: EVIDENCE FROM DNA METHYLATION IN THE U.S. HEALTH AND RETIREMENT STUDY

RELATORE:

CH.MO PROF. MARCO BERTONI

LAUREANDA: ELENA SARTORE

MATRICOLA N. 2121802

ANNO ACCADEMICO 2024 – 2025

Dichiaro di aver preso visione del "Regolamento antiplagio" approvato dal Consiglio del Dipartimento di Scienze Economiche e Aziendali e, consapevole delle conseguenze derivanti da dichiarazioni mendaci, dichiaro che il presente lavoro non è già stato sottoposto, in tutto o in parte, per il conseguimento di un titolo accademico in altre Università italiane o straniere. Dichiaro inoltre che tutte le fonti utilizzate per la realizzazione del presente lavoro, inclusi i materiali digitali, sono state correttamente citate nel corpo del testo e nella sezione 'Riferimenti bibliografici'.

I hereby declare that I have read and understood the "Anti-plagiarism rules and regulations" approved by the Council of the Department of Economics and Management and I am aware of the consequences of making false statements. I declare that this piece of work has not been previously submitted – either fully or partially – for fulfilling the requirements of an academic degree, whether in Italy or abroad. Furthermore, I declare that the references used for this work – including the digital materials – have been appropriately cited and acknowledged in the text and in the section 'References'.

Firma (Signature) ……Elena Sartore…

TABLE OF CONTENTS

1. **INTRODUCTION** 5
2. **LITERATURE REVIEW** 7
   - 2.1. Preliminary Conceptual Framework: Genetics and Epigenetics 7
     - 2.1.1. Genes and Genetics 7
     - 2.1.2. Epigenetics 10
   - 2.2. Aging and Epigenetics 11
   - 2.3. DNA Methylation and Epigenetic Aging 13
   - 2.4. The Role of Biomarkers: Epigenetic Clocks 15
   - 2.5. Generations of Epigenetic Clocks 17
   - 2.6. Computing Epigenetic Clocks: Horvath and Levine's Methods 21
   - 2.7. Correlations and Drivers of Epigenetic Clocks 25
   - 2.8. Predictive Power 27
   - 2.9. Alternatives to Epigenetic Clocks 28
   - 2.10. Health Insurance: Medicare 30
     - 2.10.1. Medicare in the U.S. Healthcare System 30
     - 2.10.2. Eligibility Criteria and Coverage Structure 31
     - 2.10.3. Effects on Health 32
     - 2.10.4. Financial and Other Social Implications 33
     - 2.10.5. Moral Hazard in Health Insurance 34
3. **DATA** 37
   - 3.1. HRS: Health and Retirement Study 37
   - 3.2. Study Population and Sample Weights 37
   - 3.3. Health Insurance in HRS 38
   - 3.4. Descriptive Statistics 38
   - 3.5. Epigenetic Clocks Included in the Analysis 39
   - 3.6. Age Measurement & Outcome Variable 40
4. **METHODOLOGY** 41
   - 4.1. Correlation and Causation: The Problem of Causality 41
   - 4.2. Identification Strategy 42
   - 4.3. Regression Discontinuity Design: General Framework 43
   - 4.4. Estimation in Fuzzy RDD 44

4.5. Practical Implications and Bandwidth Selection 45
5. **EMPIRICAL ANALYSIS** 46
5.1. Tests and Caveats 46
5.1.1. Manipulation Test 46
5.1.2. Covariate Balance Test 47
5.1.3. Graphical Inspection 48
5.2. Main Results 50
5.2.1. Clock-Specific Estimates 51
5.2.2. MEGA Clock 52
5.3. Donut Regression Discontinuity 52
5.4. Robustness Checks 54
5.5. Stratifications 55
5.6. Other Health Outcomes 55
6. **DISCUSSION** 58
7. **CONCLUSIONS** 61
8. **APPENDIX** 62
9. **REFERENCES** 64

1. INTRODUCTION

The benefits of health insurance are widely recognised. Promoting health, facilitating access to medical care, offering preventive services, and providing risk insurance, as well as protecting individuals financially from extraordinary healthcare costs, are just some of the objectives and benefits that health insurance can deliver. A substantial body of evidence suggests that health insurance coverage has a positive impact on individuals. Most analyses have found that coverage in the near-elderly slowed health decline and decreased mortality (Woolhandler & Himmelstein, 2017). Still, the predominantly private nature of the U.S. healthcare system always leaves room for consideration of health insurance in policy debates. In light of this, multiple government efforts have been made to reduce the number of uninsured individuals, with the goal of improving access to healthcare.

Within this setting, epigenetic biomarkers provide a valuable tool to capture the potential effects of Medicare: improvements in access to care and health outcomes may not only translate into better clinical conditions but also into a deceleration of epigenetic age.

This underscores the importance of understanding the link between health insurance and biological aging, not only because research on the connection between health insurance and epigenetic aging remains limited (Nwanaji-Enwerem et al., 2025), but also because the structure of the U.S. health services system may influence individuals' health outcomes and, ultimately, their lifespan. In this context, according to the World Health Organisation (WHO), global life expectancy at birth reached 73.3 years in 2024 – an increase of 8.4 years since 1995 – highlighting essential implications for public health systems. Specifically for the U.S., projections suggest that by 2030, roughly one in five people will be aged 65 or older, posing significant societal and healthcare challenges (Lipsky & King, 2015). Looking further ahead, by 2050, the global population is expected to achieve parity between the young and the elderly, with an estimated 2 billion people aged 60 or older (Harper, 2014).

Simultaneously, epigenetic clocks have emerged as DNA methylation-based biomarkers of healthspan and lifespan that are associated with several social, environmental and lifestyle factors. Few studies have investigated the relationship between Medicare and Epigenetic Age Acceleration. This work aims to address the gap by developing an identification strategy that provides causal estimates of the effects of Medicare as health insurance on Epigenetic Age Acceleration, using a cross-sectional sample from 2016. The focus on this relationship is motivated by the key role health insurance plays as a determinant of subjective well-being, a potential tool to prevent future health risks, and its connection in determining overall health outcomes across the life course.

As populations age, understanding the modifiable biological pathways, such as those influenced by healthcare access, that contribute to aging will be essential for designing effective public policies.

This analysis develops a Fuzzy Regression Discontinuity Design to investigate the potential causal effects of Medicare insurance on Epigenetic Age Acceleration, using epigenetic clocks as biomarkers. Specifically, the analysis focuses on the use of Hannum, Horvath, Levine, and GrimAge clocks, as well as the construction of the MEGA clock (Menta et al., 2025), a composite measure that integrates these previous ones.

The sample used consists of 3,784 individuals from the U.S. Health and Retirement Study (HRS), a nationally representative survey of the older households, that provides data on demographic characteristics, Medicare enrollment and genetic information.

The estimated effects around 65 years – the eligibility cutoff for Medicare – vary slightly depending on the chosen clock, but are generally not statistically significant. Consistently, chronic conditions show no discontinuity at 65. Stratification analyses confirm similar null results, suggesting that Medicare eligibility at age 65 does not affect age acceleration, at least in the short run. This finding is consistent with the notion that it takes time for access to health insurance to have an impact on health outcomes, including biological age. Future research based on medium- or long-run longitudinal data on Epigenetic Age Acceleration will permit to shed light on this matter.

The thesis unfolds as follows. Section 2 provides a review of the relevant literature and existing evidence; Section 3 presents the dataset and the sample used, and Section 4 discusses the methodology adopted and its general framework. Section 5 presents the primary empirical analysis, including results, robustness checks, and further specifications on health outcomes. Finally, Section 6 evaluates the results and their implications, followed by conclusions in Section 7.

2. LITERATURE REVIEW

2.1. Preliminary Conceptual Framework: Genetics and Epigenetics

2.1.1. Genes and Genetics

This research contributes to the literature examining how access to health insurance may influence biological aging, using indicators derived from epigenetic measures. Specifically, it focuses on the U.S. population and explores the potential impact of Medicare on biological aging processes. To support this analysis and to ensure a clear understanding of the terminology used throughout this study, some key concepts are introduced at the outset.

First and foremost, life depends on the ability of cells to store, recover and translate the genetic instructions required to build and sustain a living organism (Alberts et al., 2015). Hereditary information is passed from one cell to its daughter cells during cell division. This hereditary code is stored in deoxyribonucleic acid (DNA), which is the fundamental molecule of heredity in all living organisms.

To understand how this information is encoded and transmitted, it is essential to examine the molecular structure of DNA itself. Specifically, each DNA molecule consists of two long polynucleotide chains, called strands. Each strand is made up of four building blocks, called nucleotides: adenine, thymine, guanine and cytosine.

From a chemical point of view, each nucleotide consists of a five-carbon sugar, a phosphate group and a nitrogenous base[1]. They pair specifically and consistently: adenine always pairs with thymine, and guanine always pairs with cytosine. The two strands run in opposite directions (antiparallel) and are held together by weak hydrogen bonds between the paired nucleotides.

In eukaryotes, DNA is organised into structures called chromosomes, which reside in the cell nucleus[2]. Within chromosomes, genes are the specific segments of DNA that contain the instructions required to synthesise proteins[3]. Each gene represents a unit of hereditary information, contributing to both the traits of an individual and the general characteristics of the species. However, the entire DNA sequence carries vast amounts of information: a portion

[1] The five-carbon sugar is a sugar molecule with five carbon atoms (e.g., the deoxyribose in DNA); the phosphate group is a molecule that links nucleotides together in a chain; the nitrogenous base is one of four molecules (adenine, cytosine, guanine, thymine).

[2] To fit inside the nucleus, DNA is tightly packed in units called nucleosomes.

[3] Protein synthesis is fundamental for any organism, as proteins are highly versatile macromolecules that carry out most cellular and physiological functions. They regulate gene expression, facilitate cell communication, and enable cellular movement (Alberts et al., 2015). A more detailed explanation follows right below.

corresponds to protein-coding genes, while the remaining regions include regulatory elements and non-coding sequences that play other roles in cellular processes. During replication, one strand of DNA serves as a template for the synthesis of a complementary strand, ensuring accurate duplication of genetic material.

Figure 1 is provided to clarify the main concepts.

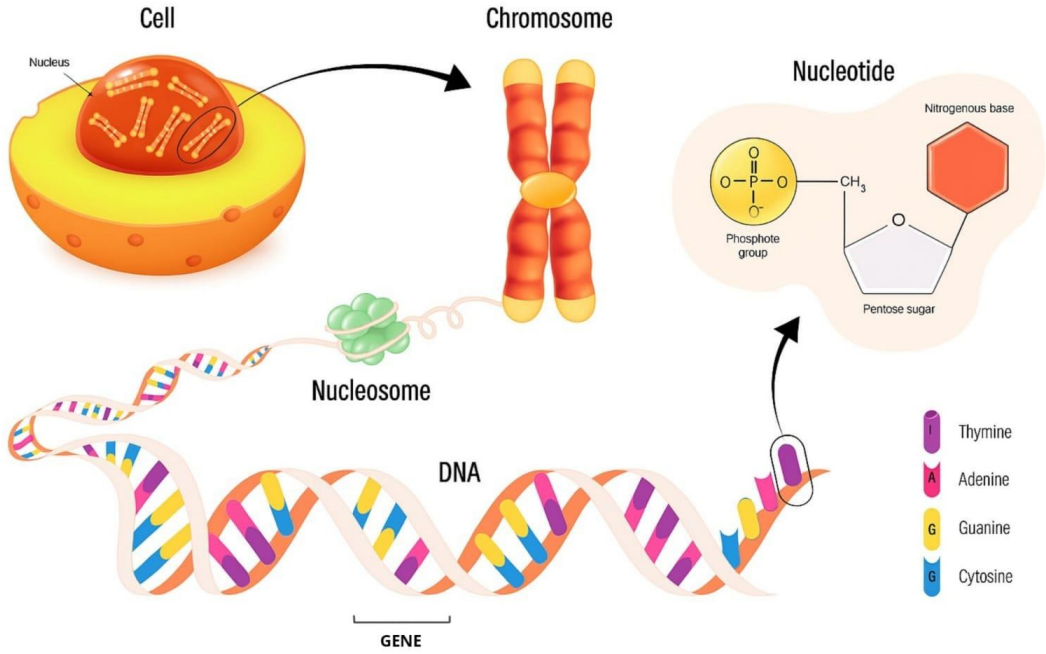

Figure 1: Preliminary Concepts (Google Images)

The entirety of an organism's DNA sequence, including all genes and regulatory elements, constitutes its genome, which is precisely defined as *"the entire set of DNA instructions found in a cell"* (National Human Genome Research Institute, NHGRI) for all the molecules and proteins the organism will ever produce.

However, it is essential to note that DNA in the genome does not directly control protein synthesis, but instead utilises ribonucleic acid (RNA), another crucial biological molecule. When the cell needs a specific protein, it first makes a copy of the appropriate portion of DNA into RNA. This copying process is called transcription. Then, the RNA copy is used as a guide to build the protein, a process known as translation. As a result, genetic information flows from DNA to RNA, and then to protein (Alberts et al., 2015).

A significant milestone in genome research was the Human Genome Project (HGP) – an international research initiative launched in 1990 and completed in 2003, which successfully sequenced approximately 92% of the human genome.

This achievement laid the groundwork for advances in whole-genome sequencing. Indeed, studying the genome as a whole is essential for understanding the complex network of DNA elements involved in gene expression, as analysing a single gene alone offers only limited insight into its function and regulation.

However, only in recent decades has genome sequencing become feasible at this broad scale. This shift was especially significant given the structural differences between prokaryotic and eukaryotic genomes. Indeed, while prokaryotic genomes are relatively small and contain few or no repetitive sequences, eukaryotic genomes are larger and rich in such sequences, making sequencing more complex (Saraswathy et al., 2016, pp. 95-107).

Still, even with complete genomic information, predicting an individual's phenotype – defined as observable traits such as height, eye colour, or blood type – remains a complex task.

This limitation has motivated researchers across both the biological and social sciences to investigate the underlying factors that drive individual differences in human traits by leveraging genetic information – particularly to study the genetic architecture of complex traits.

A commonly used approach to study genetic risk factors is the candidate gene method, which tests whether variations in a specific gene are associated with a given trait or disease (Kwon & Goate, 2000). This approach involves selecting genes thought to be relevant to the trait of interest and then examining whether specific alleles[4] may be involved in the disease. However, its application is limited by the reliance on a priori knowledge of the phenotype under investigation (Zhu & Zhao, 2007), making it vulnerable to biased selection.

Another widely used method is the polygenic risk score (PRS), which represents a mathematical aggregate of the effects of many DNA variants into a single measure. It estimates an individual's predisposition to a specific outcome, such as the risk of developing a disease (Sugrue & Desikan, 2019). While PRS are not equivalent to a diagnostic test, they offer insight into an individual's likelihood of developing a condition. Additionally, they are only modestly correlated with family history of disease, meaning they are not redundant with conventional risk indicators, and instead provide incremental predictive information (Kullo et al., 2022).

In this context, polygenic risk scores have been widely used to study educational attainment. Specifically, many genetic variants are associated with education, and PRS aim to combine them into a single measure that estimates an individual's genetic predisposition to educational success.

[4] An allele is one of two or more versions of a genetic sequence at a particular region on a chromosome. An individual inherits two alleles for each gene, one from each parent (National Cancer Institute). Similarly, an allele refers to a different version of the same gene, typically resulting in variations in how a trait is expressed.

These scores explain approximately 12 to 16% of the variance in educational attainment (Okbay et al., 2022). However, this is substantially lower than estimates from traditional twin studies, which rely on comparing within-pair correlations between identical (monozygotic) and fraternal (dizygotic) twins to generate heritability estimates (Branigan et al., 2013).

Across various twin studies, results show that a substantial portion of cognitive traits are heritable, ranging from 60 to 80% for many brain-related measures (Jansen et al., 2015).

This discrepancy highlights that, as with any predictor, PRS are not deterministic: non-genetic factors, such as lifestyle changes, family environment, socioeconomic status, school quality, and broader social influences, may significantly shape outcomes, potentially modifying the trajectory of prediction from PRS.

2.1.2. Epigenetics

It is precisely in this gap that attention turns to the contribution of other factors, and more recently to the study of epigenetic mechanisms.

Recent research in this area has primarily focused on how social and contextual variables interact with genetic factors, challenging the traditional "nature versus nurture" dichotomy, in which genes and environment are viewed as distinct explanations (Branigan et al., 2013).

Indeed, phenotypes are shaped not only by the genotype (the individual's genetic makeup) but also by environmental factors and random biological processes. As Burga and Lehner (2012) argue, the genome alone is insufficient to account for the variation observed in most phenotypic traits. Although a genotype may be directly related to a phenotype, there is rarely a one-to-one correspondence. Environmental factors, including diet, physical activity, and exposure to toxins, have a significant influence on phenotypic expression.

This means that an important component of phenotypic variance is not due to genetics. Indeed, predicting a phenotypic outcome from a genome may be very difficult. Burga and Lehner (2012) highlight some reasons why phenotypic prediction based solely on the genome is challenging. They are related to the stochastic molecular variations that can lead to different gene expressions, as well as to all the genetic variations within a single generation that may influence traits in the next. In addition, not just the environmental factors experienced by parents can affect the phenotypes of their offspring, but also the non-genetic inheritance (nutrition and metabolic disorders, immunological and hormonal factors, stress-related conditions) can further influence phenotypic outcomes across generations.

This brings us to the distinction between the genome and the epigenome. The epigenome consists of chemical compounds that modify or mark the genome, effectively guiding it on what, when, and where to express genes. These epigenetic marks do not alter the DNA sequence

itself but can be inherited during cell division and, in some cases, passed on to the next generation.

Given its influential role in gene regulation, the potential of epigenomic research is far-reaching. In addition to advancing our understanding of fundamental biological processes, epigenetic information has immediate applications in diagnostics. Since epigenetic changes are potentially reversible, they also hold promise for the development of novel therapeutic strategies (Murrell et al., 2005). Building on these theoretical foundations, it is now crucial to explore their application in practical fields and disciplines, particularly to understand how regulatory mechanisms extend beyond the genetic code —a concept at the core of the field of epigenetics.

The term *"Epigenetics"* was coined by Conrad Waddington in the 1940s: *"Epi"* means *"Over"*, while *"Genetics"* refers to genes, thus suggesting that the study goes beyond the pure genetic code.

While a universally accepted definition of epigenetics has yet to be established (Jablonka & Lamb, 2003), by the end of the 20th century, epigenetics had grown to become widely recognised as a distinct subdiscipline of biology (Jablonka & Lamb, 2003), particularly a branch of genetics.

It is crucial to distinguish epigenetics from genetics. While genetics deals with the transmission and processing of inherited characteristics encoded in the DNA, epigenetics is "*the study of changes in gene function that are mitotically and/or meiotically heritable and that do not entail a change in DNA sequence"* (Dupont et al., 2009).

In simpler terms, epigenetics is the study of how the body can switch genes "on or off" without altering the DNA itself, which helps explain why certain traits or changes in appearance or function occur predictably.

These changes include spontaneous self-organising processes driven by the physical and chemical properties of both internal and external environments, alongside mechanisms shaped by genes. Since environmental conditions strongly influence epigenetic changes, they are particularly relevant to policy variables such as access to healthcare.

2.2. Aging and Epigenetics

At the heart of the matter lies a central question: Why do we age? (Kirkwood & Austad, 2000) Even more puzzling is the evolutionary dilemma: if aging is harmful to the individual, why has natural selection not eliminated it? (Lipsky & King, 2015)

The answer could be found in the history of biological aging theories. One of the earliest came from August Weismann in the 1890s. He argued that aging, though harmful to the individual, serves a broader evolutionary purpose by eliminating older members of a population, so that they no longer compete for resources with younger generations.

Later, in 1952, Peter Medawar redefined aging as a process driven by the accumulation of cellular damage over time. He argued that natural selection mainly focuses on traits that influence early-life stages and reproduction. In contrast, harmful effects that appear later in life are less likely to be eliminated through evolution and therefore persist across generations.

More generally, theories of biological aging are typically categorised into two main groups: genetic or programmed theories and stochastic or damage-based theories (Lipsky & King, 2015).

Programmed theories suggest that humans possess an internal biological clock that regulates development, and so aging. According to this view, despite being able to delay death, a built-in biological upper limit to human lifespan exists that cannot be surpassed.

In contrast, stochastic theories claim that aging is the result of the inevitable accumulation of molecular and cellular damage over time, due to the failure of the body's repair mechanisms.

However, it is widely accepted in the literature that no single gene has been identified as solely responsible for aging. On the contrary, it is commonly believed that the genetic process of aging is multifactorial, involving complex interactions between several genes and environmental influences.

To provide answers to these core questions, over the last decades, scientists have attempted to develop quantifiable measures of aging.

One such measure is biological age, a concept that has been explored in scientific papers for the past 40 years (Klemera & Doubal, 2005). Sometimes referred to as physiological age, organismal age, or phenotypic age, biological age reflects an individual's biological condition, in contrast to chronological age, which measures the calendar time that has passed since birth (Horvath & Raj, 2018).

The measures used to compute biological aging come from epigenetics and epigenetic aging refers to a method of measuring biological age using various epigenetic biomarkers. One prominent example is provided by epigenetic clocks, which rely on patterns of DNA methylation (DNAm) observed at specific genomic sites. (Chervova et al., 2024). Epigenetic aging is also thought to have a comprehensive aging score (Crimmins et al., 2021) that can be measured during a person's lifetime using a single blood or saliva sample.

2.3. DNA Methylation and Epigenetic Aging

A growing body of research is focused on identifying biological markers that can better capture the aging process than chronological age alone. Several types of predictors have emerged, including telomeres, proteomic forecasters, transcriptomic and metabolomics markers and composite biomarker indicators. Among these, DNA methylation-based measures have become particularly prominent, mainly due to their reproducibility across populations.

The following paragraph will focus on DNA methylation, while the other predictors will be described in a subsequent section.

DNA methylation is an endogenous process by which the body activates or silences genes, depending on their function and the stage of development. This gene-switching mechanism is inherently dynamic and adaptable, enabling the organism to respond to internal and external changes.

From a chemical point of view, DNAm consists of the addition of a methyl group ($–CH_3$) to the carbon atom at the fifth position of the cytosine nucleotide, converting it to 5-methylcytosine. To clarify, each molecule of cytosine has atoms placed in specific positions, with carbon at the fifth position, serving as the attachment site for the methyl group. Figure 2 provides a graphical illustration of this concept. Crucially, the addition of the methyl group does not affect cytosine's ability to pair with guanine.

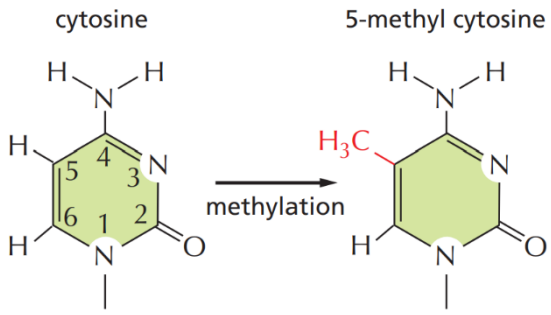

Figure 2: Methylation on Cytosine (Alberts et al., 2015)

Instead, DNAm regulates gene transcription without altering the underlying DNA sequence, often by inhibiting the transcriptional machinery (Schmitz & Duque, 2022).

Indeed, when methylation occurs, cytosine is typically silenced due to the presence of this additional methyl group. What is more, this modification creates a steric hindrance[5] which compacts the DNA structure and prevents RNA from accessing and reading the genetic

[5] Steric hindrance occurs when the presence and position of chemical groups within a large molecule prevent or hinder chemical reactions that proceed more easily in smaller molecules (Rackley, 2009).

sequence. As a result, the gene becomes transcriptionally inactive. In contrast, when methylation is absent, the DNA remains accessible, allowing regular gene expression to happen. The number of methyl groups depends on the individual's overall biological state: generally, the lower the number of methyl groups, the more advanced the biological age.
Figure 3 illustrates the difference between methylated and unmethylated cytosine.

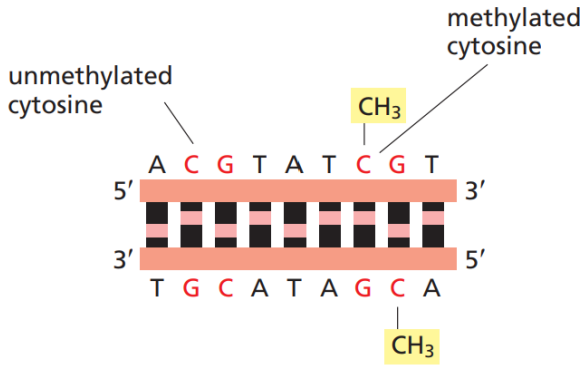

Figure 3: Methylation on the DNA strand (Alberts et al., 2015)

However, not all cytosines in the genome can be methylated: only those followed by a guanine on the same DNA strand are subject to methylation. In other words, DNA methylation occurs only at CpG sites and not on isolated cytosines.
CpG sites are regions in the DNA sequence where a cytosine is followed by a guanine in the 5′ to 3′ direction (Acharjee et al., 2023). The term *"CpG"* exactly refers to this cytosine-guanine dinucleotide, with the *"p"* standing for the phosphodiester bond between the nucleobases.
Whenever a CpG site is present on one DNA strand, it is naturally paired with another CpG site on the complementary strand, allowing for methylation to occur symmetrically on both strands. Each CpG dinucleotide is usually either fully methylated – meaning both cytosines on both strands are methylated – or fully unmethylated, with neither of the cytosines bearing a methyl group.
However, the distribution of CpG dinucleotides in the genome is not uniform: "isolated" CpG sites are pretty rare and underrepresented. Instead, there are specific regions called "*CpG islands"*, typically 1 to 2 kilobases (kb)[6] in length, where the CpG sequences are highly repeated and therefore densely clustered.

Research has shown that aging is associated with changes in the methylation status of millions of CpG sites across the human genome. In fact, the methylation patterns of many of the

[6] A kilobase is a unit of measurement used to designate the length of DNA or RNA. One kilobase is equal to 1,000 bases (NHGRI).

approximately 28 million CpG dinucleotides have been observed to shift with age (Horvath & Raj, 2018).

As previously mentioned, genes play a central role in determining lifespan, with the level of gene activation serving as a key indicator of biological age.

They are thought to be regulated by an internal biological clock that controls development, growth, maturity, and aging by sequentially activating and deactivating specific genetic pathways (Lipsky & King, 2015).

Indeed, like age itself, DNA methylation changes in a predictable, clock-like manner. In this context, the term "epigenetic clock" may refer to two distinct concepts (Horvath & Raj, 2018): an innate biological process that governs the progression of aging through molecular mechanisms and an age estimator based on DNA methylation patterns.

Technically, data on DNA methylation are collected from blood or saliva samples, through which researchers can analyse DNA methylation levels at specific sites to infer which genes are switched on or off.

The ease with which samples can be gathered highlights the value of these readily accessible tissue types for epigenetic research (Chervova et al., 2024).

DNAm age estimators are then constructed by applying supervised machine learning techniques to regress a transformed version of chronological age on a selected set of CpG sites (Horvath & Raj, 2018).

2.4. The Role of Biomarkers: Epigenetic Clocks

Biomarkers have the potential to uncover the reasons behind the aging process (Horvath & Raj, 2018). While the individual's biological age determines the long-range course of any accurate biological marker, its actual value may still fluctuate due to internal and external factors.

According to the 2016 Infoguides of the American Federation for Aging Research (AFAR), an accurate biomarker of aging must meet some specific criteria to be both accurate and valuable. It must be:

- Predictable in determining the later development of age-related conditions and diseases, independently of chronological age.
- Testable and not harmful to subjects. Alternatively, the test can be accurately reproduced without the need for specialised equipment or techniques.
- Valid for both animals and humans, since preliminary testing always involves non-human subjects.

Within this framework, epigenetic clocks have emerged as DNA methylation-based tools of both healthspan and lifespan (Nwanaji-Enwerem et al., 2025).

However, before the widespread adoption of epigenetic clocks as measures of biological age, Klemera and Doubal (2005) proposed an approach to conceptualise and compute biological age.

Their model was based on the recognition that aging is a natural process which progresses at different rates and levels across species and individuals. Crucially, they emphasised that among individuals with identical chronological ages, differences in biological age indicate heterogeneity in these underlying biological aging processes. According to their definition, a biomarker may be considered valid if it changes systematically with chronological age.

Epigenetic clocks emerged later as powerful biomarkers of biological age. However, they differ from each other in the algorithm[7] which they use to estimate a person's biological age (Chervova et al., 2024); it is crucial to remember that the sample collection should be in line with the requirements of the single chosen clocks.

A common feature of epigenetic clocks based on DNA methylation is that they estimate biological age in "epigenetic years", with the idea that ticks of the clock represent aging (HRS Epigenetic Clocks File). However, epigenetic clocks also differ in their strength of association with chronological age, among themselves, and with other various factors (Crimmins et al., 2021; Chervova et al., 2024). Assessing the correlation between different clocks is useful both to determine whether they capture similar aspects of aging and to test for multicollinearity in models with multiple clocks (Nwanaji-Enwerem et al., 2025).

Epigenetic clock estimates can also be obtained using the DNA Methylation Online Calculator[8], operated by the Epigenetic Clock Development Foundation.

Importantly, epigenetic clocks can also be used to derive another biomarker: Epigenetic Age Acceleration (EAA). EAA provides an accurate molecular readout of how much an individual's biological age deviates from their chronological age.

Mathematically, Epigenetic Age Acceleration is calculated as the residual from a regression of biological age, estimated through epigenetic clocks, on chronological age (Chervova et al., 2022). These residuals capture the component of biological age not explained by chronological age, isolating the portion of biological age not explained by the natural process of chronological

[7] The majority of the clocks are built using penalised regression models, which select a group of CpG sites that have a monotonically increasing relationship with age in a given training data (Simpson & Chandra, 2021).

[8] Webpage providing information on how to calculate DNA methylation age based on data measured using the Illumina Infinium platform through an R function (Steve Horvath & Clock Foundation Team).

aging. This approach captures all other influencing factors – such as environment, behaviour, and health conditions – that contribute to biological aging beyond simply getting older.

When epigenetic age exceeds chronological age, people are considered to be accelerated agers, suggesting a faster aging process. Conversely, individuals whose epigenetic age is lower than their chronological age are considered slow agers, indicating decelerated aging (Crimmins et al., 2024).

Epigenetic Age Acceleration is usually associated with adverse outcomes, as it implies a person is biologically older than expected. In contrast, age deceleration is generally viewed as beneficial. (Nwanaji-Enwerem et al., 2025).

However, it is important to specify that EAA should not be interpreted as a definitive indicator of health status or disease risk. Still, it should be considered with other biological and clinical markers (Chervova et al., 2024).

2.5. Generations of Epigenetic Clocks

As previously mentioned, epigenetic clocks have emerged as powerful tools based on DNA methylation, capable of estimating chronological age and predicting health outcomes.

This section offers a more in-depth examination of their development, highlighting the significant contributions that have shaped the field. In particular, Crimmins (2024) provides a comprehensive literature review that traces the evolution of epigenetic clocks, identifying four distinct generations.

At the outset, epigenetic clocks are typically constructed using statistical learning techniques to develop empirical models that reflect the evolution of DNA methylation patterns with age.

However, building such models is not straightforward, particularly when they are intended to work across multiple tissues. In this context, researchers face three main challenges (Horvath & Raj, 2018). First, there are substantial differences in DNA methylation patterns across various cell types and tissues, making it challenging to construct a universal clock. Second, age-related methylation changes are often tissue-specific, meaning that measures of aging in one tissue may not be generalised. Third, methylation patterns established early in life may differ from those that appear later.

The scientific response to these challenges began with the development of the first generation of clocks, which were trained to produce an estimate of epigenetic age, allowing for direct comparison with chronological age.

One of the earliest and most influential clocks was developed by Steve Horvath, followed by Gregory Hannum. Specifically, Horvath analysed a total of 8,000[9] DNA methylation samples obtained from newborns, children and adults (Horvath, 2013). These samples were then classified into 51 distinct tissue and cell types. This broad tissue coverage enabled the estimation of biological age across a wide range of biological systems. In contrast, Hannum's clock was explicitly trained on blood-based data (HRS Epigenetic Clocks File), making it less generalisable across tissue types. This first generation of clocks attempted to predict chronological age with methylation data directly.

Later research showed that this might not be the best way to detect accelerated biological aging and, therefore, disease risk (Michels, 2022). In turn, other clocks have been designed to examine disease and mortality risk, incorporating additional measurements specifically. As the field evolved, second-generation clocks were developed not to predict age itself, but rather lifespan and health outcomes. These clocks, such as Levine's PhenoAge (Levine et al., 2018), aim to serve as better predictors of mortality risk and physiological decline (Crimmins et al., 2021). Unlike first-generation clocks, they often incorporate sociodemographic characteristics and health behaviours, making them more broadly applicable (Chervova et al., 2024).

PhenoAge, developed in 2018, stands out for its greater power in predicting mortality, health span, and cardiovascular disease compared to earlier clocks (Levine et al., 2018).

Crimmins et al. (2024) provide further evidence that early-life socioeconomic status (SES) is significantly related to biological aging, especially when using second- and third-generation clocks.

These newer clocks – which are trained on health outcomes rather than age – indicate that higher SES is generally related to decelerated epigenetic aging, which in turn correlates with a lower risk of developing diseases.

The third generation was pioneered by DunedinPoAm38, whose novelty is based on the rate of change in epigenetic aging, offering insight into the speed of biological aging over time.

Unlike previous clocks, DunedinPoAm38 measures the pace of aging as a rate. It was trained using longitudinal data from individuals aged 26 to 28, capturing changes across this interval, rather than an absolute age (Nwanaji-Enwerem et al., 2025).

According to Horvath & Raj (2018), an age estimator deserves the label "epigenetic aging clock" only if it shows a correlation greater than r=0.80 with chronological age in a large, independent validation dataset, spanning a wide age range. Generally, first-generation clocks,

[9] The 8,000 samples do not correspond to 8,000 unique individuals, as a single person may have contributed multiple samples from different tissues. The exact number of individual donors is not specified in the original study (Horvath, 2013).

by design, tend to correlate more strongly with chronological age. Second- and third-generation clocks, such as GrimAge and DunedinPoAm38, exhibit more minor but significant correlations with age (Crimmins et al., 2021).

As already discussed, in the context of epigenetics and epigenetic clocks, most studies identify correlations between specific epigenetic marks. These associations – while informative for prediction – do not imply that changes in epigenetic markers cause those outcomes. Indeed, causality in epigenetics refers to understanding whether changes in epigenetic markers directly lead to changes in biological outcomes, which is much harder to establish and typically requires experimental or quasi-experimental designs. This is particularly true for lifestyle and environmental factors, where the association may be a correlation rather than a causal relationship (Michels, 2022). For instance, while smokers tend to have accelerated epigenetic aging as measured by certain clocks, it remains unclear whether smoking directly causes this acceleration in biological aging. It is possible that both smoking and faster epigenetic aging are jointly influenced by other confounding factors, such as genetic predispositions, stress or socioeconomic status.

To overcome this issue, recent research has increasingly relied on the fourth generation of clocks. They are called causal clocks, and they are composed of CpG sites causally linked to the aging process. They are identified through the use of the Mendelian randomisation technique, which applies to the use of genetic variation to address causal questions about how modifiable exposures influence different outcomes (Sanderson et al., 2022; Crimmins et al., 2024). It has been used, for example, by Berstein et al. (2022) to assess the causal role of epigenetic clocks in the development of multiple cancers. Their results – consistent across multiple MR analyses – suggest that accelerated biological aging, as measured by the GrimAge clock, may play a causal role in increasing the risk of colorectal cancer, particularly colon cancer. Nevertheless, while the relationship may not be causal, epigenetic clocks can still provide meaningful insights as long as they consistently predict future health outcomes (Michels, 2022).

As repeatedly mentioned, clocks do not respond equally to life conditions, raising the question of why this association with health outcomes vary that much in terms of intensity and risk.

One explanation lies in the fact that most of the clocks are only broadly correlated with each other. For example, the strongest correlation is observed between clocks that are built in a similar fashion, such as GrimAge and GrimAge2, which show a very high correlation ($r = 0.95$). At the same time, DamAge and AdaptAge exhibit a strong negative one ($r = -0.78$) (Crimmins et al., 2024).

Further differences arise from the construction of the clocks themselves. Clocks vary in the number and selection of CpG sites used. While some overlap exists, each clock includes unique sites, leading to differences in sensitivity and relevance. Additional variability stems from technical factors, such as the methylation arrays (chips) used in laboratory measurements, and the populations used for training – whether derived from single cohorts or longitudinal studies.

Other examples of epigenetic clocks include Zhang's clock, which reflects disease-related genetic markers and Levine's PhenoAge, developed in 2018. This clock is effective in differentiating morbidity and mortality and has greater predictive power for mortality, health span, and cardiovascular disease compared to the first generation of clocks. Another contribution is from Bocklandt et al. (2011), who developed the first epigenetic age estimator using DNA methylation data from saliva.

Some limitations remain in the use of epigenetic clocks. More generally, DNAm-based biomarkers will unlikely replace clinical biomarkers used to inform patient care, and it remains unknown whether it is possible to slow biological aging by directly targeting age-related DNA methylation levels (Horvath & Raj, 2018).
However, they still serve as powerful tools to explore the underlying causes of accelerated biological aging and the emergence of comorbidities, offering valuable insights into potential preventive strategies.

Table 1 summarises the main epigenetic clocks and some features.

Clock	Year	Significance
Bocklandt	2011	One of the first clocks
Garagnani	2012	Whole blood estimator
Horvath 1	2013	Most used clock
Hannum	2013	One of the first clocks
Weidner	2014	Estimation of epigenetic aging in blood
Lin	2015	Clock results are associated with mortality
VidalBralo	2016	Only a few CpG sites need
Yang	2016	Used to detect age acceleration in cancer and pre-cancerous lesions
Zhang	2017	Provides a mortality risk score
PhenoAge	2018	Used to predict mortality, cancer, healthspan and Alzheimer's disease
Horvath 2	2018	Used to predict lifespan and related age-related conditions
GrimAge	2019	Used to predict mortality and morbidity risk
GrimAge2	2019	Added DNAm components
DunedinPoAm38	2020	Third generation
DunedinPACE	2022	Third generation
CausAge	2024	Fourth generation

DamAge	2024	Fourth generation

Table 1: Summary of the Main Epigenetic Clocks
(HRS Epigenetic Clocks File; Crimmins et al., 2021; Michels, 2022; Crimmins et al., 2024)

2.6. Computing Epigenetic Clocks: Horvath and Levine's Methods

Understanding how epigenetic clocks are computed is fundamental to grasping how genetic information is translated into quantitative statistical frameworks, and especially how it is possible to discuss biological age in terms of years. The following outlines the procedures developed by Horvath and Levine.

Notably, Steve Horvath is widely regarded as a pioneer in the development of epigenetic clocks. He aimed to develop a predictor of biological age using DNA methylation patterns. To achieve this, he assembled an unprecedented collection of publicly available DNA methylation datasets. Horvath selected samples from a wide range of tissues and cell types, including blood, brain, skin, heart, lungs, kidneys, saliva and thyroid.

Methylation analysis was performed using either the Illumina Infinium Human Methylation27 BeadChip or the Illumina Infinium HumanMethylation450 BeadChip. The two platforms were merged by focusing on the approximately 26,000 CpG sites shared between them. Of these, around 21,369 high-quality CpG sites, with fewer than 10 missing values, were selected for model training.

To construct the epigenetic clock and ensure an unbiased validation, Horvath used training datasets to develop the age predictor. At the same time, its accuracy was then evaluated on an independent test set to minimise the risk of overfitting.

Specifically, as described in Additional File 2 from Horvath (2013), the selection of the training datasets was guided by the three key criteria: the training data should represent a broad spectrum of tissues and cell types; the mean age in the training data had to be comparable to that of the test data; and a substantial proportion of the samples (37%) had to be measured using the Illumina 450K platform (a standard DNA methylation sampling chip) to ensure the model's relevance for future research.

From a mathematical perspective, methylation at each CpG site is quantified using a β value. Specifically, on the Illumina platform, methylation is detected through fluorescent signals, where Signal A corresponds to the intensity of methylation (denoted as M) and Signal B to the intensity of unmethylation (denoted as U).

The β value is calculated as the ratio of fluorescent signals, as described by Equation 1:

$$\beta = \frac{Max\ (M, 0)}{[Max\ (M, 0) + Max\ (U, 0) + 100]} \quad \text{Eq. (1)}$$

Low-quality samples were excluded from the analysis. Non-cancer samples were included only if they exhibited no technical artefacts. In contrast, all cancer samples were retained regardless of technical quality to preserve biologically meaningful variation associated with cancer-related methylation patterns.

Additional processing steps were performed to normalise the data and ensure comparability across platforms[10].

Based on the training dataset, Horvath then applied a transformation to chronological age using the following function F, defined by Equations 2 and 3 as:

- For ages up to adulthood (20 years for humans):

$$F(age) = log\ (age + 1)\ - log\ (adult.age + 1) \quad \text{Eq. (2)}$$

Age is transformed logarithmically before adulthood, which emphasises early-life changes.

- For ages after adulthood (above 20 years):

$$F(age) = \frac{(age - adult.age)}{(adult.age + 1)} \quad \text{Eq. (3)}$$

Age is transformed linearly after adulthood, assuming that aging happens more steadily.

The function F is continuous and smooth, it is monotonically increasing, and it works even for prenatal (negative) ages because of the "+1" inside the log.

An Elastic Net regression model[11] was then employed to regress the transformed version of age on the approximately 21K β values in the training data. This resulted in a linear regression model, whose coefficients relate to transformed age as Equation 4 describes:

$$F(chronological\ age) = \beta_0 + \beta_1 CpG_1 + \cdots + \beta_{353} CpG_{353} \quad \text{Eq. (4)}$$

[10] Although both the Illumina 27K and 450K arrays measure DNA methylation at CpG sites, they rely on two different probe types – type I and type II – with different characteristics. Typically, methods like Beta MIxture Quantile dilation (BMIQ) – useful as a preprocessing step for any study using the Illumina Infinium 450k platform (Teschendorff et al., 2013) – are used to adjust for these differences. However, since Horvath's study only involved type II probes, he could not apply BMIQ in its standard form. Instead, he developed a custom normalisation method ("gold standard") based on the largest dataset in his study. Each sample's methylation profile was then adjusted to match this reference, ensuring consistency across the data.

[11] An Elastic Net regression is a hybrid machine learning algorithm that combines Ridge regression and Lasso regularisation. Like Lasso, Elastic net can generate reduced models by generating zero-valued coefficients. Empirical studies have suggested that the Elastic Net technique can outperform Lasso on data with highly correlated predictors (James et al., 2021).

Where 353 is the number of CpG sites automatically selected by the Elastic Net regression model. Indeed, out of the roughly 21,000 CpG sites considered, the Elastic Net identified these 353 as the most informative predictors of transformed chronological age.

The predicted transformed age $F(chronological\ age)$ is then converted back to the original scale age using the inverse function F^{-1}. This step is necessary because, before training the regression model, age was transformed using the function F to simplify the modelling process, due to different aging dynamics before and after adulthood. As a result, the regression model does not directly estimate age in regular years, but instead predicts $F(age)$. To recover biological age in regular years, the transformation must be reversed using the inverse function detailed in Equation 5:

$$DNAmAge = inverse.F(\beta_0 + \beta_1 CpG_1 + \cdots + \beta_{353} CpG_{353}) \qquad \text{Eq. (5)}$$

Thus, the regression model can be used to predict the transformed age value by inputting the β values of the selected CpG sites into the regression.

Horvath's clock was then evaluated using three measures of predictive accuracy: the Pearson correlation coefficient, median error and average age acceleration. The model performed remarkably well in most tissues and cell types (Horvath, 2013). It demonstrated high accuracy, with an age correlation of 0.96 – indicating that DNAm age closely matches actual ages – and a median error of 3.6 years, meaning that for half of the individuals, prediction deviates by less than 3.6 years.

These statistical results are illustrated in Figure 4, which displays the model's fit on both training and testing datasets. The straight line represents perfect prediction, where DNAm age equals chronological age. Most data points lie close to this line, with remarkably accurate predictions observed in blood and brain tissue samples, especially among adolescents and children.

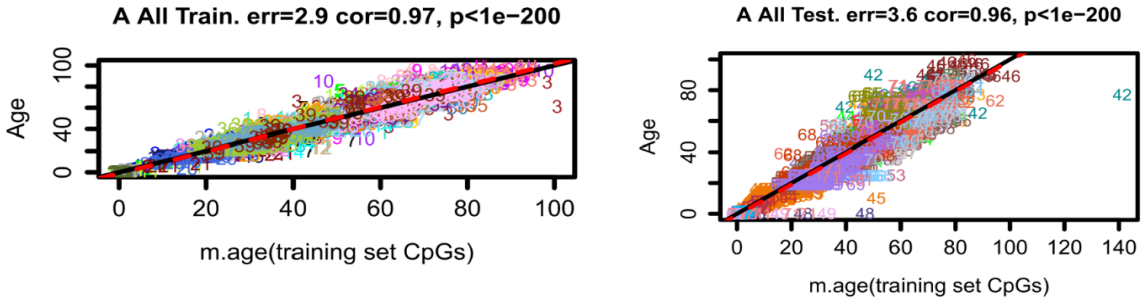

Figure 4: Horvath's train and test results (Horvath, 2013)

An interpretation from the concept of epigenetic clocks is that the "ticking rate" is faster in early life, when the body is undergoing rapid development and experiencing more frequent cell division. As individuals reach adulthood, this pace of aging stabilises into a more constant rate. Additionally, Horvath (2013) notes that *"the broad sense of heritability of age acceleration is 100% in newborns and 39% in older subjects, which suggests that non-genetic factors become more relevant later in life".* This implies that, at birth, individuals have minimal exposure to environmental and lifestyle factors; therefore, heritability and all observed variations – such as age acceleration in this case – can be attributed entirely to genetics. As people age, these other external influences accumulate, reducing the proportion of variation explained by heritability and allowing non-genetic factors to play a bigger role.

Following Horvath's work, Levine et al. (2018) developed their DNAm PhenoAge using an innovative two-step process: they created a measure of "phenotypic age" using data from the National Health and Nutrition Examination Survey (NHANES). This was first made through the application of a Cox penalised regression model[12] where the hazard of aging-related mortality was regressed on forty-two clinical explanatory variables to select those for inclusion in the phenotypic age score.

A ten-fold cross-validation was employed to select the parameter value (lambda) for the penalised regression, to use only the relevant biomarkers.

Of these forty-two, only ten variables were selected (including chronological age and nine biomarkers).

The second part of the first step consisted of including the ten variables in a parametric proportional hazards model with Gompertz distribution[13]. Based on this model, they estimated the ten-year (120 months) mortality risk of the jth individual based on the cumulative distribution function in Equation 6:

$$MortalityScore_j = CDF(120, x_j) = 1 - e^{-e^{x_j b}(expexp\,(120\times\gamma)\,-1)/\gamma} \qquad \text{Eq. (6)}$$

Where xb is the linear combination of biomarkers from the fitted model.

[12] A Cox penalised regression model is a survival model used to examine the relationship between the time until an event occurs and the explanatory variables. The hazard rate is the risk of the event happening at a particular time (Rinne, 2014).

[13] The Gompertz distribution is a probability distribution commonly used in demography with a general cumulative distribution function of Fx=1-e-ab(ebx-1) (Lenart, 2012).

Finally, once parameters are estimated, they solve for t (time), which in this case is age. So, the mortality score was converted into units of years using the formula in Equation 7:

$$PhenotypicAge_j = 141.50225 + \frac{ln\,(-.00553 * ln\,(1 - MortalityScore_j))}{0.090165} \quad \text{Eq. (7)}$$

Which provides an age-like number that reflects the person's risk of dying within the next ten years.

Once PhenoAge was built, the second step involved converting PhenoAge into a DNA methylation-based biomarker, referred to as DNAm PhenoAge. For this purpose, Levine used DNA methylation data from whole blood and applied Elastic Net regression with cross-validation. This resulted in a model where DNAm PhenoAge is expressed in Equation 8 as a weighted linear combination of methylation levels at 513[14] CpG sites:

$$DNAm\ PhenoAge = \beta_0 + \beta_1 CpG_1 + \cdots + \beta_{513} CpG_{513} \quad \text{Eq. (8)}$$

As anticipated, the key difference between Horvath and Levine's methods lies in the outcome used in the prediction, chronological vs. phenotypical age.

2.7. Correlations and Drivers of Epigenetic Clocks

Epigenetic age is not only correlated with age but is also closely linked to various health outcomes.

Numerous studies (McCrory et al., 2022; Chervova et al., 2022; Crimmins et al., 2024; Chervova et al., 2024) have documented associations between Epigenetic Age Acceleration and a wide range of factors – including physiological, environmental, social and pathological characteristics – that influence epigenetic aging.

Particularly, Crimmins et al. (2024) emphasise that it is not solely economic conditions that cause accelerated epigenetic aging, but rather their association with exposure to adverse circumstances and risky habits. Indeed, deviations in EAA have been connected to several high-risk behaviours, such as smoking, excessive alcohol consumption, poor diet, obesity, inadequate sleep, and psychological stress. These factors are subsequently reflected in adverse health outcomes, such as cancer, metabolic syndrome, and cognitive decline, all of which further contribute to accelerated aging.

Figure 5 reports the overview of the classification.

[14] The 513 CpG sites are those selected by the Elastic Net regression out of the 20,169 total ones.

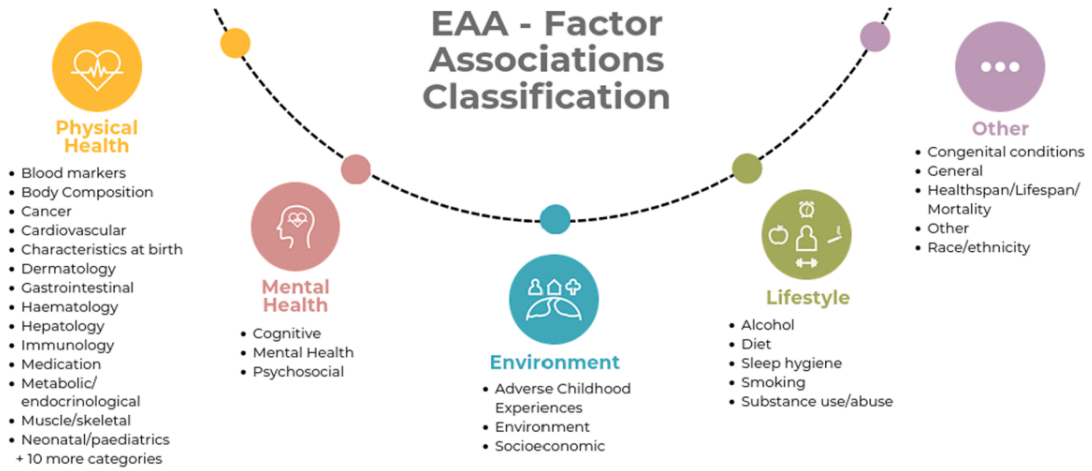

Figure 5: EAA Factor Associations Classification (Chervova, 2024)

Another study that helps clarify the analytical design behind these correlations is the one by Crimmins et al. (2021), which investigates 13 epigenetic clocks derived from the Health and Retirement Study. The authors first examined the correlations between each clock and chronological age, finding variation from 0.2 to 0.9, which suggests that the clocks differ markedly in how closely they align with aging.

They then checked the correlations among the clocks themselves and among measures of accelerated aging. Next, they perform ordinary least squares regressions (OLS), modelling EAA as a function of sex, education, race/ethnicity, obesity, and smoking status, to isolate the independent association of each variable. The main results indicate that accelerated aging is generally more pronounced among men. Several clocks show a link between lower educational attainment and faster aging, though the strength of this association varies. The majority of the clocks also associated accelerated aging with obesity and lifetime smoking. On the other hand, the associations with race and ethnicity show considerable variability across clocks.

Epigenetic aging signatures can also reflect the cumulative impact of these exposures, including those experienced early in life (McCrory et al., 2022).

For example, Schmitz and Duque (2022) demonstrate that in-utero exposure to the Great Depression had lasting effects on epigenetic markers later in life. This finding highlights the lasting impact on next-generation epigenetic signatures, further underscoring the critical importance of early-life environments.

To facilitate the association between epigenetic clocks and the wide range of external factors, the Tracker of Epigenetic Age Acceleration Association with Phenotype & Environmental Exposure (TEAPEE) offers an interactive platform for users to explore these relationships.

Based on these associations, newer EAA measures focus more on physiological dysregulation. They derive from second-generation epigenetic clocks and still maintain strong correlations with chronological age (Chervova et al., 2022).

However, comparing different epigenetic measures is crucial for an accurate interpretation. Each measure, especially epigenetic clocks, was developed based on different algorithms, which means they may reflect different dimensions of aging process.

Significantly, Epigenetic Age Acceleration can be influenced by the distribution of white blood cells (WBC) in the blood sample (Crimmins et al., 2025). Since DNA methylation profiles are derived from the whole blood, variations in leukocyte composition can introduce bias. These immune cells change in proportion over time across the organism. As a result, apparent changes in epigenetic age may sometimes be due to WBC distribution rather than proper biological aging.

2.8. Predictive Power

The accurate identification of individuals at greater risk of age-related diseases and death is a critical goal in both clinical and economic research, as it enables efficient targeting for prevention and intervention strategies. This is what has prompted growing interest in clinical biomarkers that accurately predict the risk of age-related diseases and mortality (Fransquet et al., 2019).

In this context, the documented correlation between epigenetic age and morbidity, mortality, and other factors suggests the potential of epigenetic clocks for risk prediction and identification (Duan et al., 2022).

Numerous studies highlight the quantitative predictive utility of these methods. For instance, Bocklandt et al. (2011) estimate age from a saliva sample with an average accuracy of 5.2 years, whereas Zhang's clock demonstrated accurate predictions across various tissues, including breast, liver, adipose, and muscle.

In a meta-study of four longitudinal cohort studies, Marioni et al. (2015) found that a five-year increase in Epigenetic Age Acceleration was associated with a 21% higher risk of mortality using Hannum's clock and an 11% using Horvath's, even after adjusting for age and sex.

These findings reinforce the utility of epigenetic age not only for all-cause mortality, but also for anticipating risk of neurodegenerative disease, metabolic disturbances, as well as Alzheimer's and Parkinson's diseases and cognitive decline.

Crucially, the predictive accuracy of epigenetic clocks improves with the inclusion of more age-related CpG sites, which enhances robustness and reduces noise (Duan et al., 2022).

Still, it is essential to report what Karin B. Michels states in her book "Epigenetic Epidemiology" about the use of epigenetic clocks as predictive tools. In particular: "*When the objective is to use epigenetic clocks as a general screening tool to inform risk management, causality is irrelevant if epigenetic clocks provide an easy, cost-effective means to measure facets of aging that strongly predict risk of morbidity and mortality"*. This means that, as long as the clocks are reliably correlated with future health outcomes, they are still valuable, even if the relationship is not causal. In many contexts, especially when evaluating long-term outcomes, policy interventions and risk stratification, predictive accuracy can be more relevant than identifying a causal explanation.

In this regard, epigenetic clocks offer a practical alternative to complex and costly biological assays. For example, the DunedinPoAm clock can be estimated from a single blood-based sample rather than measuring the longitudinal change in 18 biomarkers.

Moreover, epigenetic clocks show external validity, characterised by generalizability and transportability. When heterogeneity in sex, race/ethnicity, and environmental exposures is accounted for, clocks trained in one population often retain predictive accuracy across others (Michels, 2022). Additionally, many newer clocks are also validated using independent test sets, ensuring their generalizability in out-of-sample contexts.

From an econometric perspective, this broad applicability enhances the clocks' utility in comparative and cross-population research.

Perhaps more notably, recent studies have shown that epigenetic clocks also respond to changes in behaviour or medical treatments. The GrimAge clock was shown to reverse – meaning to go "younger" – by 2.5 years after one year of therapy designed to improve immune function. Similarly, Horvath's clock decreased after 16 weeks of vitamin D supplements.

These multiple functions define epigenetic clocks not only as tools for risk prediction, but also as intermediate outcome variables to quantify the biological impact of behavioural and medical interventions, making them particularly appealing for economic evaluations of health programs and policies.

2.9. Alternatives to Epigenetic Clocks

Although DNA methylation data are widely used in many fields, their reliability can vary considerably, and this variability is a fundamental part of ensuring the replicability of findings. Sugden et al. (2020) addressed this issue by linking the concept of probe reliability to the likelihood that associations between environmental exposures and specific DNAm probes

would replicate across independent studies. They found that the unreliable probes were less replicable and generated an unknown volume of false negatives.

To face this issue and improve the reliability of DNA methylation-based measures, a method based on Principal Component Analysis[15] (PCA) has been developed (Crimmins et al., 2025). Rather than relying on the use of several CpG sites, each with a raw methylation level, PCA captures patterns of co-movement across CpG sites, summarising this information into principal components. These components are then used to build a PCA-based epigenetic clock, derived from the covariance structure of the CpG sites. An early example of this approach is reported in Levine (2013).

Jylhävä et al. (2017) conducted a comprehensive literature review of potential biological age estimators. While epigenetic clocks are widely regarded as the most promising molecular estimators of biological age, it remains relevant to review alternative methods. The authors identified five main alternatives to the epigenetic clocks, which are briefly summarised here.

- Telomeres, for instance, are repetitive DNA sequences located at the ends of chromosomes. Although they do not encode genes, they play a crucial role in protecting the genome during cell division. Every time a cell divides, telomeres progressively shorten due to the limitation in DNA polymerase, the enzyme that synthesises new DNA strands. As a result, telomeres get shorter, and their length has long been thought to reflect aging. The association between short telomeres and an increased mortality risk has been repeatedly demonstrated in numerous studies, particularly in relation to cancer and cardiovascular mortality predictions (Jylhävä et al., 2017).

  However, Horvath and Raj (2018) pointed out that telomere length is only weakly and negatively associated with biological aging and life expectancy. In other words, while individuals with shorter telomeres may appear biologically older, this correlation remains relatively weak.

  Additionally, they discuss the role of TERT (Telomerase Reverse Transcriptase), a gene involved in maintaining telomere length. Evidence suggests that sustaining telomere length through TERT activity does not halt biological aging, suggesting that telomere preservation alone is insufficient as an anti-aging mechanism.
- Another alternative to epigenetic clocks is represented by transcriptomic-based predictors, which are relatively new statistical tools used to predict phenotypes, linking genetic variation and phenotypic variation (Quach et al., 2017). They appear to have a

[15] Principal Component analysis (PCA) is a technique for reducing dataset dimensionality by creating new uncorrelated variables that maximise variance, improving interpretability with minimal information loss (Jolliffe & Cadima, 2016).

low correlation with epigenetic clocks, suggesting they capture distinct biological processes. However, as noted by the authors, further validation is still required.

- A similar argument applies to proteomic predictors, based on proteins, and metabolomics-based predictors, focused on the large-scale profiling of the substrates and products of metabolism. Although these approaches offer promising insights, they remain insufficiently validated and currently lack the robustness and predictive power of epigenetic clocks.
- Finally, another approach involves the use of composite biomarker predictors. Levine (2013) proposed the construction of a biological age measure by merging multiple biomarkers into a single latent variable, designed to better reflect the complexity of the aging process. In the analysis, 21 biomarkers were grouped into seven domains and combined into one predictor.

Despite the availability of various alternatives and supporting evidence, some critics doubt the existence of accurate biomarkers of aging. They argue that aging is driven by a single biological process. Instead, aging involves various overlapping but distinct changes. Additionally, since different body systems age at different rates, measuring a single "rate of aging" is challenging. (Infoguides, AFAR).

2.10. Health Insurance: Medicare

2.10.1. Medicare in the U.S. Healthcare System

The U.S. health insurance structure has long been the subject of an intense debate regarding both its efficiency and its ethical foundations.

Health insurance coverage in the United States is primarily divided into three categories: Medicare, Medicaid and private insurance, which is either individually purchased or provided by an employer.

Because individual states largely determine eligibility criteria and coverage levels, significant disparities often arise in access to care. These disparities can lead to *financial toxicity*, where patients face difficult choices between affording medical treatment and covering other essential living expenses (Cavaliere, 2018).

In 2010, the Affordable Care Act (ACA), a comprehensive healthcare reform introduced under President Obama, aimed to expand access to healthcare and reduce the number of uninsured individuals. According to the most recent Current Population Survey (CPS) report, in 2023, 92% of people, equivalent to 305.2 million individuals, had health insurance coverage for at

least part of the year. Both private and public coverage rates remained statistically unchanged compared to 2022, while Medicare coverage rose slightly, increasing by 0.2 percentage points. Despite these gains, the United States continue to undergo rapidly growing healthcare costs and a substantial number of non-elderly people without insurance. This raises the critical question of whether high healthcare spending leads to an improved quality of care (Harrington, 2010).

2.10.2. Eligibility Criteria and Coverage Structure

Medicare is the primary insurance program analysed in this context and by far the most extensive government insurance program in the U.S. It was enacted in July 1965 and implemented nationwide the following year. As a government-provided scheme and a federally administered program, it offers uniform coverage across all U.S. states. Unlike Medicaid, Medicare is not tied to lower incomes or a state of poverty; instead, its eligibility is primarily based on age, for individuals aged 65 and older.

It is possible to get earlier enrollment for individuals with certain conditions, such as a qualifying disability, end-stage renal disease (ESRD) or Lou Gehrig's disease (ALS).

Card et al. (2008) suggest that Medicare eligibility causes a sharp increase in the use of healthcare services, with patterns of gain varying across groups depending on the type of service.

In this context, the U.S. Department of Health and Human Services (HHS) outline its four major programs:

- Part A (Hospital Insurance)

Individuals aged 65 and older are automatically enrolled in Medicare Part A, which covers inpatient care in hospitals, critical access hospitals, and skilled nursing facilities (excluding custodial or long-term care). It also provides coverage for hospice care and certain home healthcare services. Most people do not pay a monthly premium for Part A, as eligibility is typically earned through payroll tax contributions made by the individual or their spouse during employment.

- Part B (Medicare Insurance)

Medicare Part B covers a range of medical services and items not supplied by Part A coverage, such as physician visits, preventive care, laboratory tests and diagnostic screenings. Enrollment in Part B is optional and requires a monthly premium, with high-income individuals paying more.

- Part C (Medicare Advantage Plus)

Medicare Part C, also known as Advantage Plan or MA Plan, allows beneficiaries to receive Medicare benefits through private insurance plans, such as HMOs or PPOs[16].
These plans are required to cover all services under Part A and Part B coverage and often include additional services, such as vision, hearing, dental, and/or health and wellness programs, as well as prescription drug coverage (Part D).
Medicare pays a fixed amount per enrollee to plan providers, although costs and service rules may vary by plan and are subject to annual changes.

- Part D (Drug Coverage)

Medicare Part D is Prescription Drug Coverage. It was introduced in 2006 and made available to all Medicare beneficiaries regardless of income, health status, or existing medication needs. Premiums for Part D vary depending on the specific plan chosen by the enrollee.

2.10.3. Effects on Health

A substantial body of research indicates that individuals with health insurance tend to experience better health outcomes than those without coverage. In particular, numerous studies reported in the comprehensive literature by McWilliams (2009) use data from the Health and Retirement Study to examine the relationship between insurance status and health outcomes.
For instance, Polsky et al. (2006) followed participants after the age of 65 and compared self-reported health trajectories between insured and previously uninsured near-elderly adults. They found that health declines became less common after Medicare eligibility began, with the share of previously uninsured adults reporting excellent or very good health increasing by 7.7 percentage points more than their pre-65 trend.
Similarly, Card et al. (2004) analysed the longitudinal health scores of participants with cardiovascular disease or diabetes. Their study showed both immediate health improvements within the first three years of Medicare eligibility and an attenuation in the rate of health decline.
However, as Levy and Meltzer (2008) argue, some of these studies face a significant methodological limitation: the endogenous status of health insurance in observational studies, which prevents the drawing of practical causal inference. While much of the existing literature focuses on the causal effects of health insurance on health outcomes, there remains a scarcity of research examining the relationship between health insurance and epigenetic age.
A recent contribution to this emerging field is the study by Nwanaji-Enwerem et al. (2025), which analyses data from the NHANES. Their results suggest that more comprehensive

[16] HMO (Health Maintenance Organisation) is an in-network provider, meaning it establishes networks of providers who have agreed to provide care to their members at negotiated, reduced rates; PPO (Preferred Provider Organisation) is an out-of-network provider (WebMD).

insurance plans are more beneficial to the health of the insured and contribute to a potentially lower epigenetic age. In particular, the authors examined biomarkers associated with having any form of insurance, but further stratified the analysis by private and public health insurance. Private health insurance was marginally associated with a 1.32-year increase in SkinBloodAge and significantly associated with both a lower GrimAge2 (-2.70 years) and a lower DunedinPoAm (-0.05). For public insurance, similar positive trends were observed for GrimAge2 and DunedinPoAm, although only SkinBloodAge results reached statistical significance. When the analysis was extended to more comprehensive insurance plans – including, for example, dental coverage – the estimated effects were greater in magnitude and remained statistically significant for private insurance. For public insurance, the association with DunedinPoAm reached -0.09.

However, they caution that these findings describe statistical associations, not causal effects, and could be partly driven by the higher socioeconomic status, typically associated with individuals enrolled in more extensive insurance plans. Nevertheless, the authors found that marginal associations of health insurance with lower epigenetic age persist even after adjusting for these and other health and lifestyle variables. The authors also emphasise the importance of considering the broader meaning that more comprehensive insurance coverage may be associated with greater health benefits. While approximately 8% of American adults remain uninsured, it is estimated that 23% are underinsured, meaning they possess an insurance that does not provide sufficient financial protection to make healthcare meaningfully accessible (Collins & Gupta, 2024).

As a result, underinsured individuals, despite having some coverage, may face barriers to care and financial strain similar to those experienced by the uninsured.

2.10.4. Financial and Other Social Implications

The impacts of Medicare have been extensively studied. Card et al. (2008), using survey data from the National Health Interview Survey (NHIS), found that Medicare coverage increases by 60 percentage points at the age of 65, from a base level of 12% among 63-to 64-year-olds. Additionally, they found that Medicare eligibility reduces disparities in insurance coverage, as groups with lower coverage rates before 65 experience larger gains at age 65. Furthermore, the 65-year threshold is associated with an increase in hospitalisation rates, but the gains are more pronounced among better-educated white individuals, compared to other groups.

Finkelstein and McKnight (2008) analysed two key advantages of public health insurance: the direct benefits of risk protection and the indirect health benefits. They state that they were not able to detect a clear or immediate effect on elderly mortality following the implementation of

Medicare. However, they found that during Medicare's first five years (using data on healthcare expenditures from 1963 to 1970), out-of-pocket spending declined by 40% for individuals in the top quartile of out-of-pocket expenditures. In other words, after Medicare was introduced, those who had previously spent the most out of pocket on healthcare saw their personal medical costs fall by nearly half. This highlights Medicare's effectiveness in protecting the highest-cost individuals from a substantial financial burden.

This effect of public health spending in reducing individuals' financial exposure is consistent with Lichtenberg's (2002) findings, where he estimated that each additional dollar of public health expenditure crowds out approximately 43 cents of private health spending. He also highlights that the introduction led to an increase in inpatient care, primarily driven by a rise in the frequency of physician visits among individuals aged 65. Additionally, the Medicare-induced increase in health utilisation was associated with a reduction in the number of days spent in bed.

Medicare coverage has been detected to have an impact on the use of basic clinical services for previously uninsured adults (McWilliam et al., 2003). More precisely, the use of cholesterol testing, mammography and prostate examination increased substantially, especially among previously uninsured adults, contributing to a marked reduction in disparities in access to preventive care.

It goes without saying that understanding how access to health insurance may influence biological aging is crucial to anticipate the future implications of healthcare systems and policies for both younger and older generations. Incorporating indicators derived from epigenetic measures, alongside additional biological measures and socioeconomic metrics, is a critical step toward defining and evaluating the long-term impact of insurance coverage on population health.

The present study aims to contribute to the analysis and description of the relationship between health insurance and epigenetic aging among U.S. adults. Its ultimate goal is to examine how epigenetic aging, as measured by epigenetic clocks, may help better understand the effects of health insurance on the population. These insights could support policy development, particularly in addressing the needs of an aging population, as increasing life expectancy will require more effective strategies for managing the health and care of elderly individuals.

2.10.5. Moral Hazard in Health Insurance

Asymmetric information is a key feature of all markets, including the health insurance market. As insurers cannot fully observe individuals' health-related behaviours, coverage may create incentives for people to engage in riskier activities or to consume more healthcare services.

Asymmetric information gives rise to both moral hazard and adverse selection, and in the insurance literature, both mechanisms are expected to generate a positive correlation between risk and coverage (Chiappori et al., 2006).

More specifically, adverse selection refers to the fact that people with higher risk are more likely to purchase insurance than those with lower risk. As confirmed by evidence, individuals with a family history of illness are more predisposed to buy health insurance than those whose relatives have lived long and healthy lives.

With respect to moral hazard, the term dates back to Arrow (1963), where it originally referred to a situation in which one party engages in risky behaviour knowing that they are protected against the risk and the associated costs, which another party instead bears.

In health insurance, however, the concept is typically associated with Pauly (1968), who argued that under an insurance contract that reduces the price of care, there is nothing unethical or immoral in the insured individual's decision to consume more services than they would when uninsured. The central point of moral hazard theory is precisely that insurance removes part of the incentive to restrict the use of insured services. For example, under full insurance, there is no financial reason for someone not to visit the doctor.

This perspective laid the foundation for later formal definitions, such as Einav and Finkelstein's (2018) definition of moral hazard as the notion that insurance coverage, by lowering the marginal cost of care to the individual (often referred to as the out-of-pocket price of care), may increase healthcare utilisation.

Indeed, health insurance, by design, lowers the price individuals pay for their medical care. Since demand curves generally slope downward, making something cheaper tends to increase its consumption. However, in the context of healthcare, some argue that healthcare cannot be analysed like any other good. According to this view, demand for healthcare is driven by "needs" not by economic factors, making demand completely inelastic with respect to price (Einav & Finkelstein, 2018).

An alternative perspective is that health insurance improves people's health by facilitating timely and adequate medical care, and that this enhanced health may, in turn, reduce future healthcare utilisation.

Moral hazard in health insurance has been studied in its ex-ante form by Dave and Kaestner (2009), who found that health coverage may reduce preventive activities, because it lowers the expected costs of medical care. Specifically, the receipt of Medicare among the elderly was associated with a 39.7% decrease in the probability of engaging in physical exercise, an 18%

lower likelihood of quitting smoking, and a 14.8% increase in the probability of daily alcohol consumption.

While ex ante effects are important, most of the literature emphasises ex post moral hazard, defined as the responsiveness of healthcare demand to the price faced by the consumer, conditional on health status (Einav & Finkelstein, 2018).

Konetzka et al. (2019) analyse ex post moral hazard in long-term care insurance, which they define as the additional utilisation of long-term care services due to insurance coverage. They argue that moral hazard is socially inefficient when individuals consume more only because insurance makes out-of-pocket costs cheaper. By contrast, moral hazard can also be socially efficient: if people would still choose to buy the care even when given cash instead of insurance, then the extra use reflects better access to needed care rather than wasteful spending.

This distinction between efficient and inefficient moral hazard motivates policy discussions on optimal insurance design. For instance, Einav and Finkelstein (2018) suggest that a declining out-of-pocket price schedule is a natural way to balance two objectives optimally: on the one hand, reducing inefficient moral hazard through higher cost-sharing when spending is low and, on the other hand, providing risk protection by lowering cost-sharing when spending is high. Since the value of insurance increases with total expenditure, it is optimal to design policies that offer greater protection for larger health expenses.

3. DATA

3.1. HRS: Health and Retirement Study

The data employed in this study come from the freely available U.S. Health and Retirement Study (HRS), which has been ongoing since 1992. This biannual, longitudinal survey is representative of the American population aged 50 and older, including their spouses.

The HRS was the first study to collect epigenome-wide data on DNA methylation. This process was part of the 2016 HRS Venous Blood Study (VBS), managed by Hooper Holmes Wealth & Wellness. The 2016 wave included an in-home core interview and blood collection performed by a phlebotomist during a home visit (Crimmins et al., 2025). DNA methylation assays were conducted on a non-random subsample of 4,104 participants. After applying quality control procedures, 4,018 samples were retained for analysis.

Data are then generated using the assay of the Infinitum Methylation EPIC BeadChip[17] completed at the Advanced Research and Diagnostics Laboratory at the University of Minnesota (HRS Epigenetic Clocks File; Crimmins et al., 2024). The protocol follows a harmonised pipeline to perform quality control and eliminate unreliable probes (Crimmins et al., 2025). From this sample, thirteen epigenetic clocks have been trained on HRS data.

The HRS is the dataset that best serves the purposes of this study because it collects high-quality data on health insurance and epigenetic data, and it records the age of all household members, which is crucial for implementing the RDD method in this context.

3.2. Study Population and Sample Weights

Considering that this study is focused on understanding the relationships between Medicare insurance and DNA methylation-based epigenetic clocks, I utilised data from the Epigenetic Clocks File in conjunction with 2016 Core Interviews from the Health and Retirement Study.

The epigenetic clocks were developed using data from the 2016 Venous Blood Study in which respondents with at least one valid venous blood measurement were assigned a sampling weight (VBSI16WGTRA[18]) to account for differential participation probabilities.

Indeed, since the HRS sample is a multi-stage area probability sample of the United States, with oversamples of Blacks and Hispanics, unbiased estimates of population parameters require the use of these sampling weights (HRS Tracker Early Data Description, 2022).

---

[17] Genome-wide methylation analysis tool targeting ~930K unique methylation sites in the most biologically significant regions of the human methylome (Illumina).

[18] 2016 VBS weight for analyses including DNA methylation and epigenetic clocks, telomeres, and homocysteine (HRS Tracker Early Data Description, 2022).

These weights were derived by adjusting the 2016 HRS sample weight based on the likelihood of providing a valid venous blood sample, then post-stratified by age, sex, and race/ethnicity. The final weights available in the HRS Tracker File should be applied in analyses involving DNA methylation and epigenetic clocks.

3.3. Health Insurance in HRS

Health insurance information was collected through self-reports. Participants in the core interview answered questions regarding their insurance status in every survey (HRS 2016 Data Description, Section N "Health Services & Insurance"). Questions included, for example: "*Are you currently covered by Medicare health insurance?"* Responses were coded as "Yes", "No", "Do Not Know" or "Refused". For participants over 65, interviewers asked, *"Can you tell us why you are not covered by Medicare?"*; while those under 65, but eligible, were asked, *"Can you tell us more about the reason why you are eligible?"*. These last two were open-ended questions. Additional items also covered the different Medicare option plans.

3.4. Descriptive Statistics

Before data cleaning and final sample selection, the 2016 Epigenetic Clocks dataset was merged with the 2016 Core Interview files, incorporating information from Demographics (Section B), Physical Health (Section C), Health Services and Insurance (Section N), Employment (Section J), as well as the HRS 2022 Early Release Tracker File (Version 2).

The descriptive analysis presented in this section is based on the merged dataset, which includes 4,018 individuals from the Epigenetic Clocks File.

All data processing and analyses presented here were performed using the R statistical software environment. The analysis provides an overview of the distribution of key demographic and socioeconomic characteristics, namely:

- Age
- Gender
- Nativity (Born in the U.S.)
- Race/Ethnicity
- Medicare Coverage
- Medicare Advantage
- Private Insurance Coverage
- Dental Coverage

- Health Status
- Highest Level of Education
- Employment Status
- Marital Status

The main descriptive statistics for these variables offer a general characterisation of the study population before implementing the final inclusion criteria.

The sample has a mean age of 69.4 years, with a predominance of female respondents (58.5%). The majority of individuals were born in the United States, and most identify as White. Approximately 66% of the sample has Medicare coverage, and about 50% of the sample members are retired.

In addition, 41.8% are enrolled in a Medicare Advantage plan, 52% hold private insurance, and 43% have dental coverage. The largest share of respondents holds a high school diploma, and the majority are married.

Observations without sample weights, as well as information on Medicare and age, were eliminated, resulting in a final sample of 3,874 individuals.

3.5. Epigenetic Clocks Included in the Analysis

Among the thirteen available clocks in the HRS, four were selected: two first-generation clocks (the Hannum and the Horvath clocks), trained to predict chronological age; and two second-generation clocks (PhenoAge[19] and GrimAge), trained to predict lifespan/functional stage. In addition, a composite measure was constructed based on one procedure presented by Menta et al. (2025), referred to as MEGA Clock. This clock was obtained by aggregating the four clocks using a weighted index approach, as described by Anderson (2008), to reduce measurement error and improve efficiency. Specifically, the clocks were combined into a weighted sum, with weights corresponding to the row sums of the inverse variance-covariance matrix. Generally, a clock is assigned greater weight if it captures less variance or provides non-redundant information. In this specification, GrimAge receives approximately 50% of the total weight.

[19] Here labelled as “Levine”.

3.6. Age Measurement & Outcome Variable

The HRS offers information on respondents' age at the time of the interview, along with the exact year and month of both the interview and the respondent's birth. This allows age to be measured in different units, including monthly, bimonthly, quarterly, semi-annually, or yearly. Following Christelis et al. (2020), age is measured here in bimonthly intervals rather than in whole years, to achieve greater precision in identifying observations close to the cutoff.

Indeed, measuring the running variable in units that are too narrow may lead to excessively noisy estimates, while using intervals that are too vast risks including observations that are further off from the discontinuity threshold (Christelis et al., 2020).

Additionally, since multiple individuals can share the same age value, the running variable is technically discrete and exhibits mass points. To be more precise, a mass point is a value where the probability mass function[20] (PMF) is greater than zero. However, when the discrete score has many mass points, meaning that there is a sizable number of distinct values, it is considered to have an approximately continuous score. This approach is considered common in practice (Cattaneo et al., 2023), which supports the assumption of continuity of the running variable's density at the cutoff. Accordingly, I rely on the standard RDD methods, which require the continuity assumption, to estimate causal effects around the threshold.

Finally, chronological age will then be included in the regression model to compute Epigenetic Age Acceleration, which will serve as the outcome variable. It is computed as the residual from a regression of biological age – estimated through epigenetic clocks – on chronological age. In this case, the resulting residuals are expressed in years to remain consistent with the measurement of epigenetic clocks. The analysis relies on the assumption that, in the absence of Medicare, the EAA trend would be expected to follow a smooth trajectory with no "jump" at the cutoff point.

[20] The PMF assigns probabilities to each possible outcome in a discrete distribution.

4. METHODOLOGY

4.1. Correlation and Causation: The Problem of Causality

This analysis aims to establish a causal effect of Medicare on epigenetic age. From a statistical perspective, when estimating the causal effect of an intervention or a treatment, researchers face the so-called "fundamental problem of causal inference": for any given unit, we cannot observe both potential outcomes under treatment and control at the same time. Let $D_i = 1$ indicate that the unit is exposed to the treatment and $D_i = 0$ otherwise. Denote $Y_i(0)$ and $Y_i(1)$ the pair of potential outcomes for unit i, where $Y_i(0)$ is the outcome without exposure to the treatment and $Y_i(1)$ is the outcome under treatment.

If we could observe both outcomes for the same units, the individual treatment effect would be:

$$\Delta_i = Y_i(1) - Y_i(0)$$

However, because it is not possible to observe the pair simultaneously, Δ_i cannot be identified at the individual level. Therefore, causal inference typically focuses on average causal effects, most commonly the Average Treatment Effect (ATE), defined as:

$$\tau = E[Y_i(1) - Y_i(0)]$$

The observed outcome for each unit can be written as:

$$Y_i = (1 - D_i) \cdot Y_i(0) + D_i \cdot Y_i(1) = \begin{Bmatrix} Y_i(0) & if & D_i = 0 \\ Y_i(1) & if & D_i = 1 \end{Bmatrix}$$

To correctly identify the causal effect, units must be as good as randomly assigned to treatment and control groups. Random assignment ensures that the only systematic difference across groups is the treatment itself, which implies: $Y_i(1), Y_i(0) \perp D_i$.

However, in many observational studies, this assumption rarely holds without adjustments. Therefore, it is replaced by the weaker conditional independence or unconfoundness assumption:

$$Y_i(1), Y_i(0) \perp D_i, X_i$$

Where X_i is a vector of covariates. Adjusting for X_i allows identification of the ATE, but this identification entirely depends on the validity of the underlying assumptions of the chosen method.

4.2. Identification Strategy

When relying on observational data, a careful research design is needed to tackle the multiple sources of endogeneity. Indeed, the endogenous status of health insurance in observational studies prevents drawing practical causal inferences. Therefore, the use of appropriate techniques permits the exploitation of the exogenous variation created by natural experiments (Levy & Meltzer, 2008).

This study develops an identification strategy to provide causal estimates of the effects of health insurance on biological age, using data from the Health and Retirement Study. Building on the fact that the transition to Medicare coverage – except for those eligible earlier – occurs abruptly at age 65, there is a clearly defined threshold for Medicare eligibility. This makes it well-suited for applying a Regression Discontinuity Design (RDD) framework (Card et al., 2008). The RDD approach uses age trends before 65 to predict utilisation beyond this age and attributes any sudden deviations from these trends occurring at age 65 to Medicare coverage (McWilliams, 2009). Indeed, Medicare leads to a sharp increase in the use of healthcare services.

As has been just established, the basic idea behind the RDD method is that eligibility for Medicare medical services is determined, at least in part, by the value of a forcing variable – individual age – relative to the fixed threshold of 65 years (Christelis et al., 2019).

However, as illustrated in Figure 6, the probability of having Medicare does not shift abruptly from 0 to 1 at age 65.

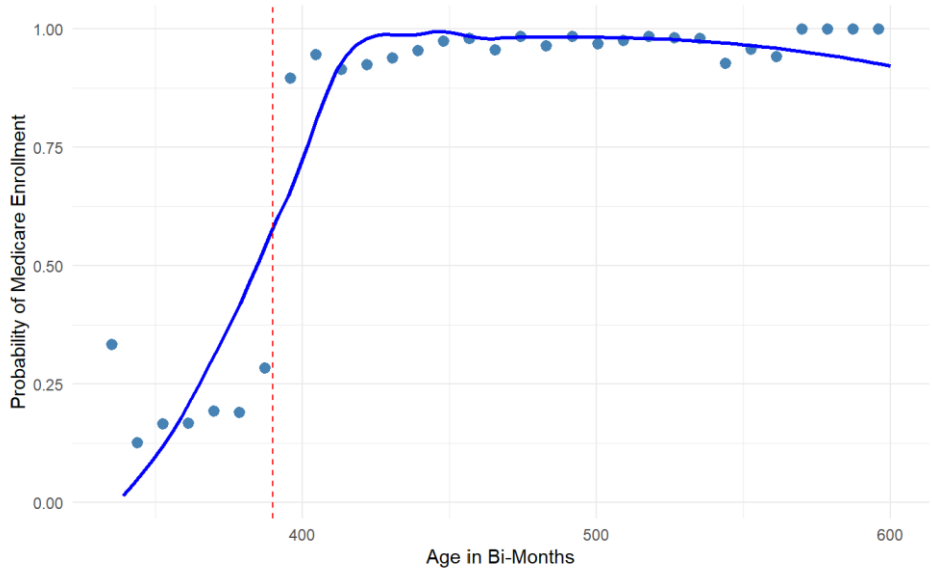

Figure 6: Medicare Enrollment vs Age in Bi-Months

Actually, approximately 17.1% of the individuals in this sample have Medicare coverage[21] before reaching age 65, indeed, there are additional conditions that allow respondents to become eligible at ages younger than 65, as illustrated in the previous section.
Additionally, the "jump" in the enrollment probability is 61.4%. Given that there is no one-to-one correspondence between the treatment (Medicare coverage) and the forcing variable (age), the appropriate approach is a Fuzzy Regression Discontinuity Design (FRD).

4.3. Regression Discontinuity Design: General Framework

To provide a clearer understanding of the Regression Discontinuity Design (RDD), with particular attention to the Fuzzy Regression Discontinuity method, this section outlines its key principles and framework.
According to Cattaneo et al. (2024), a RDD relies on three elements: a running variable X_i (also referred to as the score or forcing variable), a cutoff c, and a treatment rule that assigns units to treatment if their running variable exceeds the cutoff and to control otherwise: $T_i = 1[X_i > c]$. The binary variable D_i denotes whether the unit i actually receives the treatment.
The assignment rule creates a discontinuous change in the probability of the treatment at the cutoff. As long as units cannot manipulate their score, those just below the threshold serve as valid counterfactuals for those just above it. Hence, RDD relies on the assumption that individuals on either side of the cutoff are comparable in all respects except for treatment status, thereby approximating randomisation. Such designs typically occur in those circumstances when an administrative or organisational rule, such as the age threshold for Medicare enrolment, determines eligibility.

In the simplest case of Regression Discontinuity, the Sharp RDD, treatment assignment and treatment received coincide ($T_i = D_i$). By contrast, the FRD case allows for non-compliance: $T_i \neq D_i$. This means that some individuals do not take up the treatment even when assigned ($D_i = 0$ despite $T_i = 1$) or receive it without being assigned ($D_i = 1$ despite $T_i = 0$). Thus, while the assignment rule remains $T_i = 1[X_i > c]$, compliance is imperfect: at the cutoff, the probability of receiving the treatment increases discontinuously, but by less than one (Cattaneo et al., 2024). In the end, the FRD estimate is obtained as the ratio between the discontinuity in the outcome and the discontinuity in the probability of the treatment received:

[21] Medicare coverage is here considered as a binary variable (0 = No, 1 = Yes), excluding respondents who answered "Do not know" (DK) or "Refused" (RF).

$$\tau_{FRD} = \frac{\lim_{x \downarrow c} E(Y_i|X_i = x) - \lim_{x \uparrow c} E(Y_i|X_i = x)}{\lim_{x \downarrow c} E(D_i|X = x) - \lim_{x \uparrow c} E(D_i|X = x)} = \frac{\tau_Y}{\tau_D}$$

The analogy here is to the distinction between a randomised trial with perfect and imperfect compliance, where in the latter case we can instrument actual take-up with eligibility. As in that case, the Fuzzy Regression Discontinuity isolates the impact of compliers – those whose treatment status is altered by crossing the cutoff, but would not have changed otherwise. This is akin to the Local Average Treatment Effect (LATE) interpretation of the standard instrumental variable estimator. However, as the expression above makes clear, the FRD parameter is local in another sense as well, as it only captures an average treatment effect for compliers *at the cutoff*, defined as:

$$\tau_{FRD} = E[Y_i(1) - Y_i(0)|i \; is \; a \; complier, X_i = c]$$

4.4. Estimation in Fuzzy RDD

This lack of deterministic assignment introduces endogeneity, as the actual treatment status is correlated with unobserved factors, resulting in biased comparisons. This issue is addressed by relying on an instrumental variable (IV) setup (Angrist & Pischke, 2008).

In this IV framework, $T_i = 1[X_i > c]$ serves as an instrument for the treatment assignment D_i, as it captures the discontinuous change in treatment probability at the threshold, which provides exogenous variation for the identification of the Local Average Treatment Effect.

Indeed, the indicator variable T_i is exogenous, as it affects the outcome only through its effect on the treatment, and it also strongly correlates with the treatment D_i itself.

Additionally, in the Regression Discontinuity Design, there is no need for covariates beyond the running variable, as it fully captures the assignment mechanism.

Formally, the τ_{FRD} can be estimated through a two-stage least squares (2SLS) procedure:

First stage (the parameter τ_D)*:* $D_i = \pi_0 + \pi_1 Z_i + f(X_i) + \varepsilon_i$

Second stage (the parameter τ_{FRD}): $Y_i = \beta_o + \beta_1 \widehat{D}_i + f(X_i) + u_i$

Here, $f(X_i)$ denotes a smooth function of the running variable (age), π_1 measures the discontinuity in treatment probability, and β_1 identifies the LATE for compliers at the cutoff.

In R, the *rdrobust* package facilitates the estimation, specifically through the primary function rdrobust, which implements a bias-corrected robust inference procedure. This offers bias-corrected robust confidence intervals for average treatment effects at the cutoff for Regression Discontinuity Designs.

4.5. Practical Implications and Bandwidth Selection

The FRD has important implications for identification and interpretation.

First, the standard assumption of unconfoundness $Y_i(1), Y_i(0) \perp D_i, X_i$ cannot be applied (Imbens & Lemieux, 2008), because in this setting, the treatment assignment T_i differs from the treatment received D_i. Identification instead relies on the discontinuous change in treatment probability at the cutoff, which allows estimation of the effect only for compliers.

Second, this feature limits the external validity of FRD estimates. Both Sharp and Fuzzy designs identify a Local Average Treatment effect at the cutoff, but the FRD restricts it even further to the subpopulation of compliers. Without strong assumptions justifying extrapolation to other subpopulations, these designs never allow the researcher to estimate the overall average effect of the treatment. In that sense, the design has only a limited degree of external validity.

Finally, the characteristics of the running variable and the nature of the treatment are critical for identification. In the Medicare setting, age served as the running variable and cannot be manipulated, which strengthens credibility. At the same time, treatment effects are unlikely to appear instantaneously at eligibility, but instead take place over time, meaning the observed discontinuity in the outcome is likely to be attenuated. Consequently, the observed causal effect in outcomes at the cutoff may understate the longer-term impact of Medicare coverage (Christelis et al., 2019).

Another key component of Regression Discontinuity is the choice of the bandwidth (h), which determines how close observations must be to the cutoff to be included in the estimation. Bandwidth selection involves a trade-off between bias and variance: a narrower bandwidth reduces bias by focusing on observations very close to the threshold, but increases variance due to the smaller sample size; conversely, a wider bandwidth includes more observations, reducing variance but potentially introducing bias.

According to Imbens and Lemieux (2008), bandwidth choice can be guided by a cross-validation procedure. However, more recent literature has introduced faster methods that automatically select optimal bandwidths.

In particular, the *rdbwselect* function of the *rdrobust* package[22] (Calonico et al., 2014) selects bandwidths by minimising the asymptotic mean squared error (AMSE).

To assess the sensitivity of the results, it is standard practice to re-estimate the model using bandwidths that are half or twice the size of the originally selected ones.

[22] Specifically for this analysis, treatment estimates are obtained with *rdrobust*, leaving the automatic selection of the optimal bandwidth to the function. The figures instead use, for each specific clock, the optimal bandwidth computed with *rdbwselect*.

5. EMPIRICAL ANALYSIS

5.1. Tests and Caveats

5.1.1. Manipulation Test

As previously noted, age is an immutable characteristic, as individuals cannot manipulate whether they fall below or above the cutoff. This "already verified" assumption simplifies identification by ruling out the possibility of manipulation in this context.

Nonetheless, it is worth reporting the core idea behind manipulation testing, which can detect non-random sorting of units around the cutoff. Cattaneo et al. (2018) outline that, within the RDD framework, a manipulation test seeks to formally determine whether the density of observations is discontinuous at a known cutoff. Such discontinuities are usually interpreted as evidence of self-selection or non-random sorting of units into control and treatment groups.

If agents can influence the running variable, treatment assignment may no longer be as good as random at the threshold (McCrary, 2007). On the other hand, in the absence of manipulation, the density of observations should be continuous near the cutoff. Testing for manipulation, therefore, involves estimating the density on either side of the threshold and conducting a hypothesis test for discontinuity. It must be noted that, in some cases, "false positive" jumps are also plausible. For instance, a discontinuity density may show up in survey data as an artefact of sampling design, suggesting caution in the interpretation of positive results.

McCrary (2007) introduced an approach based on the nonparametric local-polynomial density estimator of Cheng et al. (1997), which requires pre-binning[23] the data and introduces additional tuning parameters. Cattaneo et al. (2018) improved on this by developing a manipulation test that does not require pre-binning and is constructed on kernel functions.

Fitzgerald (2025) offers a more recent contribution, arguing that the conventional approach – testing for continuity of the density – can be misleading. In particular, an underpowered estimate may be substantially large even if it is not statistically significantly different from zero, leading to biased inferences. His methodology extends the *DCdensity* procedure (McCrary, 2007) by incorporating bootstrap methods for finite-sample, cluster-robust inference, thereby improving the detection of meaningful manipulation near the cutoff.

For this analysis, manipulation testing was conducted manually, following the procedure presented by McCrary (2007). Although the commands developed by Cattaneo et al. (2018) are designed explicitly for manipulation density testing, they do not allow for the incorporation of

[23] Data binning is a data pre-processing technique used to group a set of continuous values into a smaller number of "bins".

sample weights in R. Therefore, the density test was computed as a Wald test of the null hypothesis that the discontinuity equals zero.

Formally:

$$\hat{\theta} = log\ \hat{f}^{+}\ (c) - log\ \hat{f}^{-}\ (c);$$

$$\widehat{SE}(\hat{\theta}) = \sqrt{\frac{\widehat{SE}(\hat{f}^{+})^2}{(\hat{f}^{+})^2} + \frac{\widehat{SE}(\hat{f}^{-})^2}{(\hat{f}^{-})^2}}$$

$$z = \frac{\hat{\theta}}{\widehat{SE}(\hat{\theta})}\ ;\ p = 2\{1 - \Phi(|z|)\}$$

The two-sided p-value is 0.73, confirming there is no manipulation in the running variable. Figure 7 reports the graphical representation of the density test. Despite being statistically insignificant, the histogram suggests a mild irregularity just to the left of the threshold. In particular, the bars immediately preceding the cutoff are lower than the surrounding ones, creating a small gap.

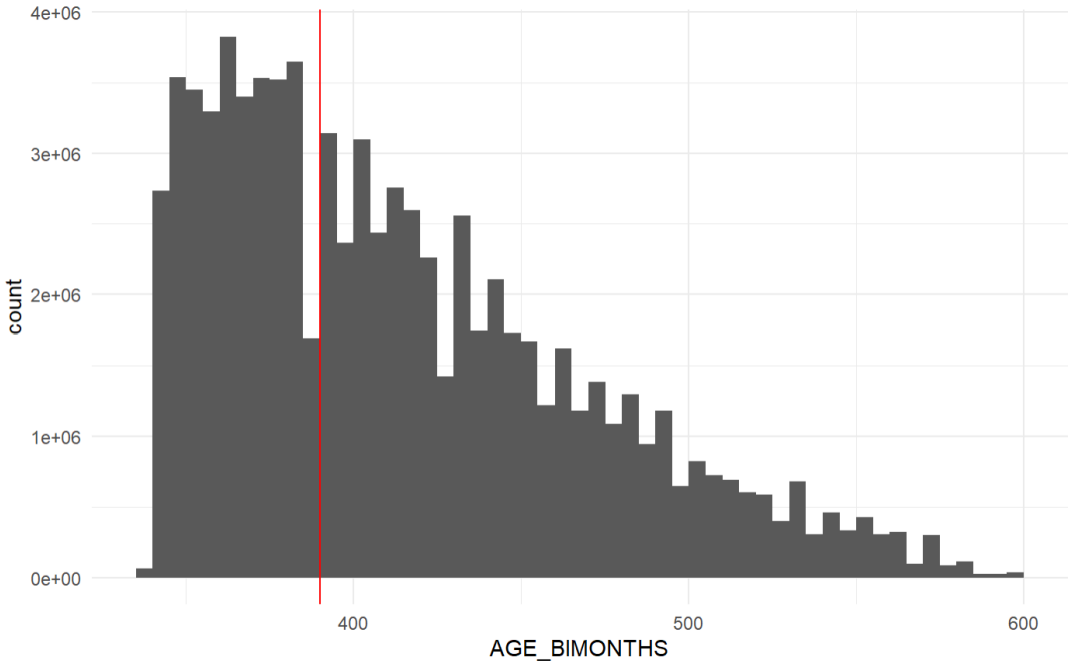

Figure 7: Density Test – Graphical Representation

5.1.2. Covariate Balance Test

Covariate balance is used to assess whether the distribution of covariates is similar across treatment levels. Indeed, under proper randomisation, covariate balance is expected to occur automatically. However, in observational data, balance is not guaranteed and must be tested. In

the RD design, pre-treatment characteristics should not exhibit systematic discontinuities at the cutoff, meaning they should be balanced immediately before and after the threshold.
Assessing balance involves comparing the distributions of covariates between treated and control units. This is typically done by examining univariate balance summary statistics for each covariate; however, there is no single best way to assess balance or weight these statistics. The underlying characteristics of the outcome data-generating process determine the appropriate degree and structure of covariate balance.

For this analysis, covariate balance was evaluated using the *rdrobust* function. Indeed, although it is primarily designed to estimate the LATE, its functionality can be extended to covariate balance. Using *rdrobust* to test for such discontinuities ensures that balance checks are carried out in the same local neighbourhood of the cutoff as the main causal estimation, employing the exact local polynomial specification and bandwidth selection procedures. This approach avoids misleading results from global mean comparisons (see CRAN Project – Assessing Balance for details on implementation in R). Using the automatic bandwidth selection, the covariates were found to be balanced around the cutoff.
Only the education level displayed a p-value slightly below 0.10 (0.093), yet it remained above the conventional 5% significance threshold.
Figure 8 in the Appendix reports the graphical representations of the results.

5.1.3. Graphical Inspection

Prima facie evidence on the effect of interest can be gathered by a graphical inspection of the evolution of the outcome of interest across the relevant age cutoff. The optimal bandwidth was selected using the function rdbwdensity (Cattaneo et al., 2018), which provides data-driven methods for selecting the bandwidth based on asymptotic mean squared error (AMSE) minimisation, allowing for different bandwidths on either side of the cutoff.
This procedure was repeated for all five clocks. Visual inspection of scatter plots between biological age and chronological age reveals that all clocks exhibit a positive association. The graphs here are visualised at their optimal bandwidth.
For all clocks, we observe a consistently increasing relationship between biological and chronological age, consistent with the idea that individuals are expected to grow older biologically as their chronological age advances. Furthermore, a visible discontinuity is observed around the age threshold. Notably, Horvath exhibits no apparent shift, whereas Hannum and GrimAge show more pronounced discontinuity, with Levine falling intermediate.

However, this appears to be primarily driven by a potential blip in the sampling process, as reflected in a transient, sharp downward movement of the line immediately before the cutoff.

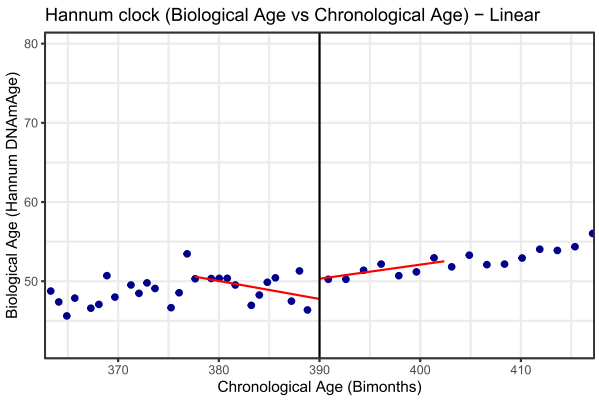

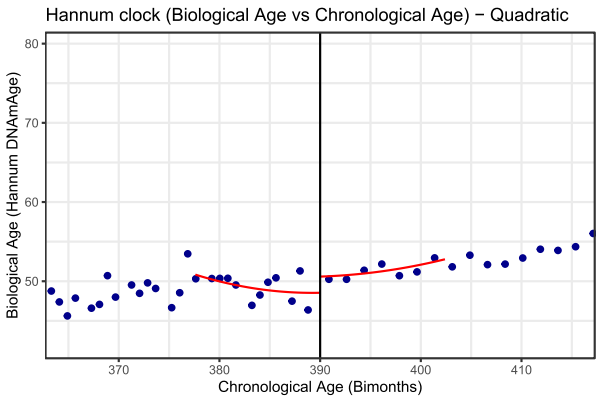

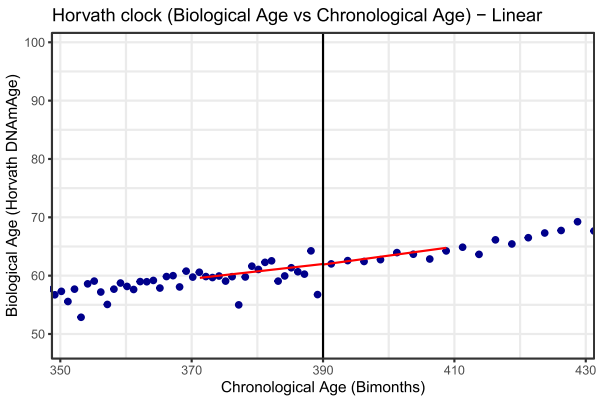

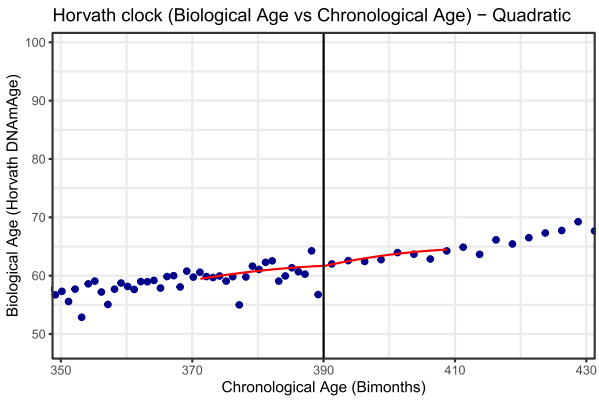

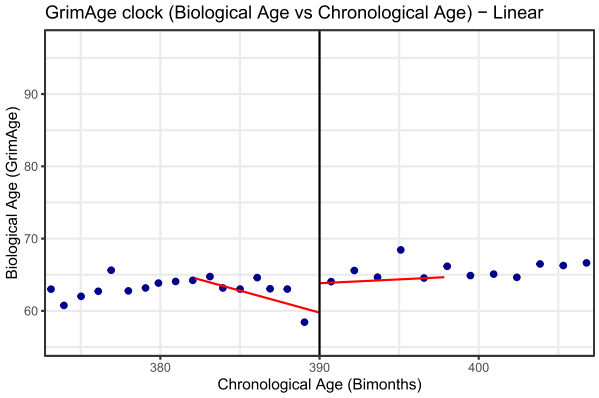

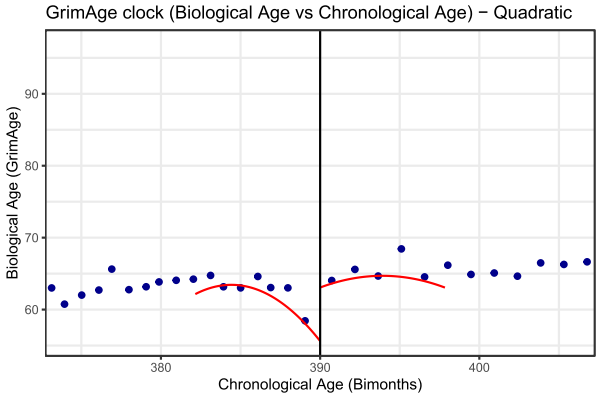

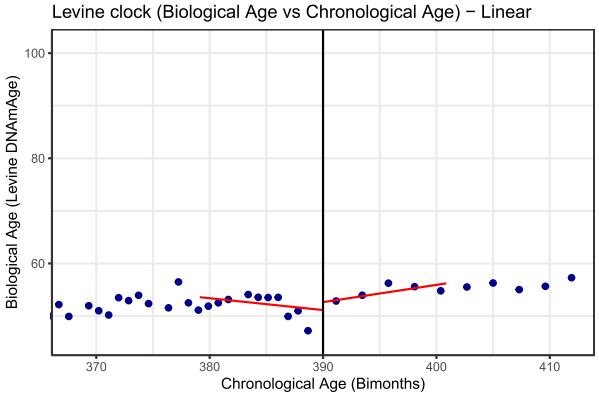

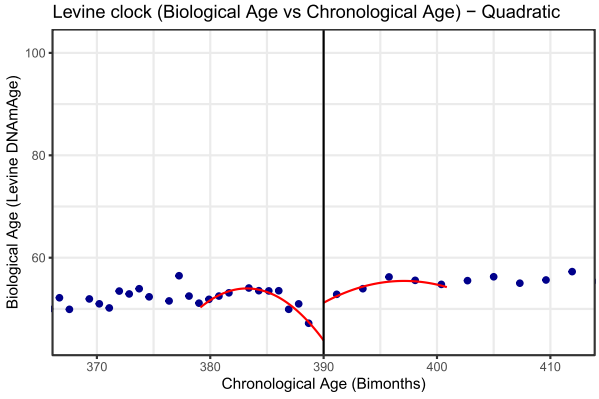

Figure 9: Relation between Chronological and Biological Age in Specific Clocks

The MEGA Clock also displays a similar irregularity just to the left of the cutoff, where the line appears to dip downward. This pattern aligns with the results from the GrimAge clock – that is heavily weighted in the construction of the MEGA clock – and with the small gap highlighted in the density test histogram, suggesting that the graphical evidence may partly reflect local sampling fluctuations rather than genuine biological discontinuities.

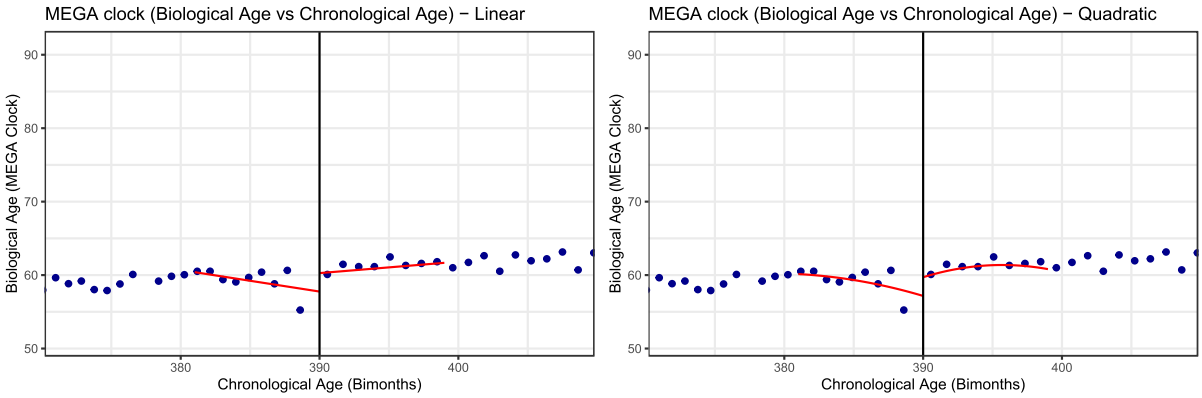

Figure 10: Relation between Chronological and Biological Age in MEGA Clock

5.2. Main Results

As previously discussed, the *rdrobust* package is designed to estimate the average causal effect in the Regression Discontinuity Design. It does not require the explicit implementation of the 2SLS procedure, but its internal logic is equivalent to the instrumental variables approach. Being a semi-parametric approach, its estimations are generally more robust and flexible.

An important consideration concerns the choice of the global polynomial order. When the order is too high, global polynomial regressions risk overfitting the data (Cattaneo et al., 2022). In line with Gelman and Imbens (2019), who recommend restricting to polynomials of at most second degree, the estimations were conducted using both linear and quadratic specifications. Tables 2 and 3, respectively, report the results for HRS clocks and MEGA Clock.

5.2.1. Clock-Specific Estimates

First Stage					
Clock	**Specification**	**BW**	**Estimate**	**p-value**	**CI**
Hannum	Linear	11.1	0.590***	0.000	[0.365, 0.826]
	Quadratic	21.3	0.641***	0.000	[0.387, 0.888]
Horvath	Linear	16.9	0.659***	0.000	[0.478, 0.844]
	Quadratic	20.1	0.630***	0.000	[0.353, 0.896]
GrimAge	Linear	9.2	0.583***	0.000	[0.320, 0.855]
	Quadratic	13.2	0.500**	0.021	[0.071, 0.864]
Levine	Linear	11.8	0.605***	0.000	[0.385, 0.833]
	Quadratic	14.6	0.517***	0.006	[0.140, 0.847]
Treatment Effect					
Clock	**Specification**	**BW**	**Estimate**	**p-value**	**CI**
Hannum	Linear	11.1	3.711	0.164	[-1.797, 10.584]
	Quadratic	21.3	4.505	0.114	[-1.177, 10.999]
Horvath	Linear	16.9	-0.034	0.949	[-4.785, 4.480]
	Quadratic	20.1	0.777	0.797	[-6.453, 8.400]
GrimAge	Linear	9.2	6.775**	0.038	[0.450, 15.607]
	Quadratic	13.2	10.992*	0.055	[-0.279, 27.375]
Levine	Linear	11.8	4.045	0.314	[-4.511, 14.059]
	Quadratic	14.6	10.789	0.158	[-5.079, 31.261]

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 2: Clock-Specific Estimates

First-stage results indicate that the RDD is well identified, as the instrument T_i is strong and relevant across all the specifications. All clocks display a similar coefficient for the discontinuity in the instrument (around 0.6), confirming that the probability of receiving the treatment shifts clearly at the cutoff.

Regarding the treatment effect, the effects are generally not statistically significant. Significance emerges only for GrimAge, where the positive coefficient suggests that epigenetic age tends to increase after the 65-year-old cutoff, despite the introduction of Medicare. This finding aligns with the existing literature, which suggests that epigenetic age changes typically capture modest deviations (Porter et al., 2022). Additionally, the sign of the coefficients suggests that changes in biological aging are unlikely to materialise in the short run, but may instead emerge over a longer horizon. This is consistent with the idea that biological aging responds only gradually to improved access to healthcare, indicating that the effect of Medicare is likely to be and cumulative.

5.2.2. MEGA Clock

First Stage					
Clock	**Specification**	**BW**	**Estimate**	**p-value**	**CI**
MEGA	Linear	13.5	0.637***	0.000	[0.436, 0.841]
	Quadratic	13.8	0.497**	0.015	[0.090, 0.851]
Treatment Effect					
Clock	**Specification**	**BW**	**Estimate**	**p-value**	**CI**
MEGA	Linear	13.5	3.267*	0.067	[-0.253, 7.641]
	Quadratic	13.8	6.434	0.161	[-3.146, 18.987]

*** p < 0.01, ** p < 0.05, * p < 0.10.

Table 3: MEGA Clock Estimates

For the MEGA clock, first-stage results are in line with those of the other clocks, showing a strong and relevant instrument. The treatment effect is positive and marginally significant in the linear specification.

5.3. Donut Regression Discontinuity

As already noted, the graphical inspections suggest potential local irregularities around the cutoff. To mitigate this concern, a Donut Regression Discontinuity Design (Barreca et al., 2011) was implemented.

In empirical applications, the Donut RDD is often employed as a robustness or falsification exercise to support the credibility of the RDD assumption, as it repeats estimation and inference after excluding observations within a specific range of the treatment threshold (Noack & Rothe, 2023).

Beyond its role as a robustness check, the approach is also motivated by concerns that systematic sorting of units or other data irregularities in some neighbourhood of the cutoff may bias the estimation of the treatment effect (Noack & Rothe, 2023).

Donut designs assume that potential manipulation is confined to a specific region around the threshold from which all observations are excluded. In practice, this involves dropping all individuals within a distance *d* of the cutoff (Dowd, 2020). While the drawback of this method is that it removes data at the threshold – including those observations that would otherwise identify the LATE – and may lead to wider confidence intervals, it remains consistent with the underlying motivation for RD designs: rely on a relatively narrow sample that provides more credible identification of unbiased treatment effects.

Noack and Rothe (2023) further note that the design is commonly applied symmetrically, using the same bandwidth h and donut size d on both sides of the threshold, though asymmetric specifications are also feasible.

In this present analysis, given that the observed irregularity was located in the pre-65 region, an asymmetric Donut RDD was applied, excluding only the observations immediately to the left of the cutoff, while preserving enough observations nearby to avoid compromising the estimation of the LATE.

In the interval 380-400, each value of the running variable (bimonthly age) was considered as a separate bin, with its corresponding number of observations counted. Bins with fewer than 15[24] observations were labelled as weak. Starting from the bin immediately before the cutoff, I identified the sequence of consecutive weak bins. The Donut length L, defined as the distance from the cutoff to the last consecutive weak bin, was capped at a maximum of four bins to prevent data loss. In this case, $L = 2$, leading to the removal of 389 and 388, corresponding to 4 months of data.

To streamline computations across the large number of models, the analysis was restricted to the MEGA Clock. The graphical results[25] are presented in Figure 11, and the estimates are listed in Table 4. With respect to the evidence in Figure 10 and Table 3, they display insignificant and minor treatment effects. Overall, the evidence is consistent with null effects.

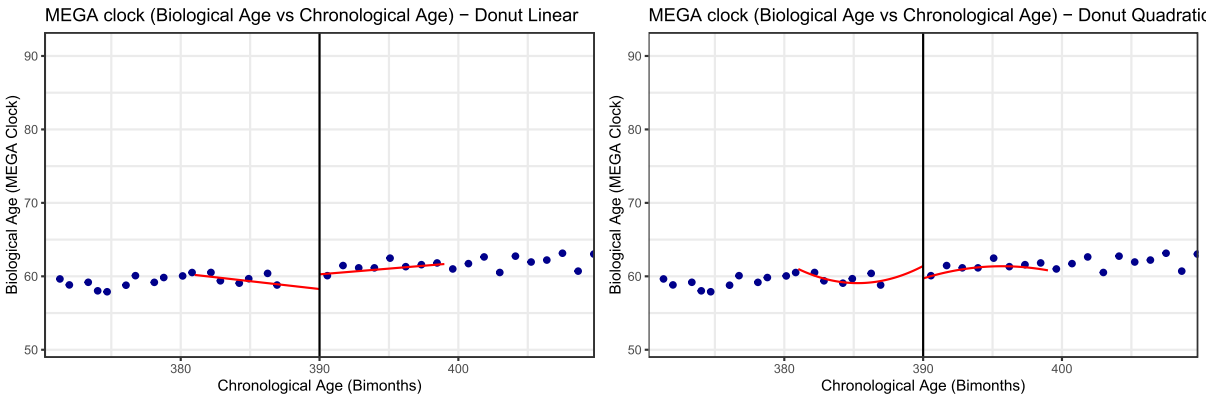

Figure 11: MEGA Clock on Donut Left Sample

[24] The threshold of 15 was chosen based on the observed distribution of observations: the median bin count is 17, and the first quartile is 13-14, making 15 a reasonable benchmark to exclude bins with insufficient mass.

[25] Graphical inspection in the Donut sample is done using the bandwidth selected previously for the entire sample.

First Stage					
Clock	**Specification**	**BW**	**Estimate**	**p-value**	**CI**
MEGA	Linear	9.0	0.665***	0.001	[0.327, 1.339]
	Quadratic	9.0	0.833	0.184	[-0.960, 4.997]
Treatment Effect					
Clock	**Specification**	**BW**	**Estimate**	**p-value**	**CI**
MEGA	Linear	9.0	1.978	0.893	[-13.382, 15.357]
	Quadratic	9.0	1.187	0.109	[-7.691, 76.731]

*** p < 0.01, ** p < 0.05, * p < 0.10.

Table 4: MEGA Clock Estimates for Donut Left Sample

5.4. Robustness Checks

Using the linear specification, the following robustness checks were performed to validate the analysis.

- Bandwidth sensitivity: estimations with half the optimal bandwidth.
- Covariates inclusion: addition of pre-treatment covariates, which Medicare does not influence. These are: Gender, Race/Ethnicity, Marital Status, Degree, Born in the USA and Employment Status.
- Combined check: estimation with half bandwidth and added covariates.

Results are reported in Table 5 and generally confirm the main findings: no significant effect. Only GrimAge and Levine become marginally statistically significant when covariates are included.

Robustness Specifications					
Clock	**Specification**	**BW**	**Estimate**	**p-value**	**CI**
Hannum	Half BW	5.5	7.396	0.492	[-38.460, 79.928]
	Covariates	11.1	3.878	0.668	[-10.506, 16.393]
	Half BW + Covariates	5.5	6.260	0.503	[-31.376, 63.970]
Horvath	Half BW	8.4	-1.150	0.989	[-16.995, 16.750]
	Covariates	16.9	-0.573	0.934	[-7.013, 7.627]
	Half BW + Covariates	8.4	-2.764	0.760	[-19.080, 13.936]
GrimAge	Half BW	4.6	20.883	0.210	[-40.944, 186.087]
	Covariates	9.2	6.730**	0.033	[1.159, 28.325]
	Half BW + Covariates	4.6	17.620	0.138	[-17.239, 124.783]
Levine	Half BW	5.9	17.313	0.232	[-21.763, 89.880]
	Covariates	11.8	4.508*	0.054	[-0.252, 32.366]
	Half BW + Covariates	5.9	18.148	0.204	[-19.449, 91.268]
MEGA	Half BW	6.7	5.966	0.165	[-7.518, 44.141]
	Covariates	13.4	3.294	0.132	[-1.653, 12.588]
	Half BW + Covariates	6.7	5.470	0.146	[-6.065, 40.741]

*** p < 0.01, ** p < 0.05, * p < 0.10.

Table 5: Robustness Checks – Half BW & Covariates

5.5. Stratifications

To understand whether effects are heterogeneous across different population subgroups, stratification analysis was carried out, distinguishing by the following dimensions:

- Gender: Male, Female.
- Ethnicity: White, Black, Other.
- Polygenic Risk Scores for Myocardial Infarction and Type II Diabetes. Polygenic Risk Scores are indicators of genetic liability for certain diseases. Here, they are used to capture heterogeneity in terms of the risk of poor health. This approach is adopted instead of using self-reported health, as in the data there is a slight variation in this variable prior to age 65, with virtually all respondents reporting good health.
- Dental Coverage: Yes, No. This analysis aims to explore whether additional forms of health insurance may play a role, specifically whether more comprehensive coverage contributes to a deceleration in epigenetic aging.

Because of the small sample size in the subsamples, the results revealed insignificant estimates throughout the board and are not reported here to save space.

5.6. Other Health Outcomes

To gain a better intuition behind the null effect of Medicare on Epigenetic Age Acceleration, the analysis was extended to other selected health outcomes at the cutoff.

The objective was to assess whether medical conditions remain unchanged at the eligibility threshold, confirming that access to Medicare does not produce immediate effects on health status. This would support the interpretation that changes in biological aging are unlikely to materialise in the short run, but may instead emerge over a longer horizon.

According to a 2018 US national survey, 88% of adults aged 65 and older reported at least one chronic condition, and 64% reported multiple chronic conditions (Watson et al., 2025).

For this study, the medical conditions considered were heart conditions, cancer (excluding skin) and diabetes (registered in the 2016 Core Interview), recognised mainly as some of the age-related diseases and consistent with the Centres for Disease Control and Prevention (CDC), which identifies these as the leading causes of death and disability in the United States.

Results in Table 6 confirm that both in the full sample and the Donut RDD, estimates for diabetes, cancer, and heart conditions are minimal and also statistically not significant (except

for a marginal significance in cancer). Medical conditions remain unchanged at the eligibility threshold, supporting the interpretation that Medicare does not have an immediate effect on these health outcomes.

Full sample			
Outcome	**Estimate**	**p-value**	**CI**
Diabetes	0.126	0.207	[-0.087, 0.401]
Cancer	0.112*	0.076	[-0.017, 0.347]
Heart condition	0.094	0.337	[-0.118, 0.345]
Donut sample			
Outcome	**Estimate**	**p-value**	**CI**
Diabetes	0.025	0.735	[-0.181, 0.257]
Cancer	0.011	0.694	[-0.270, 0.406]
Heart condition	0.153	0.202	[-0.098, 0.461]

*** p < 0.01, ** p < 0.05, * p < 0.10.

Table 6: Estimates for Chronic Conditions at the Cutoff

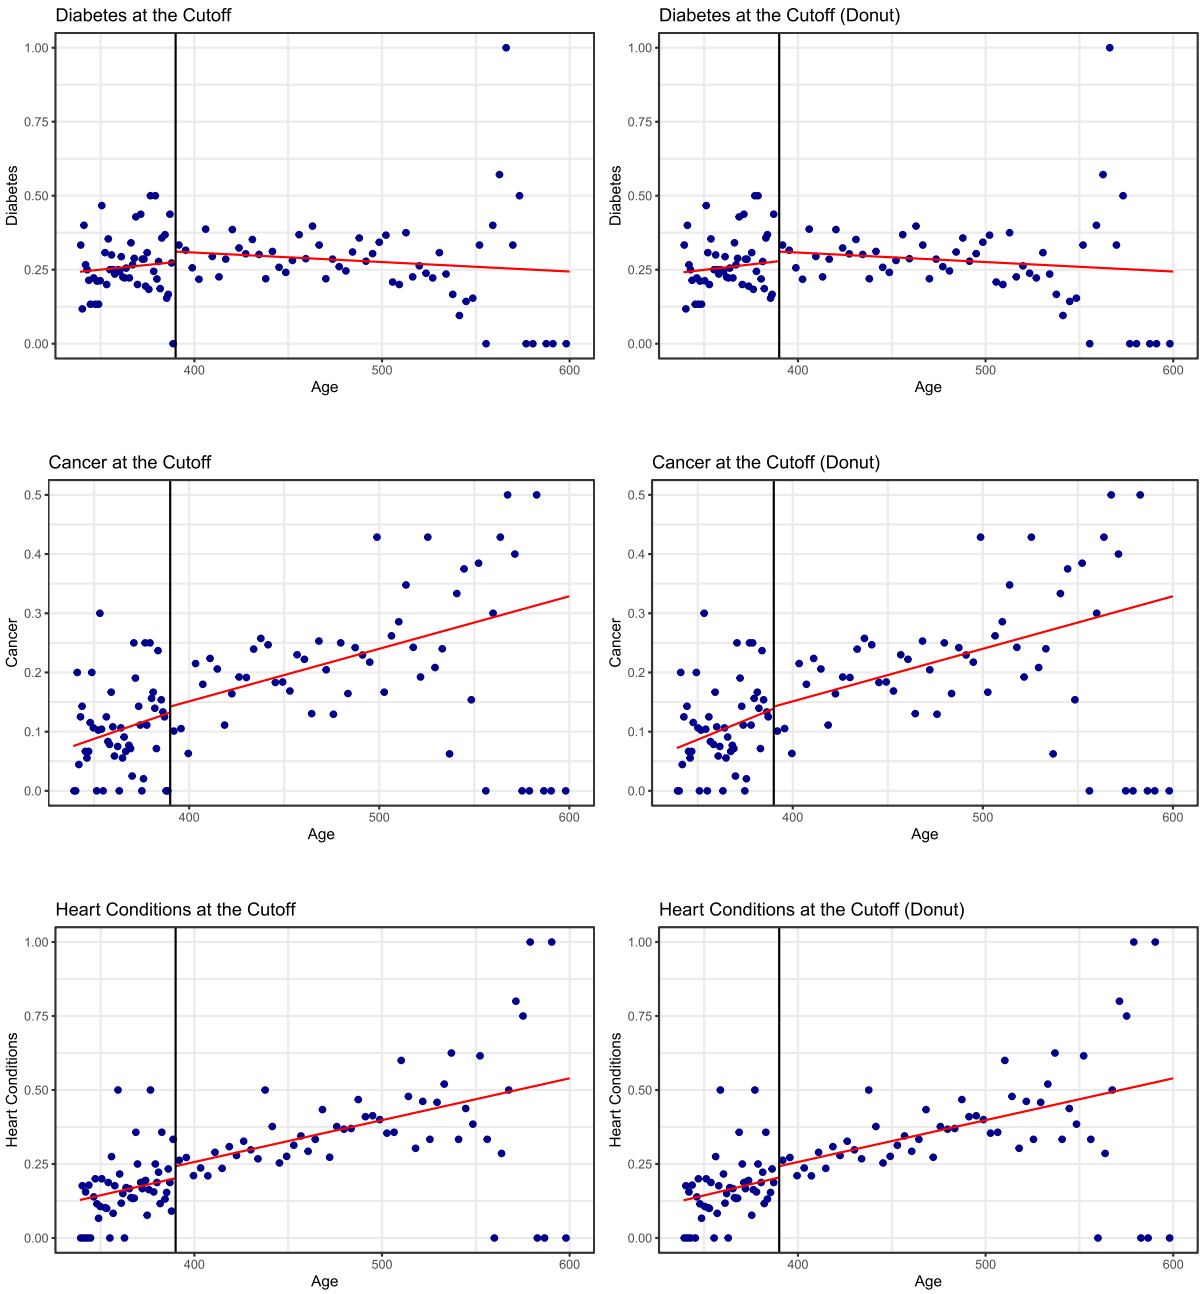

Figure 12: Chronic Conditions at the Cutoff

6. DISCUSSION

The present study investigated the causal effect of Medicare coverage on Epigenetic Age Acceleration, using epigenetic clocks as measures. These findings suggest that there is no evidence to support the notion that Medicare eligibility at age 65 affects age acceleration, at least in the short run. Around the cutoff point, estimates are statistically indistinguishable from zero, and the Donut specification that excludes observations near the threshold (to address a potential sampling blip) does not alter the conclusions. Subgroup analyses by gender, ethnicity, health status, polygenic scores and dental coverage produce similar null results. Consistently, indicators of chronic medical conditions show no discontinuity at 65.

Few studies have examined the link between Medicare and Epigenetic Aging Acceleration, making direct comparison challenging. A nationally representative NHANES analysis (Nwanaji-Enwerem et al., 2025) reports associations between insurance coverage and slower epigenetic aging, with more comprehensive plans contributing to potentially lower epigenetic age. The differences in the methodological approach (simple associations in a large population survey vs. an age-based Regression Discontinuity Design) contribute to explaining the different findings.

The null RDD findings here suggest that any effects may require time to materialise. Consistent with medical evidence, changes in EAA typically emerge over months or years. For instance, Nwanaji-Enwerem et al. (2021) tested metformin[26] and weight loss interventions in overweight postmenopausal breast cancer survivors as strategies for promoting anti-aging processes and improvements in overall health and lifespan. They found that six months of therapy was insufficient to produce measurable deceleration of epigenetic age, implying the need for longer-duration trials. Similarly, Esteban-Cantos et al. (2021) observed EAA reductions only two years after antiretroviral therapy initiation in adults with HIV. Fiorito et al. (2021) reported slower DNA methylation-based biomarkers of aging after a two-year diet and physical activity intervention trial. Although these settings differ from the one presented here, they underscore that epigenetic age responds slowly to exogenous changes.

It is plausible that the effect of gaining Medicare coverage is not statistically detectable at the cutoff. Indeed, a sizeable share of near-elderly adults is already insured prior to age 65, and the Affordable Care Act broadened coverage of preventive services in many non-Medicare plans; thus, the marginal change in access at 65 may be modest. Additionally, for those already insured, transitioning to Medicare often involves switching plans, learning about benefits, and

[26] Metformin is a medicine used to treat type 2 diabetes and gestational diabetes.

provider-network adjustments, which create short-run frictions that attenuate any immediate jump at the threshold. Consistent with this interpretation, Park and Marcondes (2024) report little to no change in diagnoses of chronic conditions at age 65.

In more detail, Park and Marcondes (2024) studied over 43,000 individuals from the Medical Expenditure Panel Survey (MEPS). They found little to no change in diagnoses of chronic conditions at the eligibility threshold. Similarly, Baicker et al. (2013), in their "The Oregon Experiment", did not find any significant improvement in measured physical health within the first two years. These results align with the null findings for medical conditions at the cutoff.

Moreover, Polsky et al. (2009) analysed eight biennial waves of the Health and Retirement Study to assess whether Medicare enrollment at age 65 influences the health trajectory of the near-elderly uninsured. They concluded that any change in overall health attributable to Medicare was small and not statistically significant. In addition, McWilliams et al. (2007) employed a longitudinal design from the HRS, reporting significant medium-run improvements in health among previously uninsured adults with cardiovascular disease or diabetes, with gaps reduced by about 50% by age 70, within five years of Medicare eligibility, rather than immediately at the cutoff. These findings suggest a potential role for Medicare in detecting and managing chronic conditions over the long term.

Interpreting the null at the cutoff thus requires a biological lens. Indeed, from a pure biological perspective, aging results from the accumulation of a wide variety of molecular and cellular changes over time, and these shifts are neither linear nor consistent. According to the World Health Organisation, most factors are beyond biology and genetics, being due to people's physical and social environments. Apart from their personal characteristics, individuals are influenced by their homes, neighbourhoods and communities, and even early-life environments have lasting effects on aging. This view is confirmed in two HRS-based studies: Hillmann et al. (2023) found that higher support and contact frequency with friends were prospectively associated with a slower Pace of Aging and a lower GrimAge, while greater contact with children and family was linked to lower Hannum age. Rentscher et al. (2023) showed that the absence of close relationships and lower support within existing relationships with family and friends relate to accelerated epigenetic aging in older adulthood. Together, perceived support and frequent contact with close others may have promoted healthier biological aging.

In line with the role of social ties, retirement is a transition that can reconfigure contact patterns and support, which may moderate the effects found in this analysis, even if evidence on epigenetic aging post-retirement is mixed, with effects varying by sociodemographic and health characteristics (Das, 2023). In this specific sample, about 50.6% are retired and 17% are not working, which could dilute average effects. Overall, the absence of a significant downward

discontinuity at the cutoff is compatible with minor, short-run effects and with impacts that, if any, unfold over a longer horizon.

Now the question may be how it is possible to slow down the ticking rate of epigenetic aging. The World Health Organisation emphasises a supportive physical and social environment, alongside individual-level approaches. For instance, examples of favourable contexts include the availability of safe and accessible public buildings and transportation, as well as places that are easy to navigate on foot.

From a biomedical standpoint, whether epigenetic aging can be prevented, slowed, or reversed remains uncertain (Mendelson, 2018).

Recent reviews (Galow & Peleg, 2022) highlight alteration of the environment (e.g. caloric restriction, intermittent fasting and altered diet), drug administration (Sir2 factor or Trichostatin A (TSA))[27] and genetic intervention – aimed at attenuating or counteracting specific age-associated alterations in the epigenome – as potential strategies where derived benefits may not be restricted to or manifested in total lifespan, but instead reflected in a prolonged healthy lifespan. The empirical results reported herein should be interpreted in light of some limitations. The lack of longitudinal data limits the analysis to a single cross-section, preventing the assessment of longer-run effects of Medicare on Epigenetic Age Acceleration. Panel data would allow to model within-person changes in clock measures relative to changes in chronological age. Additionally, the sample is relatively small compared to other HRS analyses, as it is constrained to respondents who participated in both the Venous Blood Study and reported Medicare information. A potential sampling blip near the cutoff exacerbates this issue. Furthermore, the scarcity of prior research on this specific question necessitated the development of a novel empirical approach.

These gaps could be addressed by future work extending the time horizon. Additionally, examining the effect of heterogeneity through social interrelationships and support networks would extend this study in line with recent literature. Additionally, the analysis could be expanded to link Medicare costs and utilisation, which was not feasible here due to limited access in HRS to relevant variables. Another promising avenue is to study epigenetic clocks under exogenous, high-intensity adversities, such as wars, recessions, natural disasters, displacement or conflict. Such settings may yield clearer signals because shocks are larger and more persistent, likely impacting epigenetic age modifications more significantly.

[27] Sir2 helps protect DNA by tightening chromatin, keeping the genome orderly. Trichostatin A (TSA) loosens that packaging to unlock helpful genes.

7. CONCLUSIONS

In conclusion, this study focused on the relationship between Medicare and Epigenetic Age Acceleration, motivated by the key role health insurance plays as a determinant of subjective well-being, as a potential tool for preventing future health risks and its connection to overall health outcomes across the life course. With an aging population, it becomes essential to understand how modifiable biological pathways, including those influenced by healthcare access, contribute to aging in order to inform sound and effective public decisions.

Among the 3,874 individuals from the 2016 Venous Blood Study, the estimated effects were small and statistically indistinguishable from zero. Specifically, although biological age increases with chronological age, Medicare eligibility shows no detectable short-run impact on epigenetic age. This is consistent with findings for medical conditions at the cutoff, which do not change in the immediate short run. In medical research, changes in Epigenetic Age Acceleration typically emerge over months or years.

Taken together, these results suggest that Medicare effects on epigenetic aging likely unfold over longer horizons, meaning that any measurable impact may require time to materialise.

Additionally, few studies have examined the link between Medicare and epigenetic aging acceleration, leaving clear room for further research. Some effects may be easier to detect in more adverse settings, where epigenetic age is more sensitive to sustained stressors or substantial environmental/behavioural changes.

Finally, these estimates are plausibly conservative: some residual confounding may remain, and epigenetic clocks are slow-moving biomarkers, making immediate impacts at age 65 unlikely.

Still, despite the limitations underlined above, this analysis highlights the emerging role of epigenetic clocks as promising tools not only for detecting changes In Epigenetic Age Acceleration but also for assessing broader aspects of health aging and longevity.

8. APPENDIX

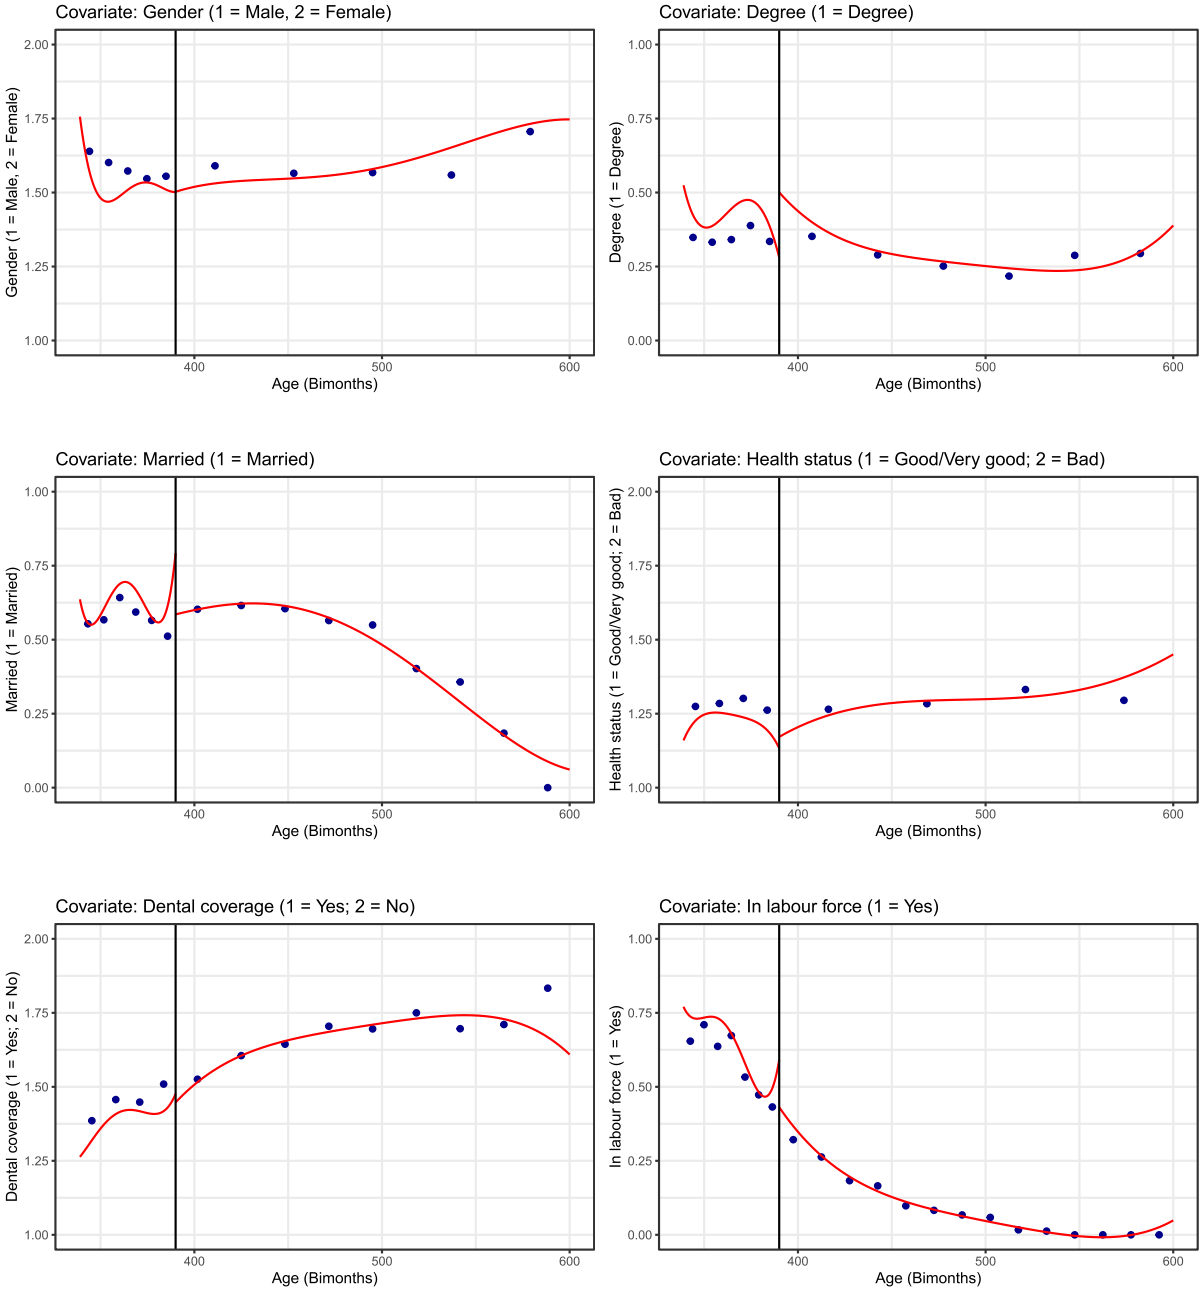
Covariate: Gender (1 = Male, 2 = Female)
Gender (1 = Male, 2 = Female)
Age (Bimonths)
Covariate: Degree (1 = Degree)
Degree (1 = Degree)
Age (Bimonths)
Covariate: Married (1 = Married)
Married (1 = Married)
Age (Bimonths)
Covariate: Health status (1 = Good/Very good; 2 = Bad)
Health status (1 = Good/Very good; 2 = Bad)
Age (Bimonths)
Covariate: Dental coverage (1 = Yes; 2 = No)
Dental coverage (1 = Yes; 2 = No)
Age (Bimonths)
Covariate: In labour force (1 = Yes)
In labour force (1 = Yes)
Age (Bimonths)

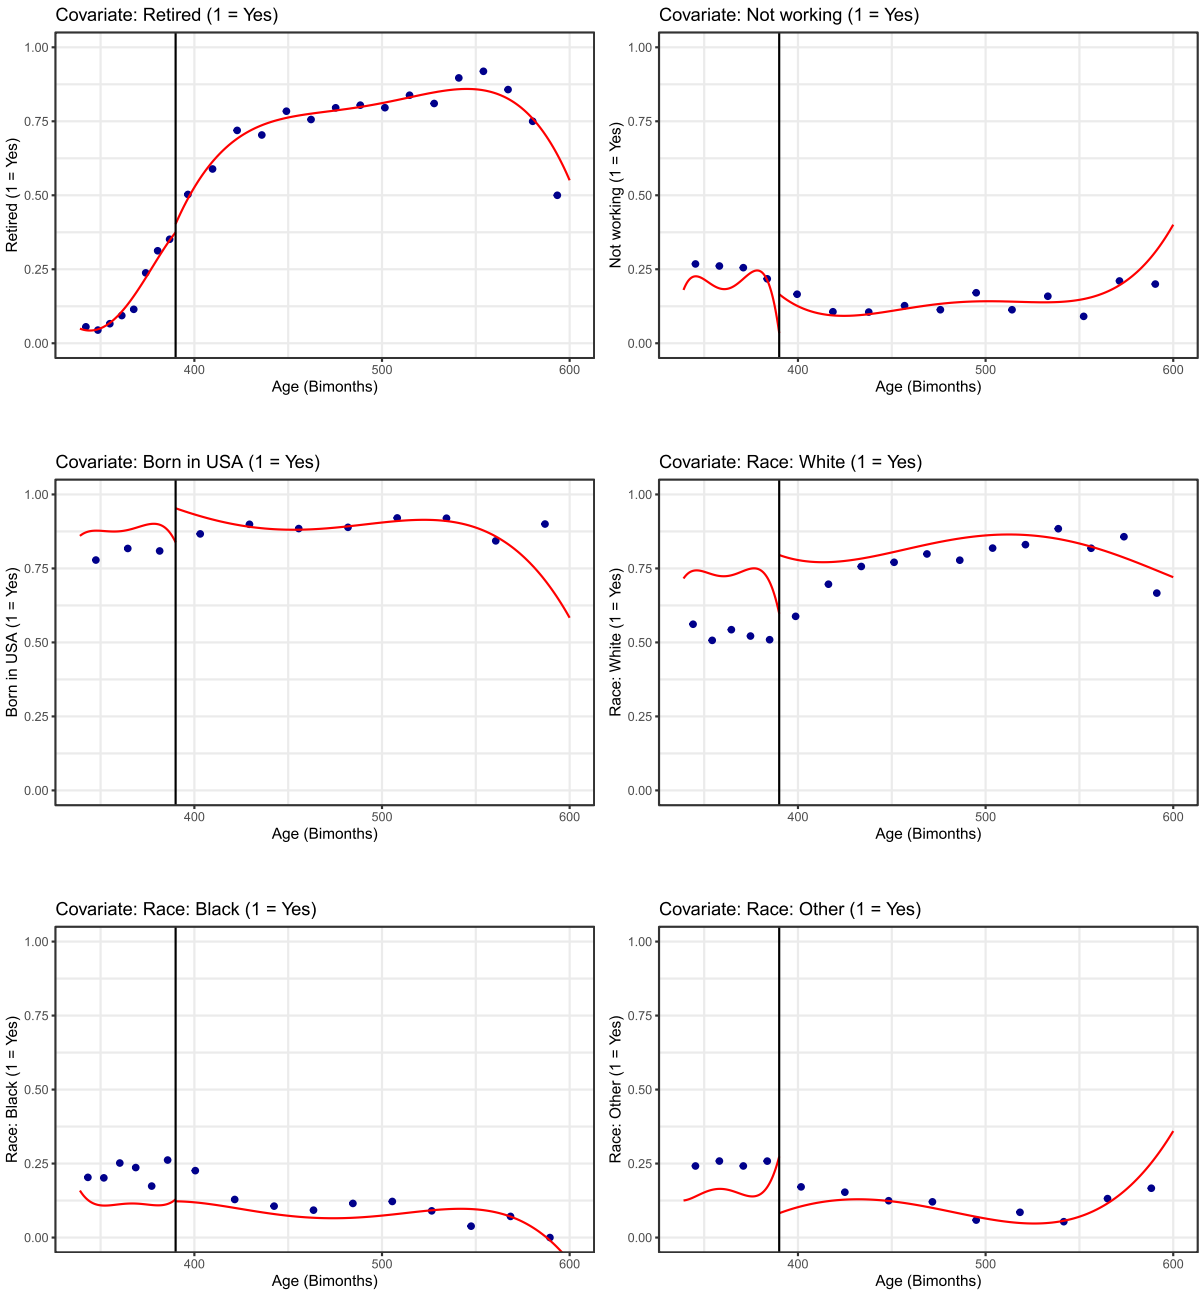

Figure 8: Covariate Balance Test – Graphical Representation

9. REFERENCES

23andMe Research Team, COGENT (Cognitive Genomics Consortium), Social Science Genetic Association Consortium, Lee, J. J., Wedow, R., Okbay, A., Kong, E., Maghzian, O., Zacher, M., Nguyen-Viet, T. A., Bowers, P., Sidorenko, J., Karlsson Linnér, R., Fontana, M. A., Kundu, T., Lee, C., Li, H., Li, R., Royer, R., … Cesarini, D. (2018). Gene discovery and polygenic prediction from a genome-wide association study of educational attainment in 1.1 million individuals. *Nature Genetics*, *50*(8), 1112–1121. https://doi.org/10.1038/s41588-018-0147-3

Acharjee, S. (n.d.). Mechanisms of DNA methylation and histone modifications. In *Progress in Molecular Biology and Translational Science* (Vol. 197, pp. 51–92). Vijai Singh, Indra Mani,.

Aggarwal, R., Yeh, R. W., Dahabreh, I. J., Robertson, S. E., & Wadhera, R. K. (2022). Medicare eligibility and healthcare access, affordability, and financial strain for low- and higher-income adults in the United States: A regression discontinuity analysis. *PLOS Medicine*, *19*(10), e1004083. https://doi.org/10.1371/journal.pmed.1004083

Alberts, B. (2015). *Molecular biology of the cell* (Sixth edition). Garland Science, Taylor and Francis Group.

Angrist, J. D. (n.d.). *Mostly Harmless Econometrics: An Empiricistís Companion*.

Baicker, K., Taubman, S. L., Allen, H. L., Bernstein, M., Gruber, J. H., Newhouse, J. P., Schneider, E. C., Wright, B. J., Zaslavsky, A. M., & Finkelstein, A. N. (2013). The Oregon Experiment—Effects of Medicaid on Clinical Outcomes. *New England Journal of Medicine*, *368*(18), 1713–1722. https://doi.org/10.1056/NEJMsa1212321

Barreca, A. I., Lindo, J. M., & Waddell, G. R. (n.d.). *NBER WORKING PAPER SERIES*.

Bertoni, M., Cavapozzi, D., Pasini, G., & Pavese, C. (2025). The causal effect of working from home on mental health of 50+ Europeans. *The Journal of the Economics of Ageing*, *31*, 100560. https://doi.org/10.1016/j.jeoa.2025.100560

Bocklandt, S., Lin, W., Sehl, M. E., Sánchez, F. J., Sinsheimer, J. S., Horvath, S., & Vilain, E. (2011). Epigenetic Predictor of Age. *PLoS ONE*, *6*(6), e14821. https://doi.org/10.1371/journal.pone.0014821

Branigan, A. R., McCallum, K. J., & Freese, J. (2013). Variation in the Heritability of Educational Attainment: An International Meta-Analysis. *Social Forces*, *92*(1), 109–140. https://doi.org/10.1093/sf/sot076

Burga, A., & Lehner, B. (2012). Beyond genotype to phenotype: Why the phenotype of an individual cannot always be predicted from their genome sequence and the environment that they experience. *The FEBS Journal*, *279*(20), 3765–3775. https://doi.org/10.1111/j.1742-4658.2012.08810.x

Calonico, S., Cattaneo, M. D., & Titiunik, R. (2014). Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs: Robust Nonparametric Confidence Intervals. *Econometrica*, *82*(6), 2295–2326. https://doi.org/10.3982/ECTA11757

Calonico, S., Farrell, M. H., & Cattaneo, M. D. (2025, July 19). *Package 'rdrobust'*.

Cameron, A. C., & Trivedi, P. K. (n.d.). *Microeconometrics: Methods and Applications*.

Card, D., Dobkin, C., & Maestas, N. (2008). The Impact of Nearly Universal Insurance Coverage on Health Care Utilization: Evidence from Medicare. *American Economic Review*, *98*(5), 2242–2258. https://doi.org/10.1257/aer.98.5.2242

Cattaneo, M. D., Idrobo, N., & Titiunik, R. (2024). *A Practical Introduction to Regression Discontinuity Designs: Extensions* (1st edn). Cambridge University Press. https://doi.org/10.1017/9781009441896

Cattaneo, M. D., Jansson, M., & Ma, X. (2018). Manipulation Testing Based on Density Discontinuity. *The Stata Journal: Promoting Communications on Statistics and Stata*, *18*(1), 234–261. https://doi.org/10.1177/1536867X1801800115

Cattaneo, M. D., Keele, L., & Titiunik, R. (2023). A guide to regression discontinuity designs in medical applications. *Statistics in Medicine*, *42*(24), 4484–4513. https://doi.org/10.1002/sim.9861

Cavaliere, R. (2018). Health insurance and the ongoing debate of quality and quantity. *Neuro-Oncology*, *20*(10), 1287–1288. https://doi.org/10.1093/neuonc/noy121

Centres for Disease Control and Prevention (CDC). (n.d.). *About Chronic Diseases*. https://www.cdc.gov/chronic-disease/about/index.html

Chen, B. H., Marioni, R. E., Colicino, E., Peters, M. J., Ward-Caviness, C. K., Tsai, P.-C., Roetker, N. S., Just, A. C., Demerath, E. W., Guan, W., Bressler, J., Fornage, M., Studenski, S., Vandiver, A. R., Moore, A. Z., Tanaka, T., Kiel, D. P., Liang, L., Vokonas, P., … Horvath, S. (2016). DNA methylation-based measures of biological age: Meta-analysis predicting time to death. *Aging*, *8*(9), 1844–1865. https://doi.org/10.18632/aging.101020

Chervova, O., Chernysheva, E., Panteleeva, K., Widayati, T. A., Hrbkova, N., Schneider, J., Maximov, V., Ryabikov, A., Tillmann, T., Pikhart, H., Bobak, M., Voloshin, V., Malyutina, S., & Beck, S. (2022). Evaluation of Epigenetic Age Acceleration Scores and Their Associations with CVD-Related Phenotypes in a Population Cohort. *Biology*, *12*(1), 68. https://doi.org/10.3390/biology12010068

Chervova, O., Panteleeva, K., Chernysheva, E., Widayati, T. A., Baronik, Ž. F., Hrbková, N., Schneider, J. L., Bobak, M., Beck, S., & Voloshin, V. (2024). Breaking new ground on human health and well-being with epigenetic clocks: A systematic review and meta-analysis of epigenetic age acceleration associations. *Ageing Research Reviews*, *102*, 102552. https://doi.org/10.1016/j.arr.2024.102552

Chiappori, P.-A., Jullien, B., Salanié, B., & Salanié, F. (2006). Asymmetric information in insurance: General testable implications. *The RAND Journal of Economics*, *37*(4), 783–798. https://doi.org/10.1111/j.1756-2171.2006.tb00057.x

Christelis, D., Georgarakos, D., & Sanz-de-Galdeano, A. (2020). The impact of health insurance on stockholding: A regression discontinuity approach. *Journal of Health Economics*, *69*, 102246. https://doi.org/10.1016/j.jhealeco.2019.102246

Collins, S. R., & Gupta, A. (2024). *The State of Health Insurance Coverage in the U.S.: Findings from the Commonwealth Fund 2024 Biennial Health Insurance Survey*. https://doi.org/10.26099/BYCE-QC28

Crimmins, E., Faul, J., Thyagarajan, B., & Weir, D. (n.d.). *Venous blood collection and assay protocol in the 2016 Health and Retirement Study 2016 Venous Blood Study (VBS)*.

Crimmins, E. M., Hernandez, B., Potter, C., Kim, J. K., Higgins-Chen, A., Kenny, R. A., O'Halloran, A. M., McGuinness, B., Smyth, L. J., Hill, C., Fiorito, G., Faul, J., McKnight, A. J., & McCrory, C. (2025). Epigenetic Clocks Relate to 4 Age-related Health Outcomes Similarly Across 3 Countries. *The Journals of Gerontology, Series A: Biological Sciences and Medical Sciences*, *80*(7), glaf036. https://doi.org/10.1093/gerona/glaf036

Crimmins, E. M., Klopack, E. T., & Kim, J. K. (2024). Generations of epigenetic clocks and their links to socioeconomic status in the Health and Retirement Study. *Epigenomics*, *16*(14), 1031–1042. https://doi.org/10.1080/17501911.2024.2373682

Crimmins, E. M., Thyagarajan, B., Levine, M. E., Weir, D. R., & Faul, J. (2021). 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. https://doi.org/10.1093/gerona/glab016

Das, A. (2023). Retirement and Epigenetic age Acceleration Among Older U.S. Adults. *Adaptive Human Behavior and Physiology*, *9*(3), 264–283. https://doi.org/10.1007/s40750-023-00221-2

Dave, D., & Kaestner, R. (2009). Health insurance and ex ante moral hazard: Evidence from Medicare. *International Journal of Health Care Finance and Economics*, *9*(4), 367–390. https://doi.org/10.1007/s10754-009-9056-4

Dowd, C. (2020). Donuts and Distant LATEs: Derivative Bounds for RD Extrapolation. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3641913

Duan, R., Fu, Q., Sun, Y., & Li, Q. (2022). Epigenetic clock: A promising biomarker and practical tool in aging. *Ageing Research Reviews*, *81*, 101743. https://doi.org/10.1016/j.arr.2022.101743

Dupont, C., Armant, D., & Brenner, C. (2009). Epigenetics: Definition, Mechanisms and Clinical Perspective. *Seminars in Reproductive Medicine*, *27*(05), 351–357. https://doi.org/10.1055/s-0029-1237423

Einav, L., & Finkelstein, A. (2018). Moral Hazard in Health Insurance: What We Know and How We Know It. *Journal of the European Economic Association*, *16*(4), 957–982. https://doi.org/10.1093/jeea/jvy017

Einav, L., Finkelstein, A., Ryan, S. P., Schrimpf, P., & Cullen, M. R. (2013). Selection on Moral Hazard in Health Insurance. *American Economic Review*, *103*(1), 178–219. https://doi.org/10.1257/aer.103.1.178

Esteban-Cantos et al. (2021). *Epigenetic age acceleration changes 2 years after antiretroviral therapy initiation in adults with HIV: a substudy of the NEAT001/ANRS143 randomised trial*.

Finkelstein, A., & McKnight, R. (2008). What did Medicare do? The initial impact of Medicare on mortality and out of pocket medical spending. *Journal of Public Economics*, *92*(7), 1644–1668. https://doi.org/10.1016/j.jpubeco.2007.10.005

Fiorito, G., Caini, S., Palli, D., Bendinelli, B., Saieva, C., Ermini, I., Valentini, V., Assedi, M., Rizzolo, P., Ambrogetti, D., Ottini, L., & Masala, G. (2021). DNA methylation-based biomarkers of aging were slowed down in a two-year diet and physical activity intervention trial: The DAMA study. *Aging Cell*, *20*(10), e13439. https://doi.org/10.1111/acel.13439

Fitzgerald, J. (2025). *Manipulation Tests in Regression Discontinuity Design: The Need for Equivalence Testing*. MetaArXiv. https://doi.org/10.31222/osf.io/2dgrp_v1

Fransquet, P. D., Wrigglesworth, J., Woods, R. L., Ernst, M. E., & Ryan, J. (2019). The epigenetic clock as a predictor of disease and mortality risk: A systematic review and meta-analysis. *Clinical Epigenetics*, *11*(1), 62. https://doi.org/10.1186/s13148-019-0656-7

Galow, A.-M., & Peleg, S. (2022). How to Slow down the Ticking Clock: Age-Associated Epigenetic Alterations and Related Interventions to Extend Life Span. *Cells*, *11*(3), 468. https://doi.org/10.3390/cells11030468

Garber, A. M. (2002). *Frontiers in health policy research. Vol. 5*. MIT.

Gelman, A., & Imbens, G. (2019). Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs. *Journal of Business & Economic Statistics*, *37*(3), 447–456. https://doi.org/10.1080/07350015.2017.1366909

Grechyna, D. (2023). Parenthood and political engagement. *European Journal of Political Economy*, *76*, 102238. https://doi.org/10.1016/j.ejpoleco.2022.102238

Harper, S. (2014). Economic and social implications of aging societies. *Science*, *346*(6209), 587–591. https://doi.org/10.1126/science.1254405

Harrington, S. E. (2010). The Health Insurance Reform Debate. *Journal of Risk and Insurance*, *77*(1), 5–38. https://doi.org/10.1111/j.1539-6975.2009.01345.x

Health and Retirement Study. (2019, December). *2016 Core, Data Description and Usage*.

Health and Retirement Study. (2024, November). *2022 Tracker, Data Description and Usage*.

Hillmann, A. R., Dhingra, R., & Reed, R. G. (2023). Positive social factors prospectively predict younger epigenetic age: Findings from the Health and Retirement Study. *Psychoneuroendocrinology*, *148*, 105988. https://doi.org/10.1016/j.psyneuen.2022.105988

Horvath, S. (n.d.). *Additional File 2 'Materials and Methods' for 'DNA methylation age of human tissues and cell types'*.

Horvath, S. (2013a). DNA methylation age of human tissues and cell types. *Genome Biology*, *14*(10), 3156. https://doi.org/10.1186/gb-2013-14-10-r115

Horvath, S. (2013b). DNA methylation age of human tissues and cell types. *Genome Biology*, *14*(10), 3156. https://doi.org/10.1186/gb-2013-14-10-r115

Horvath, S., & Raj, K. (2018). DNA methylation-based biomarkers and the epigenetic clock theory of ageing. *Nature Reviews Genetics*, *19*(6), 371–384. https://doi.org/10.1038/s41576-018-0004-3

Illumina Chip. (n.d.). https://emea.illumina.com/products/by-type/microarray-kits/infinium-methylation-epic.html

Imbens, G. W., & Lemieux, T. (2008). Regression discontinuity designs: A guide to practice. *Journal of Econometrics*, *142*(2), 615–635. https://doi.org/10.1016/j.jeconom.2007.05.001

Infoaging guides. (n.d.). *BIOMARKERS OF AGING*. https://www.afar.org/imported/AFAR_BIOMARKERS_OF_AGING_2016.pdf

Insler, M. (2014). *The Health Consequences of Retirement*. *1*.

Jablonka, E., & Lamb, M. J. (2002). The Changing Concept of Epigenetics. *Annals of the New York Academy of Sciences*, *981*(1), 82–96. https://doi.org/10.1111/j.1749-6632.2002.tb04913.x

James, G., Witten, D., Hastie, T., & Tibshirani, R. (2022). An introduction to statistical learning with applications in R: By Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani, New York, Springer Science and Business Media, 2013, $41.98, eISBN: 978-1-4614-7137-7. *Statistical Theory and Related Fields*, *6*(1), 87–87. https://doi.org/10.1080/24754269.2021.1980261

Jansen, A. G., Mous, S. E., White, T., Posthuma, D., & Polderman, T. J. C. (2015). What Twin Studies Tell Us About the Heritability of Brain Development, Morphology, and Function: A Review. *Neuropsychology Review*, *25*(1), 27–46. https://doi.org/10.1007/s11065-015-9278-9

Johnson, W., Turkheimer, E., Gottesman, I. I., & Bouchard, T. J. (2009). Beyond Heritability: Twin Studies in Behavioral Research. *Current Directions in Psychological Science*, *18*(4), 217–220. https://doi.org/10.1111/j.1467-8721.2009.01639.x

Jolliffe, I. T., & Cadima, J. (2016). Principal component analysis: A review and recent developments. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *374*(2065), 20150202. https://doi.org/10.1098/rsta.2015.0202

Jylhävä, J., Pedersen, N. L., & Hägg, S. (2017). Biological Age Predictors. *EBioMedicine*, *21*, 29–36. https://doi.org/10.1016/j.ebiom.2017.03.046

Kiefer, J. C. (2007). Epigenetics in development. *Developmental Dynamics*, *236*(4), 1144–1156. https://doi.org/10.1002/dvdy.21094

Kirkwood, T., & Austad, S. (2000). NATURE. *Why Do We Age?*, *408*.

Klemera, P., & Doubal, S. (2006). A new approach to the concept and computation of biological age. *Mechanisms of Ageing and Development*, *127*(3), 240–248. https://doi.org/10.1016/j.mad.2005.10.004

Konetzka, R. T., He, D., Dong, J., & Nyman, J. A. (2019). Moral hazard and long-term care insurance. *The Geneva Papers on Risk and Insurance - Issues and Practice*, *44*(2), 231–251. https://doi.org/10.1057/s41288-018-00119-1

Kullo, I. J., Lewis, C. M., Inouye, M., Martin, A. R., Ripatti, S., & Chatterjee, N. (2022). Polygenic scores in biomedical research. *Nature Reviews Genetics*, *23*(9), 524–532. https://doi.org/10.1038/s41576-022-00470-z

Kwon, J. M. (n.d.). *The Candidate Gene Approach*.

Lee, D. S., & Lemieux, T. (n.d.). *NBER WORKING PAPER SERIES*.

Lee, S., Nishimura, R., Burton, P., & McCammon, R. (n.d.). *HRS 2016 Sampling Weights*.

Lenart, A. (2012). *The Gompertz distribution and maximum likelihood estimation of its parameters—A revision* (No. WP-2012-008; 0 edn, p. WP-2012-008). Max Planck Institute for Demographic Research. https://doi.org/10.4054/MPIDR-WP-2012-008

Levine, M. E. (n.d.). *Additional File for 'An epigenetic biomarker of aging for lifespan and healthspan'*.

Levine, M. E. (2013). Modeling the Rate of Senescence: Can Estimated Biological Age Predict Mortality More Accurately Than Chronological Age? *The Journals of Gerontology Series A: Biological Sciences and Medical Sciences*, *68*(6), 667–674. https://doi.org/10.1093/gerona/gls233

Levine, M. E., Lu, A. T., Quach, A., Chen, B. H., Assimes, T. L., Hou, L., Baccarelli, A. A., Stewart, J. D., Li, Y., Whitsel, E. A., Wilson, G., Reiner, A. P., Aviv, A., Lohman, K., Liu, Y., & Ferrucci, L. (2018). *An epigenetic biomarker of aging for lifespan and healthspan*.

Levy, H., & Meltzer, D. (2008). The Impact of Health Insurance on Health. *Annual Review of Public Health*, *29*(1), 399–409. https://doi.org/10.1146/annurev.publhealth.28.021406.144042

Lichtenberg, F. R. (n.d.). *The Effects of Medicare on Health Care Utilization and Outcomes*.

Lipsky, M. S., & King, M. (2015). Biological theories of aging. *Disease-a-Month*, *61*(11), 460–466. https://doi.org/10.1016/j.disamonth.2015.09.005

Marioni, R. E., Shah, S., McRae, A. F., Chen, B. H., Colicino, E., Harris, S. E., Gibson, J., Henders, A. K., Redmond, P., Cox, S. R., Pattie, A., Corley, J., Murphy, L., Martin, N. G., Montgomery, G. W., Feinberg, A. P., Fallin, M. D., Multhaup, M. L., Jaffe, A. E., … Deary, I. J. (2015). DNA methylation age of blood predicts all-cause mortality in later life. *Genome Biology*, *16*(1), 25. https://doi.org/10.1186/s13059-015-0584-6

McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, *142*(2), 698–714. https://doi.org/10.1016/j.jeconom.2007.05.005

McCrory, C., Fiorito, G., O'Halloran, A. M., Polidoro, S., Vineis, P., & Kenny, R. A. (2022). Early life adversity and age acceleration at mid-life and older ages indexed using the next-generation GrimAge and Pace of Aging epigenetic clocks. *Psychoneuroendocrinology*, *137*, 105643. https://doi.org/10.1016/j.psyneuen.2021.105643

McWilliams, J. M. (2003). Impact of Medicare Coverage on Basic Clinical Services for Previously Uninsured Adults. *JAMA*, *290*(6), 757. https://doi.org/10.1001/jama.290.6.757

McWilliams, J. M., Meara, E., Zaslavsky, A. M., & Ayanian, J. Z. (2007). Health of Previously Uninsured Adults After Acquiring Medicare Coverage. *JAMA*, *298*(24), 2886. https://doi.org/10.1001/jama.298.24.2886

Mendelson, M. M. (2018). Epigenetic Age Acceleration: A Biological Doomsday Clock for Cardiovascular Disease? *Circulation: Genomic and Precision Medicine*, *11*(3), e002089. https://doi.org/10.1161/CIRCGEN.118.002089

Menta, G., Biroli, P., Mehta, D., D'Ambrosio, C., & Cobb-Clark, D. (n.d.). *Aggregating Epigenetic Clocks to Study Human Capital Formation*.

Michael McWILLIAMS, J. (2009). Health Consequences of Uninsurance among Adults in the United States: Recent Evidence and Implications. *The Milbank Quarterly*, *87*(2), 443–494. https://doi.org/10.1111/j.1468-0009.2009.00564.x

Michels, K. B. (Ed.). (2022). *Epigenetic Epidemiology*. Springer International Publishing. https://doi.org/10.1007/978-3-030-94475-9

Morales Berstein, F., McCartney, D. L., Lu, A. T., Tsilidis, K. K., Bouras, E., Haycock, P., Burrows, K., Phipps, A. I., Buchanan, D. D., Cheng, I., the PRACTICAL consortium, Martin, R. M., Davey Smith, G., Relton, C. L., Horvath, S., Marioni, R. E., Richardson, T. G., & Richmond, R. C. (2022). Assessing the causal role of epigenetic clocks in the development of multiple cancers: A Mendelian randomization study. *eLife*, *11*, e75374. https://doi.org/10.7554/eLife.75374

Murrell, A., Rakyan, V. K., & Beck, S. (2005). From genome to epigenome. *Human Molecular Genetics*, *14*(suppl_1), R3–R10. https://doi.org/10.1093/hmg/ddi110

National Cancer Institute. (n.d.). *Definition of 'Allele'*. https://www.cancer.gov/publications/dictionaries/genetics-dictionary/def/allele

NHGRI. (n.d.-a). *Definition of 'Genome'*. https://www.genome.gov/genetics-glossary/Genome

NHGRI. (n.d.-b). *Definition of 'Kilobase'*. https://www.genome.gov/genetics-glossary/Kilobase-kb

Noack, C., & Rothe, C. (2023). *Donut Regression Discontinuity Designs* (No. arXiv:2308.14464). arXiv. https://doi.org/10.48550/arXiv.2308.14464

Noack, C., & Rothe, C. (2024). Bias-Aware Inference in Fuzzy Regression Discontinuity Designs. *Econometrica*, *92*(3), 687–711. https://doi.org/10.3982/ECTA19466

Noah Greifer. (2025, May 29). *Assessing Balance*. https://cran.r-project.org/web/packages/MatchIt/vignettes/assessing-balance.html#summary.matchit

Nwanaji-Enwerem, J. C., Chung, F. F.-L., Van Der Laan, L., Novoloaca, A., Cuenin, C., Johansson, H., Bonanni, B., Hubbard, A. E., Smith, M. T., Hartman, S. J., Cardenas, A., Sears, D. D., & Herceg, Z. (2021). An epigenetic aging analysis of randomized metformin and weight loss interventions in overweight postmenopausal breast cancer survivors. *Clinical Epigenetics*, *13*(1), 224. https://doi.org/10.1186/s13148-021-01218-y

Nwanaji-Enwerem, J. C., Khodasevich, D., Gladish, N., Shen, H., Daredia, S., Needham, B. L., Rehkopf, D. H., & Cardenas, A. (2025). Health insurance and epigenetic aging: Trends in a

United States adult population. *SSM - Population Health*, *30*, 101806. https://doi.org/10.1016/j.ssmph.2025.101806

Okbay, A., Wu, Y., Wang, N., Jayashankar, H., Bennett, M., Nehzati, S. M., Sidorenko, J., Kweon, H., Goldman, G., Gjorgjieva, T., Jiang, Y., Hicks, B., Tian, C., Hinds, D. A., Ahlskog, R., Magnusson, P. K. E., Oskarsson, S., Hayward, C., Campbell, A., … Young, A. I. (2022). Polygenic prediction of educational attainment within and between families from genome-wide association analyses in 3 million individuals. *Nature Genetics*, *54*(4), 437–449. https://doi.org/10.1038/s41588-022-01016-z

Palaniappan, A., Blitzer, D., Takayama, H., & Sellke, F. W. (2022). Estimating the causal effect of the Medicaid expansion on heart transplant volume with a differences-in-differences model. *JTCVS Open*, *11*, 200–213. https://doi.org/10.1016/j.xjon.2022.06.011

Park, & Marcondes. (2024). The Effect of Medicare Eligibility on Diagnosis of Chronic Conditions. *The American Journal of Managed Care*, *30*(2), 96–100. https://doi.org/10.37765/ajmc.2024.89497

Park, S., Wadhera, R. K., & Jung, J. (2023). Effects of Medicare eligibility and enrollment at age 65 years on the use of h IGH-VALUE and l OW-VALUE care. *Health Services Research*, *58*(1), 174–185. https://doi.org/10.1111/1475-6773.14065

Polsky, D., Doshi, J. A., Escarce, J., Manning, W., Paddock, S. M., Cen, L., & Rogowski, J. (2009). The Health Effects of Medicare for the Near-Elderly Uninsured. *Health Services Research*, *44*(3), 926–945. https://doi.org/10.1111/j.1475-6773.2009.00964.x

Porter, H. L., Brown, C. A., Roopnarinesingh, X., Giles, C. B., Georgescu, C., Freeman, W. M., & Wren, J. D. (2021). Many chronological aging clocks can be found throughout the epigenome: Implications for quantifying biological aging. *Aging Cell*, *20*(11), e13492. https://doi.org/10.1111/acel.13492

Quach, A., Levine, M. E., Tanaka, T., Lu, A. T., Chen, B. H., Ferrucci, L., Ritz, B., Bandinelli, S., Neuhouser, M. L., Beasley, J. M., Snetselaar, L., Wallace, R. B., Tsao, P. S., Absher, D., Assimes, T. L., Stewart, J. D., Li, Y., Hou, L., Baccarelli, A. A., … Horvath, S. (2017). Epigenetic clock analysis of diet, exercise, education, and lifestyle factors. *Aging*, *9*(2), 419–446. https://doi.org/10.18632/aging.101168

Rackley, S. (2009). *Carbon Capture and Storage* (1st edn). Butterworth-Heinemann.

Rentscher, K. E., Klopack, E. T., Crimmins, E. M., Seeman, T. E., Cole, S. W., & Carroll, J. E. (2023). Social relationships and epigenetic aging in older adulthood: Results from the Health and Retirement Study. *Brain, Behavior, and Immunity*, *114*, 349–359. https://doi.org/10.1016/j.bbi.2023.09.001

Rinne, H. (n.d.). *The Hazard Rate – Theory and Inference*.

Sanderson et al. (2022). NATURE. *Mendelian Randomization*.

Santos, R. L. (n.d.). *Health Insurance Coverage in the United States: 2023*.

Saraswathy, & Ramalingam. (2016). Genome Sequencing Methods. In *Concepts and Techniques in Genomics and Proteomics* (1st edn).

Schmitz, L. L., & Duque, V. (2022). 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. https://doi.org/10.1073/pnas.2208530119

Simpson, D. J., & Chandra, T. (2021). Epigenetic age prediction. *Aging Cell*, *20*(9), e13452. https://doi.org/10.1111/acel.13452

Steve Horvath. (n.d.). *Clock Foundation*. https://clockfoundation.org/

Sugden, K., Hannon, E. J., Arseneault, L., Belsky, D. W., Corcoran, D. L., Fisher, H. L., Houts, R. M., Kandaswamy, R., Moffitt, T. E., Poulton, R., Prinz, J. A., Rasmussen, L. J. H., Williams, B. S., Wong, C. C. Y., Mill, J., & Caspi, A. (2020). Patterns of Reliability: Assessing the Reproducibility and Integrity of DNA Methylation Measurement. *Patterns*, *1*(2), 100014. https://doi.org/10.1016/j.patter.2020.100014

Sugrue, L. P., & Desikan, R. S. (2019). What Are Polygenic Scores and Why Are They Important? *JAMA*, *321*(18), 1820. https://doi.org/10.1001/jama.2019.3893

Survey Research Center. (2013). HRS Documentation Report—Epigenetic Clocks. *Genome Biology*, *14*(10), 3156. https://doi.org/10.1186/gb-2013-14-10-r115

Teschendorff, A. E., Marabita, F., Lechner, M., Bartlett, T., Tegner, J., Gomez-Cabrero, D., & Beck, S. (2013). A beta-mixture quantile normalization method for correcting probe design bias in Illumina Infinium 450 k DNA methylation data. *Bioinformatics*, *29*(2), 189–196. https://doi.org/10.1093/bioinformatics/bts680

Tracker of EAA Associations with Phenotype & Environmental Exposure (TEAPEE). (n.d.). https://ucl-medical-genomics.github.io/TEAPEE/

Unnikrishnan, A., Freeman, W. M., Jackson, J., Wren, J. D., Porter, H., & Richardson, A. (2019). The role of DNA methylation in epigenetics of aging. *Pharmacology & Therapeutics*, *195*, 172–185. https://doi.org/10.1016/j.pharmthera.2018.11.001

U.S. Department of Health and Human Services. (n.d.). *Medicare and Medicaid*. https://www.hhs.gov/answers/medicare-and-medicaid/index.html#medicare

van Wolferen. (n.d.). *Moral hazard in the insurance industry*.

Ware, E. B., Schmitz, L. L., Faul, J., Gard, A., Mitchell, C., Smith, J. A., Zhao, W., Weir, D., & Kardia, S. L. (2017). *Heterogeneity in polygenic scores for common human traits*. Genetics. https://doi.org/10.1101/106062

Watson, K. B., Wiltz, J. L., Nhim, K., Kaufmann, R. B., Thomas, C. W., & Greenlund, K. J. (2025). Trends in Multiple Chronic Conditions Among US Adults, By Life Stage, Behavioral Risk Factor Surveillance System, 2013–2023. *Preventing Chronic Disease*, *22*, 240539. https://doi.org/10.5888/pcd22.240539

WebMD. (n.d.). *HMO vs. PPO*. https://www.webmd.com/health-insurance/hmo-vs-ppo

Woolhandler, S., & Himmelstein, D. U. (2017). The Relationship of Health Insurance and Mortality: Is Lack of Insurance Deadly? *Annals of Internal Medicine*, *167*(6), 424–431. https://doi.org/10.7326/M17-1403

World Health Organization. (n.d.-a). *Ageing and health*. https://www.who.int/news-room/fact-sheets/detail/ageing-and-health

World Health Organization. (n.d.-b). *Ageing: Global population*. https://www.who.int/news-room/questions-and-answers/item/population-ageing

Zhu, M., & Zhao, S. (2007). Candidate Gene Identification Approach: Progress and Challenges. *International Journal of Biological Sciences*, 420–427. https://doi.org/10.7150/ijbs.3.420

Zurbuchen, R., Von Däniken, A., Janka, H., Von Wolff, M., & Stute, P. (2025). Methods for the assessment of biological age – A systematic review. *Maturitas*, *195*, 108215. https://doi.org/10.1016/j.maturitas.2025.108215