

RESEARCH PAPER

Preventive versus curative breakthroughs: shaping the distribution of age at death

Pablo Garcia-Sanchez and Olivier Pierrard

Département Économie et Recherche, Banque centrale du Luxembourg, Luxembourg

Corresponding author: Pablo Garcia-Sanchez; Email: pablo.garciasanchez@bcl.lu

(Received 24 January 2025; revised 18 June 2025; accepted 22 September 2025)

Abstract

How have preventive and curative medical breakthroughs shaped life expectancy and the dispersion of age at death in the United States over the past century? We address this question by developing a life-cycle model in which both health and lifespan are endogenous. The model distinguishes between preventive innovations, which reduce the incidence of disease, and curative advances, which lower mortality risks associated with existing health conditions. Our quantitative analysis shows that while both types of medical innovation have contributed to increased life expectancy since 1935, curative advances have been the primary driver of the decline in the dispersion of age at death. Medical innovations have also improved welfare – measured in terms of a consumption-equivalent metric – by an average of 0.11% per year, with curative advances representing the most significant contribution. These findings are robust across different scenarios and parametrization strategies.

Keywords: preventive care; curative care; age at death; life-cycle model

JEL classification: C60; D15; D81; I12; I18

1. Introduction

Over the past century, groundbreaking medical advances have transformed healthcare in the US (and in most developed countries). (a) Improvements in preventive care, such as vaccinations and cancer screenings, have reduced the spread of diseases, while (b) breakthroughs in curative treatments, like antibiotics and chemotherapy, have lowered mortality risk for a given health condition (left panel of Fig. 1). During this period, health outcomes improved dramatically: (i) life expectancy surged from just over

For useful comments and suggestions, we thank Patrick Fève, Paolo Guarda, Luca Marchiori, Alban Moura, Emmanuel Thibault and BCL colleagues. This paper should not be reported as representing the views of the BCL or the Eurosystem. The views expressed are those of the authors and may not be shared by other research staff or policymakers in the BCL or the Eurosystem.

© The Author(s), 2025. Published by Cambridge University Press in association with Université catholique de Louvain. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.




Figure 1. Selected medical advances and age at death in US: trends.

Notes: Left panel: Based on selected key medical advances detailed in Appendix A. Preventive advances and curative advances correspond to the trends (a) and (b) mentioned in the text. Right panel: Based on data from the Human Mortality Database.

60 years in 1935 to nearly 80 years in 2019, and (ii) the dispersion of age at death dropped from 25 years to 17 years (right panel of Fig. 1).

How did medical trends (a) and (b) contribute to health outcomes (i) and (ii)? The answer is not straightforward, as both medical advances may have differing effects. For instance, advancements in preventive care not only enhance the benefits of prevention but also encourage its use, improving overall health across the population. This should increase the mean age at death, though its impact on the dispersion of age at death remains uncertain. Meanwhile, curative breakthroughs weaken the link between an agent's health status and mortality risk, potentially narrowing the dispersion of age at death. However, this weaker link could also reduce agents' incentives to use preventive care, making the effect of improved curative care on the mean age at death ambiguous.¹

To address our question, we present a simple theoretical setup that considers the impact of both preventive and curative medical advances on the distribution of age at death. To the best of our knowledge, this is the first model to differentiate between these two types of innovations. Moreover, while prior research has focused on increased longevity (see e.g., Nichols, 2004; Acemoglu and Johnson, 2007; Chernew et al., 2016; Ewbank and Preston, 1990; Hu et al., 1995; Cutler and Miller, 2005; Ferrie and Troesken, 2008), the role of medical innovations in reducing the dispersion of age at death remains largely unexplored. This paper fills this gap, arguing that while both preventive and curative technologies have increased life expectancy, it is mainly curative innovations that have reduced the dispersion of age at death.

Ours is a simple life-cycle model in which both health and lifespan are endogenous. As individuals age, they experience adverse health events, which accumulate over time

¹This paper adopts a strict distinction between preventive and curative technologies. In practice, however, some medical advances may possess characteristics of both, contributing to both disease prevention and treatment. Our approach recalls the framework proposed by Murphy and Topel (2006), who differentiates between technologies that impact life expectancy and those that enhance quality of life. Section 2 elaborates on our relationship to Murphy and Topel (2006)'s work.

and are collectively referred to as the health deficit. This deficit incurs associated costs, reflecting real-world expenses such as long-term care (LTC) as well as curative and palliative treatments. However, individuals can slow the accumulation of health deficits by consuming preventive care, which includes measures to avoid diseases and risk factors (e.g., vaccination) or to detect diseases early (e.g., screenings). In line with recent empirical evidence suggesting that age is malleable (see Scott, 2023, and references therein), our model assumes that individuals do not die of old age per se but due to deteriorating health. Specifically, the time of death is treated as a random variable, with its probability distribution determined by the agent's health deficit. Simply put, higher deficits result in a higher probability of death. We prevent unjustifiably long lifespans by imposing an exogenous maximum health deficit; if an individual reaches it, she will pass away immediately. Such a maximum health deficit endogenously determines a maximum lifespan. Section 3 discusses the empirical validity of these assumptions underlying our theoretical framework.

The life-cycle structure is essential for studying how advances in health technologies have shaped the distribution of age at death. Without life-cycle optimization behavior, both technologies would have similar effects on mortality patterns. Indeed, improved curative care would not alter the level of health deficit but reduce the mortality risk associated with each level of deficit. In contrast, improved preventive care would not change the mortality risk at a given deficit level but lower the entire trajectory of deficit accumulation over time. Despite these differences, the resulting effect on the age-at-death distribution would be observationally equivalent. However, when prevention is endogenous over life, changes in the two technologies affect prevention choices, and hence the death distribution, differently.

We conduct our quantitative analysis in two steps. In the first step, we calibrate our life-cycle model – including the levels of preventive and curative technologies – to replicate 2019 U.S. data. In the second step, we retain the same calibration, except that we set the levels of both medical technologies (as well as income and health care prices which have also significantly evolved over the past century) to match those of 1935. We then examine how the distribution of age at death evolved between 1935 and 2019 and compare our results to the changes observed in the data.

More precisely, in the first step, we begin by selecting values for the model parameters using 2019 US data, prior to the COVID-19 pandemic. We estimate the relationship between an individual's health deficit and her instantaneous probability of death directly from the data. We use the National Health Interview Survey (NHIS) to calculate the average health deficit for individuals, assessing each age from 20 through 80. As in Mitnitski et al. (2002) and Hosseini et al. (2021), we construct our deficit measure as the ratio of accumulated adverse health effects to the total number of potential effects considered. Next, we obtain the probability of dying within the next year for individuals in this age range (20 to 80) from the Human Mortality Database. We then fit the functional form linking the probability of death to the health deficit used in our theoretical model, based on the empirical data. The remaining parameters are calibrated to match simple statistics on time-discounting, medical expenditures and the time profile of individual's health deficit.

Quantitatively, the model generates an age-at-death distribution consistent with empirical observations, despite none of its moments being targeted during parameter selection. For example, the model gives a life expectancy of 83.9 years (compared to 79.5 in the data), a standard deviation of 15 (compared to 16.7), a skewness of -1.9 (compared to -1.6), and a 0.3% proportion of individuals reaching age 105 (compared

to 0.5%). The model also captures the rising health expenditures-to-income ratio with age, which accelerates later in life due to high health deficits, in line with observed trends (De Nardi et al., 2016). Furthermore, in our setup, a 10% increase in income is associated with a one-year increase in life expectancy at age 40, which, though slightly higher than the roughly six-month increase estimated by Chetty et al. (2016), remains within a plausible range.

In our second step, we incorporate medical trends (a) and (b) into the model to assess how well it aligns with health outcomes (i)–(ii). Because these advances occurred alongside major socioeconomic shifts in the U.S., we also include (c) the sharp rise in the relative price of medical care and (d) the increase in real median income, followed by its later stagnation.

Quantifying medical progress is challenging, as measuring medical advances is notoriously difficult. To address this, we take a holistic approach, assuming that in 1935, when modern medicine was still in its early stages, the efficiency of preventive care was negligible, and individuals consumed none at that time. This allows us to derive time series for advances in both preventive and curative care. Incorporating these four trends enables the model to qualitatively capture health outcomes (i)–(ii). Life expectancy increases from 67 years in 1935 to 84 years in 2019, yielding a 17-year gain, compared to the 19 years observed in the data. As for the standard deviation of age at death, it decreases from 18 years to 15 years, reflecting a 3-year reduction, while the data shows an 8-year decline. In addition, the model replicates the surge in the share of GDP devoted to health care, as well as the rise in the use of preventive care.

Next, we perform a series of counterfactual exercises, introducing the trends cumulatively. First, we introduce advances in preventive care while holding the other three trends fixed at their 1935 levels. In this counterfactual scenario, life expectancy increases, as better preventive care improves health status across the population. This finding aligns with evidence on oral polio vaccination campaigns reducing child mortality in low- and middle-income countries (Nielsen et al., 2021). However, the dispersion of age at death increases, contrary to what is observed in the data. Offsetting the rise in the dispersion of age at death requires advances in curative care, which reduce mortality risk and lead to a more compressed age-at-death distribution. This finding aligns with evidence on penicillin's role in reducing infectious disease mortality variation in post-WWII Italy (Alsan et al., 2021). As expected, these reductions in mortality also contribute to higher life expectancy. Lastly, we incorporate economic trends (c) and (d) (price of medical care and median income), which have limited effects on the distribution of age at death but are crucial to explain the rise in health care's share of total resources. This aligns with empirical evidence pointing to rising healthcare prices as a key driver of the growing share of output devoted to healthcare spending (Anderson et al., 2003, 2019).

In sum, our model suggests that while both preventive and curative technologies have contributed to the increase in life expectancy in the US since 1935, it is curative innovations that have driven the reduction in the dispersion of age at death.

Finally, we examine the role of the four trends in the historical evolution of living standards, measured using a consumption-equivalent metric as in Jones and Klenow (2016). We find that medical advances increased welfare by an average of 0.11% per year, with curative advances making the most significant contribution. We also highlight the negative welfare impact of rising healthcare prices. This finding aligns with widespread public concern about the escalating cost of medical care, which has led many Americans to forgo treatment, as documented in the NHIS.

All these insights remain robust across a range of sensitivity analyzes, including the introduction of a risk-free asset as well as variations in the parametrization strategy. Two caveats apply. First, our four trends are treated as exogenous and uncorrelated (see subsection 6.2 for further discussion). Second, we compare the stationary equilibrium in 1935 with that in 2019, leaving aside questions related to transitional dynamics (see subsections 8.1 and 8.2 for further discussion). These features keep the numerical exercises tractable without compromising the analysis.

The remainder of the paper is organized as follows. Section 2 reviews the literature. Section 3 sets out the empirical facts justifying our modeling approach. Section 4 lays out the continuous-time setup used for quantitative analysis. Section 5 describes our estimation and calibration procedures. Section 6 presents the main quantitative results. Section 7 tests their robustness, while Section 8 discusses our model's limitations. Lastly, Section 9 concludes.

2. Literature review

Since the beginning of the twentieth century, human mortality has declined significantly (see Bergeron-Boucher et al., 2015; Pestieau and Ponthiere, 2016, for extensive discussions). Improvements in living standards, nutrition, and medical care have been the main drivers of this change (Oeppen and Vaupel, 2002; Riley, 2001). In the first half of the century, early reductions in infant, child, and young adult mortality played a crucial role in narrowing lifespan disparities and increasing life expectancy, as deaths became more concentrated at older ages (Kannisto, 2000; Edwards and Tuljapurkar, 2005; Cheung et al., 2009; Vaupel et al., 2011). This pattern led to survival curves taking on a more rectangular shape, a phenomenon referred to as *rectangularization* (Fries, 1980).

Starting in the 1950s and continuing into the 1960s, increases in life expectancy began to largely be driven by declines in mortality at old ages (Christensen et al., 2009). Meanwhile, the dispersion of age at death began to decrease at a slower rate compared to earlier periods (Fig. 1). Under this emerging pattern, improvements in life expectancy are marked by a shift in the distribution of deaths to older ages, with only slight changes in its shape (Cheung and Robine, 2007; Canudas-Romo, 2008).

These historical developments, well-known to demographers and health researchers, motivate our research question: how have breakthroughs in both preventive care and curative treatments contributed to improvements in human mortality? To our knowledge, previous works have largely overlooked the distinction between technological progress in preventive and curative care, especially in the theoretical literature. Importantly, our analysis focuses on advancements in health technologies, given their crucial role as determinants of health (Cutler et al., 2006; Fuchs, 2010). Extending our insights to frameworks that incorporate the earlier mentioned improvements in living standards and nutrition remains a topic for future research.

Murphy and Topel (2006) also examine breakthroughs in health technologies during the twentieth century, quantifying the large economic benefits of declining mortality in the US population. Their theoretical framework distinguishes between two types of health technologies: those that extend life and those that improve quality of life. Our modeling choices incorporate these two effects. Breakthroughs in preventive care improve health, extending expected lifespans and reducing the economic burden of poor health, thereby freeing resources for utility-enhancing activities. Similarly, breakthroughs in curative care extend lifespans by lowering the probability of death for a

given health condition. That said, our work extends Murphy and Topel (2006) since we also study how advancements in health technologies shape changes in the distribution of age at death over time. In addition, in our framework, individuals' health choices are endogenous and depend on the state of medical technology, unlike the exogenous approach in Murphy and Topel (2006).

This endogeneity of health choices connects our work to the seminal contributions of Grossman (1972) and Ehrlich and Becker (1972), who stressed that health is, in part, an investment good that can be actively shaped by allocating effort and resources.² In this vein, Strulik (2015) uses a life-cycle model to argue that income growth and medical advances in preventive care have reduced the biological age of middle-aged Americans by about 15 years and doubled their value of life. In subsequent work, Strulik (2022) shows that when the effectiveness of preventive care increases by 50%, the average American lives over 7 years longer, and experiences a 10% increase in the value of life at age 65. An important distinction with these papers is that, in our setup, we also introduce curative technology and consider the age-at-death dispersion. Moreover, the uncertainty surrounding the time of death plays a more significant role. While the mentioned papers assume that the survival function depends solely on the current health deficit, we assume that it is the hazard rate that depends on it. Hence, our survival function depends on the entire *history* of health deficits since birth, not just the current one. Therefore, our approach introduces the probability of remaining alive each instant as a state variable with an associated co-state, known as the "value of life-saving" (Schelling, 1968; Mishan, 1971). This is not innocuous for the optimality conditions governing the model's solution, and hence, for the role of uncertainty in determining the use of preventive care (see, e.g., Ehrlich, 2000, for an approach similar to ours, where the survival function depends on all past choices since birth). Finally, Garcia Sanchez et al. (2025) also propose a life-cycle model with endogenous prevention that affects health deficits. They examine how uncertainty regarding the monetary cost of health deficits influences prevention decisions. However, their model features a degenerate distribution of age at death, with the entire population dying at the same exogenous age, and therefore excludes curative technology.

Regarding the role of medical progress more broadly, Zhang et al. (2001) show that rising longevity has direct effects on fertility, human capital investment, and economic growth, as well as indirect effects through increasing unfunded social security contributions. However, Hwang and Kim (2023) reveal that these effects vary depending on whether the increase in longevity is expected or unexpected. Examining barriers to health care, Frankovic and Kuhn (2019) suggest that a large share of the increase in longevity inequality between high- and low-income groups results from wealthier individuals' better access to state-of-the-art medical care. Also worth noting is the extensive literature assessing the role of health insurance in rising medical expenditures, health, and life expectancy (see e.g., Pashchenko and Porapakkarm, 2017; Cole et al., 2019; Fonseca et al., 2020; Frankovic and Kuhn, 2023; Wang et al., 2023; Chen et al., 2025). All of these models differ significantly from ours, particularly in their formulation, methods, and research questions.

A key work on endogenous health investments is De Nardi et al. (2010), who explore why many elderly keep large amounts of assets until very late in life. Unlike ours, their model assumes that survival is exogenous to health investments and focuses on retired households. In addition, our work is close in spirit to Hall and Jones (2007), who study a simple setup where utility depends on quantity of life (life expectancy) and quality of life

²See Soares (2015) for a discussion of Gary Becker's contributions to health economics.

(consumption). They argue that health spending, which enables individuals to live longer and improve their quality of life, is a superior good with an income elasticity greater than one. As people become wealthier, they allocate a larger share of their resources to healthcare, since extending life does not experience the same diminishing returns as consumption. However, recent empirical research challenges this conclusion. Baltagi et al. (2017) estimate a global income elasticity using data from 167 countries between 1995 and 2012, finding that healthcare functions as a necessity with an income elasticity less than one, rather than as a superior good. Our framework reconciles these findings, suggesting that while the use of preventive care is indeed a superior good, a healthcare basket – encompassing not only preventive care, but also curative, palliative, and LTC – does feature an income elasticity less than one.

3. Motivating empirical facts

Our theoretical model will be built on four premises: first, as individuals age, they accumulate an increasing number of health problems, collectively termed the health deficit; second, as the health deficit rises, individuals become more vulnerable to death; third, although it depends on the health deficit, the time of death is uncertain; and fourth, individuals can affect the dynamic evolution of their health by consuming preventive care. This section discusses the empirical plausibility of these premises.

3.1. The health deficit increases with age

In their pioneering work, Mitnitski et al. (2002) showed that an individual's health status can be represented by an index variable, summarizing the accumulated number of health problems. The index, here termed the health deficit, is defined as the ratio of a person's accumulated health issues to the total number of conditions considered.

To construct the health deficit index used throughout our paper, we use data from the 2023 NHIS, which provides a sample of over 29,000 individuals after restricting ages to 18 to 85. We consider 17 health conditions, all requiring a medical diagnosis to prevent variations in pain thresholds from affecting the index: hypertension, high cholesterol, coronary heart disease, angina pectoris, myocardial infarction, stroke, asthma, cancer, diabetes, chronic obstructive pulmonary disease, arthritis, dementia, hepatitis, epilepsy, Crohn's disease, ulcerative colitis, and psoriasis. As in Mitnitski et al. (2002) and Hosseini et al. (2021), we weight all health conditions equally, so incurring one additional health condition increases one's deficit by $1/17$ or 6%.

After computing the health deficit for each individual in our sample, we group individuals by age and calculate the mean health deficit for the average person within each age range. To smooth out any abrupt spikes, we apply a four-year backward moving average. The left panel in Fig. 2 plots the resulting index, which is a convex function of age. In Section 5.2, we approximate this index with a parametric exponential function. We show that the fit is strong and implies an average growth rate of 3% per annum, as documented in studies involving Australia, Canada, Sweden, and the United States (Rockwood and Mitnitski, 2007).

3.2. Higher health deficits make individuals more vulnerable to death

It is a scientific fact that most health issues included in our health deficit index increase the probability of death, making this modeling assumption uncontroversial. The right

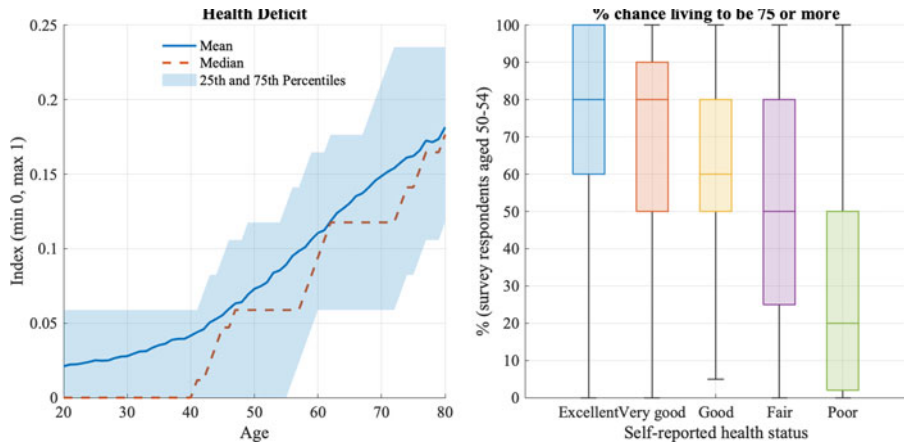


Figure 2. Selected health-related data.

Notes: The left panel displays various statistics based on a sample of over 29,000 individuals from the NHIS. An individual's health deficit is calculated as the ratio of accumulated health issues to the total number of conditions considered. The right panel plots the expected probability of living to be 75 or more as a function of self-reported health status on a sample of over 1400 individuals from the Health and Retirement Study. Each box chart displays the median, lower and upper quartiles, as well as the minimum and maximum values that are not outliers.

panel of Fig. 2 illustrates the link between health status and life expectancy using a sample of over 1400 individuals aged 50 to 54 from the 2019 Health and Retirement Study (HRS). It shows that as individuals report a deterioration in their health status – from excellent to poor – their expected probability of living to age 75 or beyond declines. Each box chart displays the median, lower and upper quartiles, as well as the minimum and maximum values that are not outliers.

3.3. *The time of death is uncertain*

While poor health increases the probability of death, uncertainty remains regarding the time of death. Knowing one's current health status does not allow for a precise forecast of when death will occur.³ The right panel of Fig. 2 illustrates this uncertainty, using data from the Health and Retirement Study. Consider the blue box chart representing individuals who self-report excellent health. The distance between the lower and upper quartiles can be interpreted as the uncertainty surrounding the median estimate of living to be 75 or more. Under this interpretation, the probability of a person aged 50 to 54 living to age 75 or more likely ranges from 60% to 100%, indicating significant uncertainty while remaining consistent with actuarial life tables. For example, the US 2019 actuarial life table in the Human Mortality Database indicates that the unconditional probability of someone aged 50 living to age 75 is roughly 70%.

³This might seem straightforward, but most macrohealth models overlook the uncertainty surrounding the time of death. In most cases, the time of death is either exogenous and known (see e.g., Cole et al., 2019; Garcia Sanchez et al., 2025) or endogenous and deterministic (see e.g., Dalgard and Strulik, 2014).

3.4. Preventive care use slows down health deficit growth

The notion that health and aging are malleable is widely acknowledged, with extensive research focusing on the social determinants of health and aging (Marmot and Allen, 2014; Hood et al., 2016). In fact, recent studies indicate that the aging process is far more adaptable than previously believed (Scott, 2023). Furthermore, healthy lifestyles, regular screening tests, and effective management of existing health conditions are known to prevent or delay serious illnesses such as strokes, heart disease, and cancer – conditions included in our health deficit index and leading causes of death.⁴ For instance, Cole et al. (2019) use survey data to show that Americans who exercise more than the cross-sectional average are more likely to maintain or improve their health statuses over the following six years.

All told, the empirical evidence supports the four premises underlying our theoretical analysis.

4. Baseline continuous-time life-cycle model

In this section, we introduce the model for evaluating the quantitative impact of breakthroughs in preventive and curative care on the distribution of age at death. Its key feature is that the time of death is a random variable, with its probability distribution determined by an individual's health deficit, which depends on their endogenous use of preventive care. Advances in preventive care technology enhance its effectiveness in slowing the accumulation of health deficits, while improvements in curative care technology reduce the probability of death for a given health deficit level.

4.1. Setup

Let $d(t)$ represent an individual's health deficit, which evolves according to the following law of motion

$$\dot{d}(t) = \gamma \left(d(t) - \frac{Ah(t)^\beta}{\beta} \right), \quad (1)$$

$$d(0) = d_0, \quad (2)$$

where d_0 is a strictly positive parameter. Equation (1) captures a simple dynamic: as an individual ages, her health deteriorates, reflected by the accumulation of the health deficit at a rate $\gamma > 0$. However, preventive care, $h(t)$, slows down this process. Preventive care includes activities aimed at preventing diseases and risk factors (e.g., vaccination) or early disease detection (e.g., screening), with further examples in Appendix A. The parameter $A > 0$ governs the effectiveness of preventive care, while $\beta \in (0, 1)$ introduces diminishing returns to scale, as in Dalggaard and Strulik (2014). This reflects the principle that as preventive care efforts increase, each additional intervention yields a smaller marginal improvement in health. For example, repeated screenings for the same cancer type within a short period provide little added value beyond the first.

⁴For more information on stroke and heart disease prevention, refer to the resources provided by the US Centers for Disease Control and Prevention at <https://www.cdc.gov/stroke/prevention.htm>, and for cancer prevention, visit <https://www.cdc.gov/cancer/dpcp/prevention/index.htm>.

The individual's lifespan extends from 0 to T , where T is a random variable. The probability law governing T is defined by the hazard rate $\lambda(d(t)) \geq 0$, which is a C^1 -function of $d(t)$ with $\lambda_d(d(t)) > 0$. It follows that the probability that the individual is alive at time t is $\Lambda(t) = e^{-\int_0^t \lambda(d(u))du}$. We prevent unjustifiably long lifespans by imposing an exogenous maximum health deficit $\bar{d} > d_0$; if the individual reaches \bar{d} , she passes away immediately. Consequently, the maximum attainable age is implicitly determined by $d(\bar{T}) = \bar{d}$.

The above specification implies that the likelihood of death increases with the individual's health deficit. Thus, changes in the function $\lambda(d(t))$ that lower its value for a given health deficit reflect improvements in curative treatments that reduce mortality risks (e.g., antibiotics, chemotherapy, and statins), with further examples in Appendix A.

At each instant $t \in [0, T]$, the individual receives a constant income $y > 0$. Her budget constraint is

$$c(t) + \theta(h(t) + Bd(t)) = y, \quad (3)$$

where $c(t)$ represents consumption and $\theta > 0$ is the relative price of medical-related activities. The term B is a positive parameter that captures the monetary cost associated with the health deficit, $\theta Bd(t)$. These costs represent pure expenditures with no direct effect on the health deficit and can be viewed as expenses related to LTC (e.g., nursing, home care), palliative care, as well as some types of curative expenditures. Equation (3) implicitly includes several simplifying assumptions: we consider only the active population, assume constant income over the life cycle, and exclude income heterogeneity. Moreover, no risk-free assets are available to smooth consumption over time. We examine and discuss these assumptions in detail in Sections 7 and 8.

The individual's expected lifetime utility is

$$\int_0^{\bar{T}} e^{-\rho t} \Lambda(t) \frac{c(t)^{1-\sigma} - 1}{1-\sigma} dt, \quad (4)$$

where $\rho \geq 0$ is the discount rate. We take a constant relative risk aversion (CRRA) utility function, where $1/\sigma > 0$ is the inter-temporal elasticity of substitution. For future reference, our numerical exercises ensure $c(t) > 1 \forall t$, so that utility flows remain strictly positive. Otherwise, the individual would prefer an earlier death, as continuing life would result in negative utility (see e.g., Dragone and Strulik, 2020, for a similar discussion).

The individual chooses sequences $\{c(t), h(t)\}_{t=0}^{\bar{T}}$ to maximize (4), subject to (1)–(3) and the endogenous hazard rate $\lambda(d(t))$.⁵

⁵We observe that agents optimize with respect to prevention and not with respect to the monetary cost of deficit, which is considered as unavoidable. We justify this approach in subsection 8.3 and propose a sensitivity analysis with alternative choices. Moreover, we do not consider here social security. A proportional subsidization of health expenditures financed through lump-sum taxes would leave the solution unchanged. Introducing non-proportional subsidies and/or distortive taxes, however, would create new dynamics, which are beyond the scope of this paper.

4.2. Solution

We solve our stochastic control problem by reformulating it as an equivalent deterministic control problem.⁶ As mentioned earlier, the probability that the individual is alive at time t is $\Lambda(t) = e^{-\int_0^t \lambda(d(u))du}$. Hence, we have $\dot{\Lambda}(t) = -\lambda(d(t))\Lambda(t)$ with $\Lambda(0) = 1$, allowing us to write the Hamiltonian function

$$H(t) = e^{-\rho t} \Lambda(t) \frac{c(t)^{1-\sigma} - 1}{1-\sigma} - \tilde{q}(t) \gamma \left(d(t) - \frac{Ah(t)^\beta}{\beta} \right) - \tilde{p}(t) \lambda(d(t)) \Lambda(t) + \tilde{\varepsilon}(t) [y - c(t) - \theta(h(t) + Bd(t))].$$

Here $-\tilde{q}(t)$ is the shadow price of health deficit and measures the value of an infinitesimal increase in $d(t)$. Similarly, $\tilde{p}(t)$ is the shadow price of the probability of survival, better known as the value-of-life-saving (Schelling, 1968; Mishan, 1971). It measures the remaining lifetime utility along the optimal path from t to \bar{T} . Lastly, $\tilde{\varepsilon}(t)$ measures the change in the optimal value of the utility function per unit of change in the budget constraint. Economic logic suggests that all three co-state variables $\{\tilde{q}(t), \tilde{p}(t), \tilde{\varepsilon}(t)\}$ should be positive.

Applying the maximum principle to $H(t)$ yields

$$\begin{cases} H_h = 0, & H_c = 0, & H_{\tilde{\varepsilon}} = 0, \\ H_d = \tilde{q}(t), & H_\Lambda = -\dot{\tilde{p}}(t), & H_{-\tilde{q}} = \dot{d}(t), & H_{\tilde{p}} = \dot{\Lambda}(t). \end{cases}$$

These necessary optimality conditions are standard in deterministic control theory. Furthermore, the concavity of the utility function ensures that these necessary conditions are also sufficient. Let $q(t) := e^{\rho t} \tilde{q}(t) / \Lambda(t)$, $\epsilon(t) := e^{\rho t} \tilde{\varepsilon}(t) / \Lambda(t)$ and $p(t) := e^{\rho t} \tilde{p}(t)$. The optimal control system must thus solve the following system of differential equations

$$\begin{cases} \dot{d}(t) = \gamma \left(d(t) - \frac{Ah(t)^\beta}{\beta} \right), & (5a) \end{cases}$$

$$\begin{cases} \dot{\Lambda}(t) = -\lambda(d(t))\Lambda(t), & (5b) \end{cases}$$

$$\begin{cases} \dot{q}(t) = (\rho + \lambda(d(t)) - \gamma)q(t) - \lambda_d(d(t))p(t) - \theta B(t), & (5c) \end{cases}$$

$$\begin{cases} \dot{p}(t) = (\rho + \lambda(d(t)))p(t) - \frac{\epsilon(t)^{\frac{1-\sigma}{1-\sigma}-1}}{1-\sigma}, & (5d) \end{cases}$$

as well as the three intratemporal conditions

$$\begin{cases} c(t) = \epsilon(t)^{-\frac{1}{\sigma}}, & (6a) \end{cases}$$

$$\begin{cases} h(t) = \left(\frac{\gamma A q(t)}{\theta(t)} \right)^{\frac{1}{1-\beta}}, & (6b) \end{cases}$$

$$\begin{cases} y = c(t) + \theta(h(t) + Bd(t)). & (6c) \end{cases}$$

Solving the above system of differential equations requires a set of boundary conditions, which in our setup are naturally given by

$$d(0) = d_0, \quad \Lambda(0) = 1, \quad d(\bar{T}) = \bar{d}, \quad p(\bar{T}) = 0.$$

⁶See Boukas et al. (1990) for mathematical proofs. Basically, we are able to transform a random setup into a deterministic problem because, apart from death, no shock arrives over time and the time-0 solution always remains unchanged. See for instance Garcia Sanchez et al. (2025) for a model in which a shock on income may happen at any moment, which in this case requires an adjustment to the solution path.

The first three conditions have been previously introduced. As for the last condition, it ensures that the remaining lifetime utility at the maximum attainable age \bar{T} is zero, which must be the case since the objective function does not include any bequest terms. Lastly, since \bar{T} is free, it must be endogenously determined by $H(\bar{T}) = 0$ (see e.g., Seierstad, 2009, for a formal derivation).

Unfortunately, there is no closed-form solution to this nonlinear boundary value problem. Therefore, after selecting all parameter values, the next section will solve it numerically using the collocation method proposed by Shampine et al. (2003).

4.3. Stationary distribution

So far, we have described the optimization problem for a single individual. However, our objective is to study the distribution of age at death across the entire population. To address this, we consider the stationary distribution implied by our model. Specifically, we assume that at each instant, a new cohort of size 1 is born, represented by $n(0) = 1$. Since all individuals are identical and face the same optimization problem, population dynamics are captured by $\dot{n}(t) = -\lambda(d(t))n(t) = \dot{\Lambda}(t)$. Thus, $\Lambda(t)$, which indicates the probability of a single individual being alive at age t , also represents the size of the population aged t in the stationary equilibrium.

Therefore, the first moment of the distribution of age at death, referred to as life expectancy, is

$$\mu_T = \int_0^{\bar{T}} t\lambda(d(t))\Lambda(t)dt + \bar{T}\left[1 - \int_0^{\bar{T}} \lambda(d(t))\Lambda(t) dt\right].$$

The first term integrates over ages up to the maximum admissible age, weighting each age by its density, $\lambda(d(t))\Lambda(t)$. The second term adjusts for the probability of surviving to the maximum age, \bar{T} . Similarly, the variance of age at death is

$$\sigma_T^2 = \int_0^{\bar{T}} (t - \mu_T)^2 \lambda(d(t))\Lambda(t)dt + (\bar{T} - \mu_T)^2 \left[1 - \int_0^{\bar{T}} \lambda(d(t))\Lambda(t) dt\right].$$

For future reference, the share of total resources spent on healthcare is

$$H_1 = \frac{\theta \int_0^{\bar{T}} (h(t) + Bd(t))\Lambda(t) dt}{y \int_0^{\bar{T}} \Lambda(t) dt},$$

while the share spent on preventive care is

$$H_2 = \frac{\theta \int_0^{\bar{T}} h(t)\Lambda(t) dt}{y \int_0^{\bar{T}} \Lambda(t) dt}.$$

5. Mapping the model to the data

As explained in the introduction, we conduct our quantitative analysis in two steps. In this Section 5, we select parameter values to align the model with key observations from the US in 2019, prior to the impact of COVID-19. Then, in the next Section 6, we keep the same calibration, except that we set the levels of both medical technologies, income and health care prices to their observed counterparts in 1935, and we examine how the distribution of age at death evolved between 1935 and 2019.

Table 1. Model parametrization

Parameter	Value	Description	Parameter	Value	Description
<i>A priori chosen parameters</i>					
β	0.5	Decreasing returns to h	θ	1	Relative price of care
y	3	Income flow	$1/\sigma$	1	IES
<i>Parameters estimated from the data</i>					
γ	0.7	Natural deficit growth	ρ	1	Discount factor
α_1	0.015	Hazard rate: scaling	α_2	24	Hazard rate: sensitivity
d_0	0.02	Initial deficit	\bar{d}	0.31	Maximum deficit
<i>Parameters calibrated within the model</i>					
A	0.013	Efficiency of prevention	B	3.6	Cost of deficit

Notes: $t = 0$ in the model corresponds to the age of 20 years, and $t = \bar{T}$ corresponds to the age of 105 years. Since our calibration implies that $\bar{T} = 4.3$, one unit of time in the model corresponds to $(105 - 20)/4.3 \approx 20$ years.

To reproduce 2019 U.S. data, we group the model parameters into three sets: a subset determined a priori, $\{\beta, \theta, y, \sigma\}$; a subset estimated from data, $\{\gamma, \lambda(d), \rho, d_0, \bar{d}\}$; and a subset calibrated within the model to minimize the distance between data targets and model outcomes, $\{A, B\}$. Table 1 summarizes the calibration.

5.1. A priori chosen parameters

To solve the model, we need to express $\epsilon(t)$ as a function of the state and co-state variables. We do this by substituting equations (6a) and (6b) into equation (6c), which yields an equation in $\epsilon(t)$. The solution to this equation gives $\epsilon(t)$, but it is guaranteed to be unique and positive only for certain values of (β, σ) . Appendix B1 shows that uniqueness holds under any of the following conditions: (i) $\sigma = 1 - \beta$; (ii) $\sigma = 0.5(1 - \beta)$; or (iii) $\sigma = 2(1 - \beta)$. Outside of these cases, uniqueness and positivity of $\epsilon(t)$ are not assured. In this paper, we choose $\sigma = 1$, which corresponds to log utility. Under this specification, only condition (iii) implies $\beta \in (0, 1)$ – that is, diminishing returns to preventive care – and leads to $\beta = 0.5$. In subsection 7.4, we conduct a sensitivity analysis on the parameter pair (β, σ) .

Next, we set the income stream y , which is a scaling parameter, to 3, ensuring that $c(t) > 1$ for all t . This choice keeps utility flows positive, ruling out any preference-for-death scenario. Lastly, we normalize the relative price of medical care, θ , to 1.

5.2. Parameters estimated from the data

We think of $t = 0$ as the age at which an individual reaches adulthood, setting its empirical counterpart to age 20. Based on the 2019 actuarial life table from the Human Mortality Database, only 0.5% of individuals live to age 105, so we set 105 as the empirical counterpart for the model's maximum admissible age, \bar{T} .⁷ Assuming an annual time discount factor of 0.95, we could normalize $\bar{T} = 1$ and choose the time discount rate ρ such that

⁷For context, 4% of individuals reach the age of 100.

$$e^{-\rho\bar{T}} \approx 0.95^{(105-20)}.$$

However, since \bar{T} is determined as part of the model's solution rather than set as a fixed parameter, we instead normalize $\rho = 1$, and retain this equation as a condition to match (i.e., matching $\bar{T} = 4.3$) when setting the third subset of parameters.

To model the health deficit process, we first set the initial value $d(0)$ and the exogenous upper bound \bar{d} using our empirical health deficit index, shown by the solid blue line in Fig. 2 (left panel). We set $d(0) = 0.02$, matching the mean health deficit for individuals aged 20 in the data. Calibrating \bar{d} , representing the mean health deficit at age 105, is more challenging, as our data only includes individuals up to age 80. To overcome this problem, we follow Mitnitski et al. (2002) and fit the exponential regression

$$d(t) = \theta_0 + \theta_1 e^{\theta_2 t}, \quad t \in [18, 80].$$

The fit is strong, with a root mean square error below 0.01, so we set $\bar{d} = \theta_0^* + \theta_1^* e^{\theta_2^* 105} \approx 0.31$.

Next, parameter γ represents the natural growth rate of the health deficit – that is, its rate of increase in the absence of preventive care. We estimate γ as follows: as discussed in the introduction, in 1935, when modern medical care was in its infancy, the effectiveness of preventive care was negligible, which translates to $A = 0$ in the model (see also subsection 6.1 for a discussion). Therefore, the evolution of the health deficit in 1935 would depend solely on biological parameters

$$d(t) = d(0)e^{\gamma t},$$

with $d(\bar{T}) = \bar{d}$ or equivalently $d(\bar{T}) = d(0)e^{\gamma\bar{T}}$. More precisely, it depends on the natural growth rate γ , the initial health deficit $d(0)$, and the maximum level of health deficits an individual can sustain, \bar{d} . The actuarial life table from 1935 suggests that taking a maximum admissible age of 98 years is reasonable, given that only 0.5% of individuals lived to that age. In model's unit, the maximum age of 98 years in 1935 corresponds to $t = \bar{T}^{1935} \approx 4$ (since $t = 0$ corresponds to the age of 20 years, $t = 4.3$ corresponds to the age of 105 years). Thus, we determine γ as

$$\gamma = \frac{1}{\bar{T}^{1935}} \ln \frac{\bar{d}}{d(0)} \approx 0.7.$$

Lastly, we address the hazard rate function $\lambda(d(t)) \in C^1$ that governs the probability law of the death process. Section 3 argued that higher health deficits increase individuals' vulnerability to death, leading us to assume $\lambda_d(d(t)) > 0$. A simple functional form in line with these requirements is

$$\lambda(d(t)) = \alpha_1 e^{\alpha_2 d(t)}.$$

Here, α_2 measures the percentage change in the hazard rate for a one-unit change in the health deficit, while α_1 scales the entire exponential distribution. In other words, α_2 governs the sensitivity of $\lambda(d)$ to changes in d , and α_1 affects the magnitude of that sensitivity at each point along the curve. Before estimating the α coefficients, we must construct a time series for the exogenous hazard rate variable $\lambda(t)$. The Human Mortality Database provides the probability $z(t)$ of death for individuals aged 0 to 110 in the coming year. In model's notation, we have

$$z(t) = \frac{\Lambda(t) - \Lambda(t + dt)}{\Lambda(t)},$$

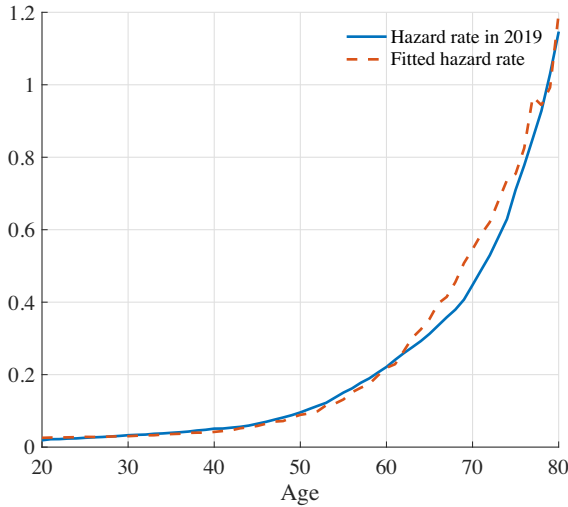


Figure 3. Hazard rate governing the probability law of the death process.

Notes: The solid blue line represents the US hazard rate observed in 2019, derived from the Human Mortality Database. The dashed red line depicts the fitted hazard rate based on the exponential regression $\lambda(d(t)) = \alpha_1 e^{\alpha_2 d(t)}$.

with $dt = 4.3/85 = 0.05$ representing one year in the data (since $t = 4.3$ in the model represents $(105 - 20) = 85$ years). Then, the hazard rate is

$$\lambda(t) = \frac{-\dot{\Lambda}(t)}{\Lambda(t)} \approx \frac{-(\Lambda(t + dt) - \Lambda(t))}{dt \Lambda(t)} = \frac{z(t)}{dt}.$$

Using the two time series for the hazard rate and health deficit index across an individual's lifespan, we estimate the α coefficients via ordinary least squares. More precisely, we estimate

$$\ln \lambda(t) = \ln \alpha_1 + \alpha_2 d(t),$$

which results in $\alpha_1 \approx 0.015$, $\alpha_2 \approx 24$, and an R-squared value of almost 0.99. Figure 3 compares the empirical hazard rate with the hazard rate obtained with our estimated functional form.

5.3. Parameters calibrated within the model

We calibrate the remaining parameters using a simple grid search method, where we evaluate parameter values over a predefined range to ensure selected statistics align with their empirical counterparts. These parameters govern the efficiency of preventive care in slowing the growth of health deficits (A) and the monetary costs associated with those deficits (B). The empirical targets used for estimation include a maximal admissible age of 105 and a ratio of total expenditures on preventive care to GDP of 0.5%, as observed in the US in 2019.⁸ Using our model notation and units, these targets read as: $\bar{T} = 4.3$ and $H_2 = 0.54\%$ (see subsection 4.3).

⁸Source: Organisation for Economic Co-operation and Development Database.

Table 2. Targeted moments: data and model

	Maximum lifespan	Preventive care to GDP
	\bar{T}	H_2
Data	105	0.54%
Model	103	0.57%

Notes: Maximum lifespan and preventive health expenditures divided by GDP, in both the 2019 US data and the model. In the model, we compute the expenditures using a 10,000-agent Monte Carlo simulation. \bar{T} is expressed in years.

Table 3. Non-targeted moments: data and model

	Distribution of age at death				Total care to GDP
	μ_T	σ_T	Skewness	$\Lambda(\bar{T})$	H_1
Data	79.5	16.7	−1.6	0.5%	17.5%
Model	83.9	15.0	−1.9	0.3%	10%

Notes: Expected value, standard deviation, skewness, share of individuals living to age 105, and total health expenditures divided by GDP, in both the 2019 US data and the model. In the model, we compute the moments using a 10,000-agent Monte Carlo simulation. μ_T and σ_T are expressed in years.

The resulting parameter values are $A = 0.013$ and $B = 3.6$. The targeted moments are well matched; the model yields $\bar{T} = 4.2$ and $H_2 = 0.57\%$. We do not achieve a perfect match for the two targets, despite having two instruments, because utility flows must remain positive to avoid the “preference-for-death” scenario discussed earlier. As a result, our grid search algorithm functions as a constrained minimization problem. Table 2 summarizes the fit.

5.4. Assessment

To assess the model’s performance regarding non-targeted moments, we first examine the first three moments of the age-at-death distribution and the share of individuals reaching the maximal admissible age. For simplicity, we compute the moments of the age-at-death distribution using a 10,000-agent Monte Carlo simulation. Table 3 shows that the model performs reasonably well in capturing the patterns observed in the data. While it slightly overestimates (underestimates) the first (second) moment of the age-at-death distribution, these estimates remain within a plausible range. Additionally, the model successfully replicates the well-known long left tail of the age-at-death distribution and the share of individuals living to age 105. Figure 4 summarizes the fit of the model regarding the age-at-death distribution.

Finally, our calibration targets the share of expenditures on prevention (H_2) but not the share of total health expenditures (H_1). The last column in Table 3 shows that our model underestimates total expenditures (17.5% in the data vs. 10% in the model). However, our model is able to reproduce the convexity of total expenditures relative to income over the life cycle, particularly the strong increase for the elderly, as described in De Nardi et al. (2016). Specifically, in our model, total expenditures as a share of income rise from 3.2% at age 20 to 14% at age 80 and 38% at age 100.




Figure 4. Age-at-death distribution: data and model.

Notes: The solid blue line and the dotted red line represent the share of deaths in 5-year intervals based on the 2019 U.S. data and the model, respectively. The model's distribution is calculated using a Monte Carlo simulation with 10,000 agents.

Next, we look at the effect of income on healthcare expenditures and life expectancy. We compare the results obtained with our model to the results from related papers (since we cannot compare to direct evidence from the data). First, we compute the income elasticity of preventive care expenditures as

$$\eta = \frac{\partial \int_0^{\bar{T}} \theta h(t) \Lambda(t) dt}{\partial \int_0^{\bar{T}} y \Lambda(t) dt} \frac{\int_0^{\bar{T}} y \Lambda(t) dt}{\int_0^{\bar{T}} \theta h(t) \Lambda(t) dt}.$$

The income elasticity of total medical expenditures ($\theta(h + Bd)$) is defined in a similar manner. The model suggests an income elasticity of preventive care use of 2.9 and an income elasticity of total medical expenditures of 0.3. These figures align with previous studies. As in Hall and Jones (2007), preventive care use in our setup enables individuals to live longer (on average) and enhances their quality of life by reducing the monetary costs associated with health deficits. As a result, preventive care use is a superior good. In contrast, the model suggests that total medical expenditures, which also encompass the monetary costs of health deficits related to curative, palliative and LTC, have an income elasticity of less than one, consistent with recent empirical evidence in Baltagi et al. (2017).

Second, we consider how much life expectancy increases with a 10% rise in income. Chetty et al. (2016) estimate that in 2014, the expected age at death for 40-year-old American men in the second quartile of the income distribution (mean \$47k) was approximately 81 years, while for those in the third quartile (mean \$83k), it was around 85 years. Assuming a linear relationship between expected age at death and income in that segment of the income distribution – an assumption supported by further results in the paper – an increase in income of 10% correlates with an increase in life expectancy of about six months. When individuals' income (y) increases by 10% in our model, the expected life expectancy at age 40 rises by slightly less than year. In our model, the income increase is permanent and begins at birth, whereas in Chetty et al. (2016), a

person at a given income quartile at age 40 could have experienced varying income levels earlier in life and might shift quartiles later. This difference, along with the model's omission of potential effects of health on income, might partly account for our model's higher estimate of income's effect on life expectancy, though it remains broadly consistent with Chetty et al. (2016).

In sum, the model replicates well empirical patterns of mortality and medical expenditures, which are key components of our analysis. This gives us confidence that the model provides a plausible laboratory for our quantitative decomposition of the roles of the trends (a)–(b) in shaping health outcomes (i)–(ii), between 1935 and 2019.

6. Quantitative exercises

We begin by constructing time series for the medical technologies that capture trends (a) and (b) – preventive and curative advances, respectively. Next, we construct two additional series reflecting major developments over the past century: (c) the sharp rise in the relative price of medical care, and (d) the increase in real median income per male worker from 1960 until the mid-1970s, followed by a period of stagnation. These four series are then incorporated into the model. Hosseini et al. (2021), Strulik (2022), De Nardi et al. (2024), and Chen et al. (2025), we do not simulate the transitional dynamics from 1935 to 2019. Instead, we compare stationary equilibria. Specifically, we solve the model using the values of trends (a) through (d) for the years 1935, 1960, 1990, and 2019. For each year, we compute the stationary distribution and derive key statistics. Lastly, we evaluate the relative contributions of each trend through counterfactual exercises.

Needless to say, what follows is conditional on the assumptions that we consider long-run equilibria and that individuals rationally choose their full life-cycle consumption and prevention paths. We discuss these modeling features in more detail in Section 8.

6.1. Advances in medical technology

Quantifying medical trends is challenging due to the difficulty in measuring medical progress. We proceed as follows: our calibration provides the 2019 values for parameters related to medical technology: A , governing preventive care efficiency, and α_1 and α_2 , governing the relationship between health deficits and mortality risk, that is, governing the curative care technology. To set the value of A in 1935, we follow Dalgaard and Strulik (2014) and assume that preventive care efficiency was virtually zero at that time, that is the role of technology in repairing health deficits was nil.⁹ Then we impose that A increases linearly to its 2019 value of 0.013 (see Fig. 5, left panel). This assumption also determines the evolution of health deficits in 1935, based solely on biological parameters

⁹This assumption is reasonable, as Appendix A demonstrates that most preventive care activities were introduced starting in the 1940s. More broadly, medicine during this era – whether preventive or curative – was often more harmful than beneficial. Medical historian Dr. Morris Fishbein, who served as editor of the *Journal of the American Medical Association* from 1924 to 1950, famously criticized the medical quackery of the time, exposing dangerous practices and fraudulent treatments. For example, many symptoms predominantly affecting women were hastily attributed to 'hysteria'. To treat these so-called 'mental illnesses', doctors frequently resorted to invasive procedures such as hysterectomies. Bucke (1898), for instance, is notorious for his theory and findings advocating surgery as a treatment for female patients. Furthermore, the lack of regulation prior to the 1906 Pure Food and Drug Act allowed harmful substances to be sold freely, highlighting the significant risks patients faced during this period.

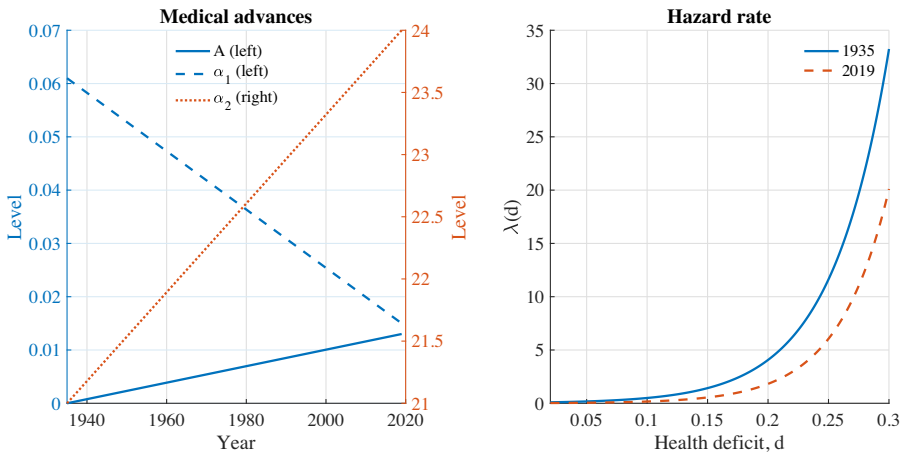


Figure 5. Effectiveness of preventive and curative care: Trends.

Notes: In the left panel, the increase in A corresponds to advances in the prevention technology and the joint decreases in α_1 and α_2 correspond to advances in the curative technology. The right panel shows that lower α 's indeed reduce the hazard rate for all levels of deficit.

(see subsection 5.2). The Human Mortality Database provides one-year-ahead death probabilities for ages 0 to 110, allowing us to derive the 1935 hazard rate as before (see again subsection 5.2). Using these two 1935 series for hazard rate and health deficit, we estimate the α parameters as we did for 2019, resulting in $\alpha_1 \approx 0.06$ and $\alpha_2 \approx 21$, with an R-squared of 0.97. Finally, we assume a gradual transition between 1935 and 2019, with each parameter adjusting by an equal amount each year (Fig. 5, left panel). The right panel of Fig. 5 shows the hazard rates in 1935 and 2019 as functions of the health deficit, based on the corresponding α coefficients. As expected, the 2019 hazard rate lies below the 1935 rate at any given health deficit level, reflecting advancements in curative treatments over this period.

There are at least two reasons to think medical technology likely evolved non-linearly. First, changes in the age-at-death distribution have shown significant non-linearity; for example, the expected age at death increased by about three years between 1935 and 1945, which matches the increase observed from 1965 to 2000. Indeed, most longevity gains in the US occurred during the early 20th century, largely due to reductions in mortality from infectious diseases, then the leading cause of death (Jayachandran et al., 2010). Second, there have been notable breakthroughs, such as the “golden age of antibiotics” in the 1950s and 1960s (Hutchings et al., 2019). However, in light of limited empirical guidance, our emphasis on simplicity, and our focus on qualitative insights over precise quantitative estimates, we opt for a gradual, linear transition. Our key insights are independent of this assumption, as our counterfactual scenarios focus solely on the initial and final years (1935 and 2019), with no attention given to the transition period between them.

6.2. Median income and healthcare prices

Together with the medical advances just described, the U.S. economy experienced two major developments: (c) a sharp rise in the relative price of medical care, and (d) an increase in real median income per male worker from 1960 until the mid-1970s,

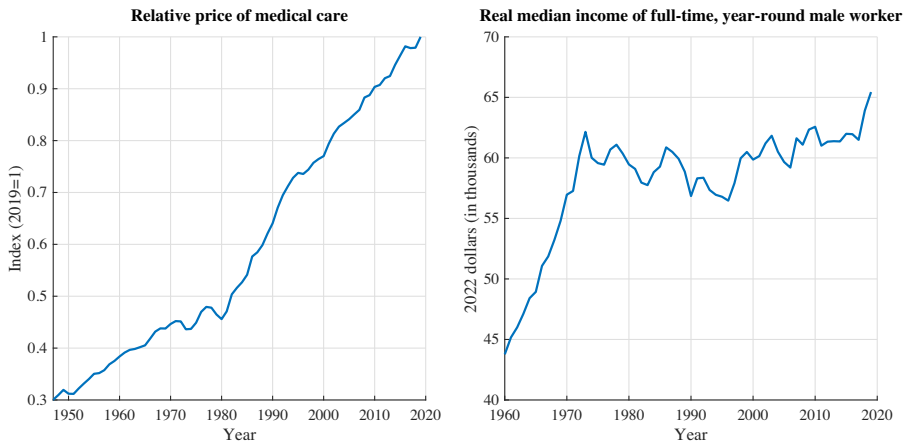


Figure 6. Economy-related trends in the US.

Notes: We define the relative price of healthcare as the ratio of the consumer price index for medical care to the overall consumer price index, both sourced from the US Bureau of Labor Statistics. Normalizing the 2019 value to 1 yields the left panel. Data on income is sourced from the Current Population Survey, conducted by the US Census Bureau and the US Bureau of Labor Statistics.

followed by stagnation (see Fig. 6). Considering medical advances (a)–(b) while ignoring these two developments may lead to an incomplete, or even misleading, understanding of the forces shaping longevity over time.

Needless to say, medical advances may have contributed to the rise in the relative price of medical care (Newhouse, 1992), suggesting that trends (a) and (b) might not be independent of trend (c). Likewise, increased life expectancy could raise income levels (Acemoglu and Johnson, 2007), indicating a potential causal link from health outcome (i) to trend (d). Despite these caveats, we treat trends (c) and (d) as exogenous – both for simplicity and because there is no compelling evidence that health outcomes (i)–(ii) and medical advances (a) and (b) are the *main* drivers of economic trends (c) and (d). More plausible explanations include structural features of the healthcare system, such as fragmentation across public and private payers (Anderson et al., 2019), as well as broader forces like foreign competition (Autor et al., 2013); shifts in technology and the relative price of capital (Acemoglu and Restrepo, 2018); the decline of unions (Stansbury and Summers, 2020); and employer concentration in local labor markets (Benmelech et al., 2022).

In terms of trend (c), we define the relative price of healthcare as the ratio of the consumer price index for medical care to the overall consumer price index, both sourced from the US Bureau of Labor Statistics. Normalizing the 2019 value to 1 for consistency with our calibration yields the left panel of Fig. 6. The relative price of medical care has tripled over the period, rising from approximately 0.3 in 1947 to 1 in 2019. Since data are available only from 1947 onward, we estimate the 1935 value by fitting a linear time trend to the available observations.

Turning to trend (d), we define income as the real median income of a full-time, year-round male worker, as shown in the right panel of Fig. 6. This data is sourced from the Current Population Survey, conducted by the US Census Bureau and the US Bureau of Labor Statistics. To align with our calibration, we normalize the 2019 value to 3. Since data is available only from 1960 onward, we estimate the 1935 value by fitting a linear time trend to the first 30 years of the sample, placing less emphasis on the later stagnation periods.

Table 4. Comparison of data and model outputs over selected years

Year	Source	Distribution of age at death		Health expenditures to GDP	
		Expected value (years)	Stand. Dev. (years)	Total, H_1 (%)	Preventive care, H_2 (%)
	Data	61.0	24.7	–	–
	Model	67.0	17.7	4.00	0.00
	Data	70.3	19.6	4.30	–
	Model	69.7	17.4	4.40	0.05
	Data	75.9	17.6	10.3	0.37
	Model	75.8	16.8	6.10	0.36
	Data	79.5	16.7	14.7	0.54
	Model	83.9	15.0	9.60	0.57

Notes: In the model, we compute the moments associated with the age-at-death distribution using a 10,000-agent Monte Carlo simulation. Data for H_1 refer to *Personal Health Care* expenditures from the National Health Expenditure Accounts (Centers for Medicare and Medicaid Services). Data for H_2 follow the System of Health Accounts 2011, and are sourced from the OECD. Expenditures are classified under “HC.6 Preventive care” if they fall within the health care boundary and aim to prevent disease, reduce risk factors, or enable early detection.

6.3. Benchmark scenario

Table 4 compares observed health outcomes with those implied by the model when we insert medical trends (a) through (d). Except for the preventive care expenditures to GDP ratio in 2019, none of the statistics in the table was targeted when selecting parameter values.

Let us begin with the model’s strengths. First, it captures the direction of changes in the distribution of age at death: life expectancy rises, while the standard deviation declines. It also replicates the increase in health care expenditure – both total and preventive – as a share of output. Finally, the model matches all statistics within a plausible range of their empirical counterparts. Nonetheless, the model overestimates the mean and underestimates the variance of the age-at-death distribution in 2019, with similar discrepancies in 1935. It also falls short in reproducing the full magnitude of the decline in the standard deviation over time.

All told, given its simplicity and limited number of parameters, the model does a good job of replicating key patterns in mortality and health expenditure. But how much does each trend (a)–(d) contribute to this performance? We turn to this next.

6.4. Counterfactual scenarios

To understand the contribution of each trend, we introduce them into the model cumulatively: first trend (a) alone, then trends (a) and (b) together, and finally the full set (a) through (d).¹⁰ Table 5 reports the resulting changes in life expectancy, dispersion in age at death, and health care spending as a share of income over the past several decades. The time span for each statistic varies based on data availability: from 1935 for the age-

¹⁰We introduce trends (c) and (d) jointly because it is not meaningful to simulate the model with 2019 income and 1935 medical prices (which leads to implausibly high life expectancy as preventive care surges), or with 1935 income and 2019 prices (which results in negative utility due to extremely low consumption).

at-death distribution, from 1960 for total health spending, and from 1990 for preventive care. As a benchmark, the first row shows the corresponding changes observed in the data.

We begin by introducing trend (a), which captures medical advances in preventive care. Specifically, we hold curative care technology, income, and the relative price of healthcare fixed at their 1935 levels. This means the α coefficients, γ , and θ remain constant at their 1935 values, while the effectiveness of preventive care, A evolves as observed. The second row of Table 5 reports the results. Crucially, the dispersion of age at death *increases* from 1935 to 2019, whereas it *decreases* in the data.

The logic is as follows. In our model, advancements in preventive care lower the average health deficit at each age by both increasing the effectiveness of current resources and boosting overall investment in preventive care (with an elasticity of total preventive care use with respect to A hovering around 2). A lower health deficit at each age then reduces the hazard rate, increasing life expectancy and extending the maximum possible lifespan, as deficits accumulate more slowly and \bar{d} is reached later. These dynamics have competing effects on the dispersion of age at death: lower mortality reduces it, while higher life expectancy and a longer maximum lifespan increase it. Under our calibration, the latter effect dominates, leading to a 1.4-year increase in the dispersion of age at death.

Offsetting this upward trend in dispersion requires advances in curative care, as the third row of the table shows. Improved curative care reduces the hazard rate of death at a given health deficit level (right panel of Fig. 5), weakening the impact of an individual's health deficit. *Ceteris paribus*, better curative care lowers the shadow price of health deficits, reducing the demand for preventive care. As a result, the average health deficit at each age rises, and the maximum lifespan shortens, with deficits accumulating more quickly and \bar{d} reached sooner. These dynamics compress the distribution of age at death: both a lower death probability and a reduced maximum lifespan contribute to this effect.¹¹

Does this prediction – that advances in curative care reduce the dispersion of age at death – align with empirical findings? Empirical evidence on this topic is limited, likely for two reasons. First, studies typically examine the impact of medical technology on life expectancy rather than on age-at-death dispersion. Second, disaggregated data on the spread of medical technologies is often unavailable, making it difficult to assess their specific effects on mortality patterns. One exception is Alsan et al. (2021), who find that penicillin reduced both the mean and the standard deviation of infectious disease mortality in Italy following World War II.¹²

Back to Table 5, the third row confirms that curative advances play a central role in raising life expectancy. By lowering the hazard rate of death at a given health deficit level,

¹¹In the limit, if all individuals reached \bar{d} alive, or if $\bar{d} \rightarrow d_0$, the dispersion in age at death would be zero.

¹²Also worth noting is the work by Jayachandran et al. (2010) on the introduction of sulfa drugs in the mid-1930s in the US. These drugs were the first effective treatment for a range of fatal bacterial infections, including pneumonia, meningitis, and scarlet fever. Since the next significant medical advancement occurred during the golden age of antibiotics in the 1950s and 1960s, the 1930s and 1940s offer a unique opportunity to examine the effects of sulfa drugs on mortality. The authors argue that sulfa drugs increased life expectancy by 0.4 to 0.7 years between 1937 and 1943. While a rigorous empirical analysis on the impact of sulfa drugs on the dispersion of age at death is beyond the scope of this paper, let us note the sharp decrease in this dispersion shown in the left panel of Figure 1 from 1935 to 1945, coinciding with the widespread distribution of these drugs in the US. This decline seems unlikely to be unrelated to the effects of sulfa drugs.

Table 5. Shock decomposition: changes in key variables over time

	Distribution of age at death		Health expenditures to GDP	
	Expected value (years)	Standard deviation (years)	Total, H_1 (%)	Preventive care, H_2 (%)
	1935–2019	1935–2019	1960–2019	1990–2019
Data	18.5	–8.0	10.4	0.2
Model				
A	3.3	1.4	0.1	0.2
A, λ	16.2	–2.9	0.7	0.1
A, λ, θ, y	17.0	–2.8	5.2	0.2

Notes: Total changes in life expectancy, age-at-death dispersion, and preventive care spending as a share of total income. The second row includes trends for both A and λ , while the third (and fourth) rows include only the trend for A (and λ , respectively). Since the model is nonlinear, a general Superposition Principle does not apply; the combined effects of trend A and trend λ need not equal the sum of their individual effects.

they allow individuals to live longer. This mechanism is consistent with empirical evidence. For example, Lichtenberg (2018) shows that cancer mortality declines significantly following the introduction of new cancer drugs, with effects appearing about five years after launch, across a sample of 36 countries.

Lastly, economic trends (c) and (d) have a limited effect on the distribution of age at death (fourth row of the table). Higher income raises preventive care, which leads to a modest increase in life expectancy relative to the previous scenario, with little change in the standard deviation. However, these trends play a central role in helping the model replicate the sharp increase in healthcare's share of total resources (H_1 , subsection 4.3), which rose by 10 percentage points from 1960 to 2019. Medical advances alone (trends a–b) account for only a 1-percentage-point increase, but adding economic trends (c) and (d) raises this to about 5 percentage points. This aligns with empirical evidence pointing to rising healthcare prices as a key driver of the growing share of output devoted to healthcare spending (Anderson et al., 2003, 2019).

To conclude, while both preventive and curative technologies have increased life expectancy in the US since 1935, it is curative innovations that have lowered the dispersion of age at death. In turn, economic factors are the primary drivers of the sharp growth in healthcare spending.

6.5. Welfare implications of preventive and curative advances

To better understand the historical role of medical breakthroughs, we now study their effects on living standards between 1935 and 2019, using a Lucas (1987)-style metric. We define

$$\mathcal{U}_j(\omega) = \int_0^T e^{-\rho t} \Lambda(t) \ln(\omega c(t)) dt,$$

as the expected discounted sum of future utility flows for an individual born at time j if her consumption at each instant was multiplied by a factor ω . To ease the notation, we immediately write our utility function as a log (we use a CRAA utility with $\sigma = 1$, see

Table 6. Annual growth rate of welfare (in %)

	g	Decomposition			
		g_{λ}	g_A	g_y	g_{θ}
–2019	0.87	0.10	0.01	–	–
–2006	0.19	0.11	0.02	0.16	–0.09

Notes: The table reports the annual average consumption-equivalent welfare growth rate g , as well as its decomposition into the four contributing trends: (a) preventive advances g_A , (b) curative advances g_{λ} , (c) relative price of medical care g_{θ} , and (d) median income g_y .

Table 1). We then compute the value of ω that would make an individual born in 1935 indifferent to living in 2019. An increase in living standards implies $\omega > 1$. Formally, we compute ω such that

$$\mathcal{U}_{1935}(\omega) = \mathcal{U}_{2019}(1),$$

which yields

$$\ln \omega = \frac{\mathcal{U}_{2019}(1) - \mathcal{U}_{1935}(1)}{\int_0^{\bar{T}_{1935}} e^{-\rho t} \Lambda_{1935}(t) dt}.$$

Hence, the average annual growth rate of living standards, in consumption-equivalent welfare terms, between 1935 and 2019 is

$$g = \frac{\ln \omega}{84}.$$

The first row of Table 6 shows that the four trends between 1935 and 2019 – curative advances λ , preventive advances A , median income y , and the relative price of medical care θ – together account for a welfare growth rate of 0.87% per year. Breaking this down, medical advances contribute 0.11% per year, with curative improvements accounting for most of this (0.10%) and preventive measures contributing the remainder (0.01%).

The reason why advances in curative care play a more prominent role in raising living standards than those in preventive care is twofold. First, as noted in the previous subsection, curative improvements were key to increasing life expectancy and reducing the standard deviation of age at death, both of which enhance welfare by lowering the risk of dying young. Second, better curative care reduces the shadow price of health deficits. *Ceteris paribus*, this allows individuals to shift resources away from prevention and toward consumption, which further boosts welfare. Lastly, the remaining growth in welfare (0.76%) stems from the joint evolution of income and prices.¹³

How do our model's estimates compare with empirical evidence? Using microdata, Jones and Klenow (2016) compute a welfare growth rate of 3.1% per year in the US between 1984 and 2006, primarily driven by income growth (2.1%) and increased life expectancy (0.9%). To facilitate comparison, we replicate our analysis for the sub-period 1984–2006. This shorter time horizon allows us to isolate the effects of income (y) from

¹³We assess the welfare effect of each trend by introducing them separately. As already mentioned in Footnote 6.4, the evolutions of income and prices during this period are significant and must be considered jointly.

those of healthcare prices (θ), since these shocks are less pronounced over 1984–2006 than over the full period from 1935 to 2019. The second row of Table 6 reports the results.

Over this limited time horizon, our model implies an annual welfare growth rate of 0.19%, far below the 3.1% estimated by Jones and Klenow (2016). Most of the gap stems from income dynamics. We focus on the real median income of full-time male workers, which grew by just 0.1% per year between 1984 and 2006. In contrast, Jones and Klenow (2016) consider a 2.1% annual income growth rate, over twenty times higher. The remainder of the gap likely reflects differences in methodology. While we rely on a simple theoretical setup with a limited set of trends, Jones and Klenow (2016) use an empirical approach that captures a broader range of factors absent from our model, including socioeconomic inequalities.

That said, our simple model complements the empirical estimates in Jones and Klenow (2016) by highlighting the negative welfare effects of rising healthcare prices, which nearly offset the gains from increasing real incomes. This result aligns with both the data and prevailing public concerns on the costs of surging healthcare prices. For example, in 1980, total healthcare premiums for the median worker accounted for 8% of take-home pay; by 2000, this share had more than doubled to 17%, and by 2020, it had reached 25% (Miller and Nyce, 2023). Rising healthcare prices lead many Americans to forgo necessary medical care, a pattern well documented in national surveys such as the NHIS. They also contribute to higher out-of-pocket spending relative to other high-income countries and are a frequent cause of financial distress, including personal bankruptcy.

7. Sensitivity analysis

We now explore the robustness of our results to alternative strategies of mapping the model to the data and to model assumptions. In subsection 7.1, we extend the model by introducing a risk-free asset, allowing agents to transfer resources safely over time. In subsections 7.2 and 7.3, we introduce age-varying income as well as income heterogeneity. Finally, in subsection 7.4, we assess how a different calibrations for (β, σ) would impact our findings.

7.1. Complete markets

For mathematical tractability, the baseline model assumes incomplete markets. We now relax this assumption, introducing a risk-free asset to allow agents transfer resources safely over time. Specifically, let an agent's asset holdings (or equivalently stock of savings) at age t be denoted as $s(t)$. Assuming no assets at birth, $s(t)$ follows

$$\dot{s}(t) = i(t), \quad (7)$$

$$s(0) = 0. \quad (8)$$

where $i(t)$ denotes the flow of savings, which is either positive or negative. Because of savings, the budget constraint (3) becomes

$$c(t) + i(t) + \theta(h(t) + Bd(t)) = y + rs(t) + z(t),$$

where $r > 0$ is the risk-free rate. Due to the uncertainty surrounding the agent's age at death, accidental bequests are certain. This means that when the agent passes away, she does so with either positive asset holdings ($s(T) > 0$) or debts ($s(T) < 0$). To tackle this issue, we

introduce an annuity market as in Yaari (1965) or Blanchard (1985). More precisely, insurance firms pay $z(s(t))$ to the individual at each instant. In return, the insurance firm collects all assets when the individual dies. The running profit of the insurance firm is

$$\pi(t) = \lambda(d(t))s(t) - z(t).$$

Free entry ensures zero profits, and hence, $z(t) = \lambda(d(t))s(t)$. The budget constraint can therefore be written

$$c(t) + i(t) + \theta(h(t) + Bd(t)) = y + (r + \lambda(d(t)))s(t). \quad (9)$$

Lastly, should the agent reach the maximum admissible lifespan, her assets must be zero, yielding $s(\bar{T}) = 0$. Appendix B2 derives the optimality conditions for the extended model, highlighting the numerical challenges brought by complete markets.

To quantify the effects of complete markets, we retain the calibration of the baseline model, and, in keeping with much of the literature, set the risk-free rate equal to the discount rate, $r = \rho$. Figure 7 contrasts the paths of consumption, health deficits, savings, and survival probability over the life cycle of an individual who reaches the maximum admissible lifespan in both the baseline and extended models. The comparison reveals that savings have a negligible effect on the accumulation of health deficits and, consequently, on survival probability. Since the latter equals the size of the population aged t in the stationary equilibrium (see subsection 4.3), savings have basically no impact on the distribution of age at death either. In other words, the incomplete market assumption is not behind our previous insights.

What is then the role of savings? Consumption smoothing. In the baseline model, consumption declines over the life cycle due to the increasing costs of accumulating health deficits and time discounting. In contrast, in a complete market economy, the standard assumption $r = \rho$ leads to a constant consumption profile. Consequently, agents accumulate assets in their younger years and deplete them in old age to shield consumption against the burden of deteriorating health.

7.2. Life-cycle income dynamics

In the baseline model, individuals received a constant income $y > 0$ at each instant $t \in [0, T]$. In reality, however, income follows a hump-shaped pattern, rising early in life, peaking in middle age, and declining as individuals reduce work hours or retire. To illustrate, Table 7 reports income per household member by age of the householder in 2019, using data from the U.S. Census Bureau's Current Population Survey.

The simplest way to capture these hump-shaped income dynamics in the model is to let income y vary over time exogenously according to

$$y(t) = \mu_0 + \mu_1 t + \mu_2 t^2,$$

where μ_0 , μ_1 , and μ_2 are parameters. We calibrate these parameters by running an OLS regression on the data in Table 7. For consistency with the baseline model, we normalize income at age 20 to 2. This ensures that total lifetime income for an individual who reaches the maximum admissible age \bar{T} , $\int_0^{\bar{T}} (\mu_0 + \mu_1 t + \mu_2 t^2) dt$, closely matches that in the baseline case with constant income, $y\bar{T}$. The resulting parameter values are $\mu_0 = 2$, $\mu_1 = 1.4$, and $\mu_2 = -0.33$, yielding a good fit with an R^2 of over 0.7.

In addition, all parameters that are either chosen a priori or estimated from the data remain unchanged. However, the two parameters calibrated within the model, A and B , must be re-estimated. Following the same procedure as before (see subsection 5.3), we

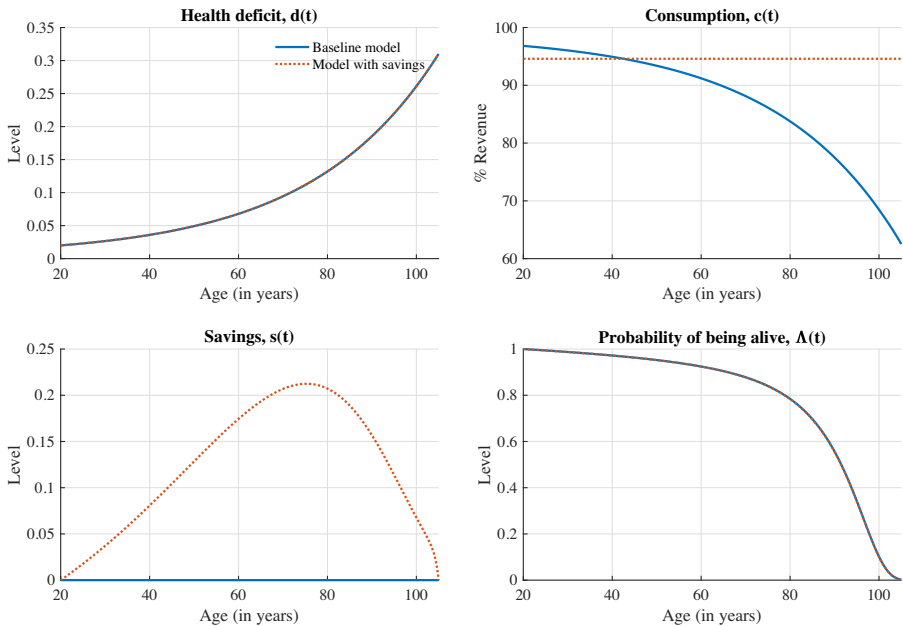


Figure 7. Paths of selected variables over the life-cycle.

Notes: The solid blue lines represent the baseline model with incomplete markets. The dashed red line represent the extended model with complete markets and $r = \rho$.

choose A and B so that the model matches a maximum admissible age of 105 years and a total expenditure on preventive care equal to 0.5% of GDP. The resulting values, $A = 0.0135$ and $B = 4$, are both close to their counterparts in the baseline model.

Figure 8 compares the life-cycle paths of consumption, health deficits, income, and survival probability for an individual who reaches the maximum admissible age, under both the baseline model and the version with time-varying income. As expected, when income follows a hump-shaped pattern (dotted red lines), consumption over the life cycle adopts the classical inverted-U shape. In contrast, the effect on health is limited: health deficits and survival probabilities remain similar across the two models. This is confirmed by the first three moments of the age-at-death distribution, which change only marginally: from 83.9, 15.0, and -1.9 in the baseline model to 84.3, 15.2, and -1.9 with time-varying income.

In sum, assuming constant income over the life cycle does not drive the earlier results. However, another income-related factor deserves attention: inequality. We turn to this next.

7.3. Income heterogeneity

Household income is highly dispersed in the U.S. In 2019, for example, households at the 90th percentile earned more than twelve times as much as those at the 10th percentile.¹⁴ The baseline model assumes a single income level for all individuals (y),

¹⁴Source: U.S. Census Bureau's Current Population Survey Annual Social and Economic Supplement.

Table 7. 2019 Income by age of householder

	Age householder						
	15–24	25–34	35–44	45–54	55–64	65–74	75+
Income (thousand \$)	24.5	33.7	34.3	42.1	47.5	42.7	35.4

Notes: The table reports average annual household income per household member, in thousands of 2020 U.S. dollars, by age group of the householder. Data are from the U.S. Census Bureau, Current Population Survey (CPS), 2020 Annual Social and Economic Supplement. Income includes all sources before taxes and is divided by household size to yield per-member values.

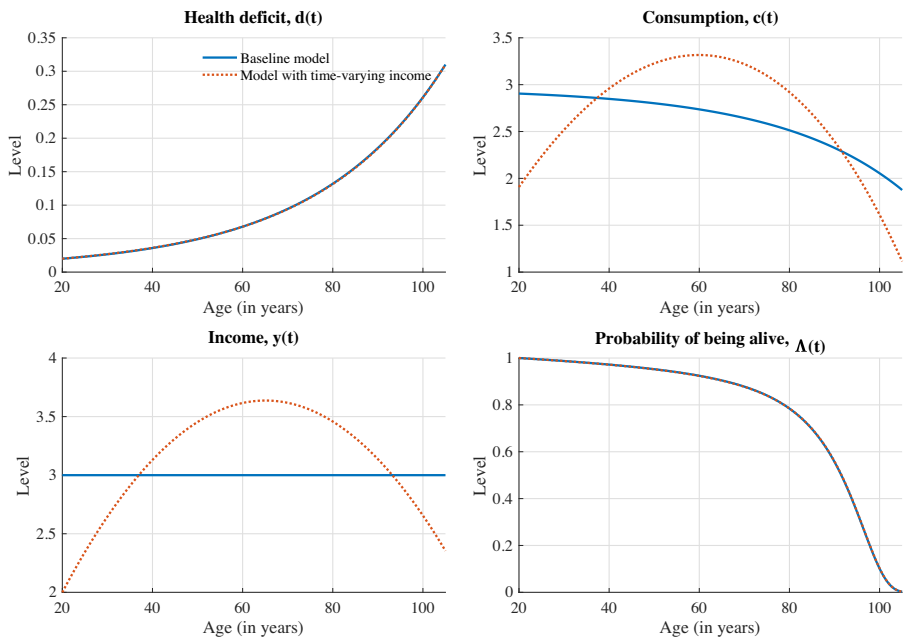


Figure 8. Paths of selected variables over the life-cycle.

Notes: The solid blue lines represent the baseline model with incomplete markets. The dashed red line represent the extended model with life-income dynamics.

abstracting from inequality. To relax this assumption in a simple way, we allow income to vary across agents while remaining constant over their lifetimes. Specifically, each individual's income, y_i , is drawn at birth from a uniform distribution: $y_i \sim U(y - \chi, y + \chi)$, where χ is a parameter. We set $\chi = 0.8$, so that even the lowest-income individuals receive strictly positive utility flows.¹⁵

¹⁵We choose this parametrization because solving our model involves a nonlinear boundary value problem with an endogenous terminal time, which is numerically challenging and does not always yield a solution. This constraint limits how far we can deviate from the parameter space used in the baseline calibration. Moreover, our life-cycle model is not built to match the full extent of income heterogeneity observed in the data. That said, our specification already implies meaningful variation in income, which will allow us to illustrate how heterogeneity might shape our model's insights.

As expected, income heterogeneity generates a gap in life expectancy between rich and poor. For instance, among 20-year-olds, individuals with the highest income ($y + \chi$) live, on average, four years longer than those with the lowest income ($y - \chi$). Indeed, higher income enables greater use of preventive care, which lowers health deficits and reduces mortality risk. The same mechanism also widens the upper bound of the life cycle: the richest individuals have a maximum admissible age, \bar{T} , that is 5 years longer than that of the poorest ones. That said, the first three moments of the age-at-death distribution remain nearly identical to those in the baseline model, with only a modest increase in dispersion of age at death of about two months.

Historically, income and wealth dispersion have increased over recent decades. In 1967, the first year for which data are available, households at the 90th percentile earned about nine times as much as those at the 10th percentile. By 2019, as previously noted, this ratio had risen to over twelve. Similarly, the share of net worth held by the top 0.1% rose from 8.5% in the late 1980s to nearly 14% in recent years.¹⁶ Since income heterogeneity increases the dispersion of age at death in our model, introducing a rising trend in inequality would partially offset the compression in longevity dispersion induced by advances in curative technologies (see Table 4).

There is, however, another trend likely to have compressed the distribution of age at death over time: better access to health care, especially among the poorest households. In the first decades of the 20th century, health insurance was rare. Free care was sometimes available at charity clinics and dispensaries, barter was not uncommon, and many doctors operated a sliding-scale fee system, treating poorer patients for little, while relying on wealthier ones to subsidize care for the community (Chapin, 2015). After World War II, prepayment plans for medical care expanded rapidly. Coverage broadened significantly to include not just hospitalization, but also surgical procedures, physician visits, laboratory services, and prescription drugs. Still, as of 1962, the year of the first NHIS, 30% of Americans remained uninsured. Coverage rates differed markedly by income: only 34% of individuals in the poorest income group had health insurance, compared to 88% among the wealthiest.¹⁷ Since then, the share of uninsured individuals has declined markedly. Landmark public programs like Medicare and Medicaid, introduced in 1965, expanded access for the elderly and low-income populations. More recently, the Affordable Care Act further broadened coverage through Medicaid expansion and subsidized private insurance. As a result, by 2023, the NHIS reported that nearly 93% of Americans had health insurance.

As explained in subsection 6.3, our baseline model falls short of reproducing the full magnitude of the decline in the standard deviation of age at death between 1935 and 2019. Accounting for improved access to medical care would likely improve the model's performance on this dimension. Since our focus is on the role of medical technology, we leave this extension outside the scope of the present paper. Nonetheless, it is a promising avenue for future empirical and theoretical research.

7.4. Alternative parameter values

So far, we have assumed a logarithmic utility function (CRRA utility function with $\sigma = 1$) for three main reasons. First, together with $\beta = 1/2$, it ensures that $\epsilon(t)$ has a unique positive solution. Second, it allows us to compare our results on annual welfare

¹⁶Source: Board of Governors of the Federal Reserve System.

¹⁷Source: National Health Survey 1963, U.S. Department of Health, Education and Welfare.

growth with empirical estimates in Jones and Klenow (2016), which also rely on log utility. Third, survey evidence from Browning et al. (1999) suggests that the inter-temporal elasticity of substitution is only slightly above one, consistent with log utility.

This subsection checks the robustness of our results to alternative values of the inter-temporal elasticity of substitution. Specifically, we retain the CRRA utility function

$$u(c) = \frac{c^{1-\sigma} - 1}{1-\sigma},$$

and vary the parameter σ to explore different levels of inter-temporal substitution.

As discussed earlier and detailed in Appendix B1, solving our model requires addressing several important non-linearities. One such non-linearity concerns the co-state variable $\epsilon(t)$, which captures how the optimal value of the utility function responds to a marginal relaxation of the budget constraint. To ensure that $\epsilon(t)$ is the *unique positive* solution to a quadratic equation, one of the following relationships must be satisfied: (i) $\sigma = 1 - \beta$; (ii) $\sigma = 0.5(1 - \beta)$; or (iii) $\sigma = 2(1 - \beta)$. Outside of these cases, $\epsilon(t)$ is no longer guaranteed to be unique and positive.

Since our baseline calibration satisfies the condition $\sigma = 2(1 - \beta)$, we restrict attention to that class of parametrizations. Specifically, we explore alternative values of σ while adjusting β accordingly to preserve the relationship. To ensure consistency with our previous analysis, we re-calibrate the parameters $\{A, B\}$ for each new pair (σ, β) so that the model continues to match two key targets: a maximum admissible age of 105 years and total preventive care expenditure equal to 0.5% of GDP.

In line with Murphy and Topel (2006), we consider values of the parameter σ ranging from 0.8 to 1.2. Figure 9 compares the life-cycle consumption path of an individual who lives to age 100 under each parametrization. As expected, a lower inter-temporal elasticity of substitution (i.e., higher values of σ) results in stronger consumption smoothing, reducing the decline in consumption later in life. That said, the differences across parametrizations remain mild.

In turn, Table 8 compares the age-at-death distribution and the average annual consumption-equivalent welfare growth (see subsection 6.5) across parametrizations. The key message is clear: none of the indicators change substantially, suggesting that the values of $\{\sigma, \beta, A, B\}$ are not critical to the model's main insights. That said, some differences do emerge across parametrizations. For instance, life expectancy increases in the high- σ scenario. Pinpointing the exact mechanism is difficult, as each parametrization involves simultaneous adjustments to four parameters. However, one consistent pattern is that higher σ values are associated with lower values of B . This reduces the monetary cost of health deficits, freeing up resources for preventive care and thereby supporting longer lifespans.

In addition to the alternative parametrizations discussed above, we explored two further modifications. First, we considered a higher annual time discount factor (0.975 instead of the baseline value of 0.95). Second, when calibrating A and B , we targeted total medical expenditures equal to 15% of GDP, as observed in the data (see Table 4), rather than focusing on preventive care expenditures of 0.5% of GDP. In both cases, our core results remain unchanged. For brevity, we do not report the full results here, but they are available upon request.

All told, we are confident that our findings are not dependent on particular parameter values, but hold instead for a wide parameter space.




Figure 9. Consumption over the life-cycle.

Notes: The solid blue line represents the baseline parametrization. The dotted red and dashed yellow lines vary σ while adjusting β according to $\sigma = 2(1 - \beta)$.

8. Discussion

In this final section, we explore how relaxing some of our modeling assumptions –such as rational expectations– might affect our main insights. By its nature, the discussion raises more questions than it answers, but it helps clarify the scope of our framework and the interpretation of our results.

8.1. Rational expectations

Our model emphasizes dynamic optimization, making individuals' expectations crucial. Following standard practice, we assume rational expectations: individuals hold probability beliefs that coincide with those implied by the model. When allocating resources between consumption and preventive care, they understand precisely how each additional unit of prevention affects mortality risk, both at time t and in all future periods $\tau \in [t, \bar{T}]$. As a result, their choices are optimal in the formal, mathematical sense.

However, rational expectations is a strong assumption. Empirical studies show that consumers often struggle to make informed choices when purchasing health insurance (see, e.g., Harris and Keane, 1998; McFadden, 2006; Frank and Zeckhauser, 2009). Difficulties are particularly pronounced among the elderly, who, for example, frequently find it hard to navigate Medicare and Medigap rules (Harris-Kojetin et al., 2001; McCormack et al., 2001). Similar challenges arise beyond insurance markets: patients face difficulties evaluating options related to physicians, drugs, or treatments, and tend to disregard technical information they do not fully understand (see Keane and Thorp, 2016, and references therein).

Replacing the rational expectations assumption with a weaker alternative may thus be worthwhile. However, it is unclear what form such an alternative should take, given the

Table 8. Age-at-death distribution and welfare growth under alternative parametrizations

	Distribution of age at death				Annual welfare growth
	μ_T	σ_T	Skewness	$\Lambda(\bar{T})$	1935–2019 (%)
Data	79.5	16.7	−1.6	0.5%	–
$\sigma = 1.0, \beta = 0.5$	83.9	15.0	−1.9	0.3%	0.9
$\sigma = 0.8, \beta = 0.6$	83.8	14.8	−1.9	0.3%	1.0
$\sigma = 1.2, \beta = 0.4$	85.3	15.2	−1.9	0.3%	0.8

Notes: Expected value, standard deviation, skewness, and share of individuals living to age 105 in both the 2019 US data and the model. In the model, we compute the moments using a 10,000-agent Monte Carlo simulation. μ_T and σ_T are expressed in years. The final column reports average annual consumption-equivalent welfare growth from 1935 to 2019, based on the methodology described in subsection 6.5.

“wilderness” of competing models of expectations formation (see Angeletos et al., 2021). A simple, naive alternative within our model is to assume that individuals do not understand how their health impacts mortality risk. While the true hazard rate satisfies $\lambda_d(d) > 0$, individuals could assume $\lambda_d(d) = 0$, ignoring the link between health and survival. As a result, their probability beliefs would no longer align with those implied by the model. Not surprisingly, this “myopic” specification leads to lower preventive care use ($H_2 = 0.1\%$ compared to $H_2 = 0.6\%$ in the baseline), as individuals fail to recognize its benefits. This lower preventive care shortens life expectancy by about three years.

Our paper focuses on how medical advances have shaped mortality risk over the past century. For this purpose, we are comfortable adopting the rational expectations hypothesis, as we see no alternative that is both well-founded and internally consistent. That said, relaxing rational expectations could be needed to study other questions, such as why most Americans consume less preventive care than recommended – an area to which our model naturally lends itself.¹⁸

8.2. Stationary equilibria

Throughout the paper, we compare stationary equilibria, that is, steady state equilibria, as in Hosseini et al. (2021), Strulik (2022), De Nardi et al. (2024), and Chen et al. (2025). Specifically, we solve the model using values of $\{A, \lambda, \theta, y\}$ for selected years, such as 1935 and 2019, and compute key statistics from the corresponding steady state. An alternative approach would be to simulate the full transitional dynamics from 1935 to 2019. Beyond its computational demands, this would require strong assumptions about how individuals form expectations over time, as well as data that do not exist – issues we briefly discuss below.

One option assumes that individuals born at time t perfectly anticipate how medical technology (A, λ), healthcare prices (θ), and real median income (y) will evolve over their entire potential lifespans $t + a$ with $a \in [0, \bar{T}]$, and the maximum attainable age \bar{T} itself depending on these trends. With this perfect foresight, they would choose at time t the optimal life-cycle paths for consumption and preventive care, with no need or re-optimizing later in life. At the other end of the spectrum, individuals could be myopic: when making

¹⁸Nearly half of all deaths each year are linked to preventable risk factors, and recent studies suggest that increased preventive care use could add over 2 million healthy life-years (see Taksler et al., 2022, and references therein).

choices at time t , they assume $\{A, \lambda, \theta, y\}$ will remain constant at their current levels indefinitely. However, at every time $t + dt$, they observe that $\{A, \lambda, \theta, y\}$ has in fact evolved, which force them to update their beliefs to reflect the new levels and re-optimize accordingly. Lastly, the most mathematically elegant, yet also most complex, approach assumes that $\{A, \lambda, \theta, y\}$ evolve stochastically rather than deterministically. This would require abandoning calculus of variations in favor of dynamic programming: at each instant t , individuals with rational expectations observe the current values of the four trends and choose optimal controls accordingly. This still requires individuals to have full knowledge of the probability laws governing $\{A, \lambda, \theta, y\}$, which is by no means a weak assumption.

Regardless of the approach, studying transitional dynamics complicates aggregation across individuals relative to the stationary case. At any given time t , the population consists of a distribution of cohorts of age a , that is, born at time $t - a$, each exposed since birth to different values of $\{A, \lambda, \theta, y\}$ and therefore facing different mortality risks.

As a result, simulating transitional dynamics requires information that does not exist. For example, the age-at-death distribution in 1935 depends on cohorts born roughly a century earlier, whose choices were shaped by the (perceived) evolution of $\{A, \lambda, \theta, y\}$ from 1835 to 1935. Likewise, the distribution in 2019 reflects individuals making decisions based on beliefs about medical advances, prices, and income over the coming century. To our knowledge, reliable historical data on these trends from approximately 1835 to 1925, along with credible forecasts out to 2120, are unavailable, making the simulation of full transitional dynamics challenging.

For all the reasons mentioned above, considering transitional dynamics over a shorter horizon (e.g., from 1980 to 2000), while still difficult, would be more manageable. However, the effects of medical advances unfold over long periods, often shaping health, longevity, and behavior across multiple decades. This motivates our focus on century-scale comparisons of stationary equilibria.

8.3. Endogenous curative expenditures

Our model distinguishes between two types of health expenditures: preventive care, which slows the accumulation of health deficits, and the monetary costs associated with the health deficit, which can be interpreted as palliative or LTC expenditures, as well as some types of curative expenditures. It is thus natural to treat preventive care as endogenous – a choice made by households, shaped by preferences and perceived health risks (Grossman, 1972; Kenkel, 1990). In contrast, palliative care, LTC and curative expenditures are often unavoidable (Gertler and Gruber, 2002; Wagstaff, 2007). As Ozkan (2025) notes, poorer individuals tend to use less preventive care early in life but more curative care later, reinforcing the idea that prevention is a choice, while treatment is a necessity. In line with this evidence, our model treats preventive care h as a control variable and assumes the monetary cost of health deficits (θBd) is fully exogenous.

For completeness, let us now assume that individuals also choose the monetary costs associated with health deficits. In practice, this captures the possibility of individuals deciding the intensity of palliative, LTC or some curative expenditures (we simply refer to all of them as LTC expenditures). To accommodate this additional control variable, we slightly modify our setup. Specifically, we assume that health deficits generate utility losses each period, which can be mitigated by LTC resources k to make daily lives easier. Mathematically, the budget constraint (3) and the expected lifetime utility (4) become

$$c(t) + \theta(h(t) + k(t)) = y,$$

$$\int_0^T e^{-\rho t} \Lambda(t) (\ln c(t) - g(d(t), k(t))) dt,$$

where the function $g(d, k)$ represents the instantaneous disutility from the deficit and satisfies: $g > 0$, $g_d > 0$, $g_k < 0$, $g_{kk} > 0$, and $g_{dk} < 0$. These properties are intuitive: the deficit increases disutility; LTC expenditures decrease disutility both directly and by mitigating the effect of d ; and there are diminishing returns to LTC expenditures. The key difference between h and k is that h permanently reduces the deficit, while k has only a temporary effect on utility.¹⁹

We present the full solution in Appendix B3, which closely resembles that of the baseline model. The key difference is that we replace the exogenous health expenditures $k = \theta Bd$ with a first-order condition

$$-g_k = \theta \epsilon, \quad (10)$$

where ϵ is the shadow price related to the budget constraint (equal to the marginal utility of consumption, see Appendix B3). This condition states that, at equilibrium, the marginal return on LTC expenditures equals their marginal cost. Fully differentiating (10) under the assumption of constant ϵ gives²⁰

$$\frac{dk}{dd} = -\frac{g_{kd}}{g_{kk}},$$

which is strictly positive given the assumed properties of $g(d, k)$. In other words, LTC expenditures increase with d , as in our baseline model.

To further illustrate this, let us define $g(d, k) = /k$, which satisfies the conditions above. The first-order condition (10) then becomes

$$k = \left(\frac{Bdc}{\theta} \right)^{\frac{1}{2}},$$

where we use that $\epsilon = 1/c$. We see that k increases with dc . Earlier, Fig. 7 showed that in the baseline model, the rise in d over the life cycle was larger than the drop in c . As a result, the optimal k would likely go up with age too, in line with the baseline setup. Moreover, in our extension with complete markets (Section 7.1), ϵ is constant, and k would unambiguously rise over time. We thus conclude that allowing LTC expenditures to be endogenous should not change our main results.

Finally, solving the model with endogenous LTC expenditures is complex, as ϵ is no longer a function of the state and co-state variables (see Appendix B3).²¹

¹⁹Since we calibrate $\sigma = 1$, we immediately simplify the CRRA utility into a log. Moreover, we use a simple additive instantaneous utility function here, but a more general multiplicative utility function $G(c, d, h)$ could also be considered.

²⁰Individuals smooth consumption over their life cycle and the assumption of constant ϵ is therefore reasonable.

²¹Another way would be to make LTC spending a choice and keep preventive spending fixed. This would let us write an exact formula for ϵ , but we think this is less realistic (see above). Also, it would not explain how to model health investments over time, which would still have to be set outside the model.

8.4. Childhood

Following Cole et al. (2019) and De Nardi et al. (2024), among many others, our setup abstracts from childhood and focuses on life-cycle dynamics from age 20 onward. We leave out children for simplicity. Including them would further link cohorts, as parents born at time t_1 would need to make decisions and incur costs for children born at time $t_2 > t_1$.

Excluding children could be problematic if healthcare use patterns shifted sharply at age 18. In practice, they do not appear to do so. According to 2019 data from the National Center for Health Statistics (Ashman et al., 2023), office-based physician visit rates rise steadily with age: from 1.5 visits per year for children aged 1 to 17 to 7.08 for those 65 and over. The reasons for visits evolve similarly: preventive care declines with age (40% of visits under 18, 13% at 65+), as do visits for new problems (38% under 18, 17% at 65+), while visits for chronic conditions increase (12% under 18, 56% at 65+), as do those related to pre- or post-surgical care (2% under 18, 7% at 65+). Injury is the only major reason for physician visits that does not exhibit a clear age trend.

Nonetheless, children have benefited the most from several major advances in preventive care over the past century, especially vaccines. For example, Shattock et al. (2024) estimate that since 1974, vaccination has averted 154 million deaths worldwide, including 146 million among children younger than 5 years of whom 101 million were infants younger than 1 year. Actuarial life tables from the Human Mortality Database also reflect the dramatic decline in child mortality, particularly for infants under age 1: in 1935, a newborn faced a 6% probability of dying within the first year of life; by 2019, that figure had fallen to just 0.5%.

Our model links the hazard rate of death not to age itself but to the health deficit, which increases monotonically with age. Therefore, to the extent that mortality risk rises with age in the data, excluding children should have limited consequences. This holds for those aged 1 to 17, whose mortality risk increases steadily with age. In this case, instead of calibrating d_0 to match the health deficit of the median 20-year-old, we could have matched that of a 5- or 10-year-old. Our findings would be unchanged. However, it does not hold for infants under age 1, who face substantially higher mortality than older children, despite (presumably) having lower health deficits. Moreover, as noted above, this group has benefited most from advances in preventive care. Excluding them may therefore lead the model to understate the full impact of such advances. One possibility would be to introduce a more sophisticated function for the hazard rate of death as $\lambda(t, A, d(t)) = \lambda_1(t, A) + \lambda_2(d(t))$. Here, $\lambda_1(\cdot)$ would only hold for the very first years in life and be decreasing and convex in t , as well as decreasing in A . $\lambda_2(\cdot)$ would be the usual hazard rate increasing in deficits. A more detailed treatment of this issue is left for future research.

9. Conclusion

Over the past century, advances in preventive care have reduced disease incidence, while innovations in curative treatments have lowered mortality risks. These medical breakthroughs might have affected the age-at-death distribution differently, for they shape both individual incentives and the link between health and mortality in distinct ways. Our quantitative analysis shows that although both forms of medical progress have boosted life expectancy since 1935, it is curative treatments that account for the narrowing of the age-at-death distribution.

In addition to the avenues for future research already highlighted in the main text, two further directions seem particularly promising. First, future work will explore the mechanisms underlying the well-documented but poorly understood relationship

between income and life expectancy (Deaton, 2016). A deeper understanding of this link could shed light on the historical evolution of the age-at-death distribution and help guide the design of more effective public health policies. Second, another important question concerns the optimal allocation of scarce resources between developing preventive and curative medical technologies. Our current framework – partial equilibrium in nature and treating medical technologies as exogenous – cannot address this issue directly. Nonetheless, it offers a useful foundation: it is tractable, empirically grounded, and well-suited for further extensions.

Supplementary material. To view supplementary material for this article, please visit <https://doi.org/10.1017/dem.2025.10009>.

Competing interests. The authors declare none.

References

- Acemoglu, D. and S. Johnson (2007). Disease and development: The effect of life expectancy on economic growth. *Journal of Political Economy* 115(6), 925–985.
- Acemoglu, D. and P. Restrepo (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review* 108(6), 1488–1542.
- Alsan, M., V. Atella, J. Bhattacharya, V. Conti, I. Mejía-Guevara, and G. Miller (2021). Technological Progress and Health Convergence: The Case of Penicillin in Postwar Italy. *Demography* 58(4), 1473–1498.
- Anderson, G. F., P. Hussey, and V. Petrosyan (2019). It's Still The Prices, Stupid: Why The US Spends So Much On Health Care, And A Tribute To Uwe Reinhardt. *Health Affairs* 38(1), 87–95.
- Anderson, G. F., U. E. Reinhardt, P. S. Hussey, and V. Petrosyan (2003). It's The Prices, Stupid: Why The United States Is So Different From Other Countries. *Health Affairs* 22(3), 89–105.
- Angeletos, G.-M., Z. Huo, and K. A. Sastry (2021). Imperfect Macroeconomic Expectations: Evidence and Theory. *NBER Macroeconomics Annual* 35(1), 1–86.
- Ashman, J. J., L. Santo, and T. Okeyode (2023). Characteristics of office-based physician visits by age, 2019. Technical Report 184, National Center for Health Statistics.
- Autor, D. H., D. Dorn, and G. H. Hanson (2013). The China Syndrome: Local Labor Market Effects of Import Competition in the United States. *American Economic Review* 103(6), 2121–2168.
- Baltagi, B. H., R. Lagravinese, F. Moscone, and E. Tosetti (2017). Health Care Expenditure and Income: A Global Perspective. *Health Economics* 26(7), 863–874.
- Benmelech, E., N. K. Bergman, and H. Kim (2022). Strong employers and weak employees. *Journal of Human Resources* 57(S), S200–S250.
- Bergeron-Boucher, M.-P., V. Canudas-Romo, and M. Ebeling (2015). Decomposing changes in life expectancy: Compression versus shifting mortality. *Demographic Research* 33(14), 391–424.
- Blanchard, O. J. (1985). Debt, deficits, and finite horizons. *Journal of Political Economy* 93(2), 223–247.
- Boukas, E., A. Haurie, and P. Michel (1990). An optimal control problem with a random stopping time. *Journal of Optimization Theory and Applications* 64(3), 471–480.
- Browning, M., L. P. Hansen, and J. J. Heckman (1999). Chapter 8. *Micro data and general equilibrium models*. Volume 1 of Handbook of Macroeconomics, pp. 543–633. Elsevier.
- Bucke, R. (1898). Surgery among the insane in Canada. *American Journal of Insanity* 55(1), 1–19.
- Canudas-Romo, V. (2008). The modal age at death and the shifting mortality hypothesis. *Demographic Research* 19(30), 1179–1204.
- Chapin, C. F. (2015). *Ensuring America's Health: The Public Creation of the Corporate Health Care System*. Cambridge University Press.
- Chen, C., Z. Feng, and J. Gu (2025). Health, health insurance, and inequality. Forthcoming in *International Economic Review*.
- Chernew, M., D. M. Cutler, K. Ghosh, and M. B. Landrum (2016). Understanding the Improvement in Disability-Free Life Expectancy in the US Elderly Population. In *Insights in the Economics of Aging, NBER Chapters*, pp. 161–201. National Bureau of Economic Research, Inc.

- Chetty, R., M. Stepner, S. Abraham, S. Lin, B. Scuderi, N. Turner, A. Bergeron, and D. Cutler (2016). The Association Between Income and Life Expectancy in the United States, 2001–2014. *JAMA* 315(16), 1750–1766.
- Cheung, S. L. K. and J.-M. Robine (2007). Increase in common longevity and the compression of mortality: the case of Japan. *Population Studies* (Cambridge) 61(1), 85–97.
- Cheung, S. L. K., J.-M. Robine, F. Paccard, and A. Marazzi (2009). Dissecting the compression of mortality in Switzerland, 1876–2005. *Demographic Research* 21(19), 569–598.
- Christensen, K., G. Doblhammer, R. Rau, and J. W. Vaupel (2009). Ageing populations: the challenges ahead. *The Lancet* 374(9696), 1196–1208.
- Cole, H. L., S. Kim, and D. Krueger (2019). Analysing the Effects of Insuring Health Risks: On the Trade-off between Short-Run Insurance Benefits versus Long-Run Incentive Costs. *The Review of Economic Studies* 86(3), 1123–1169.
- Cutler, D., A. Deaton, and A. Lleras-Muney (2006). The determinants of mortality. *Journal of Economic Perspectives* 20(3), 97–120.
- Cutler, D. and G. Miller (2005). The role of public health improvements in health advances: The twentieth-century United States. *Demography* 42(1), 1–22.
- Dalgaard, C.-J. and H. Strulik (2014). Optimal Aging And Death: Understanding The Preston Curve. *Journal of the European Economic Association* 12(3), 672–701.
- De Nardi, M., E. French, J. B. Jones, and J. McCauley (2016). Medical Spending of the US Elderly. *Fiscal Studies* 37, 717–747.
- De Nardi, M., S. Pashchenko, and P. Porapakarm (2024). The lifetime costs of bad health. Forthcoming in *The Review of Economic Studies*.
- De Nardi, M. C., E. French, and J. B. Jones (2010). Why Do the Elderly Save? The Role of Medical Expenses. *Journal of Political Economy* 118(1), 39–75.
- Deaton, A. (2016). On Death and Money: History, Facts, and Explanations. *JAMA* 315(16), 1703–1705.
- Dragone, D. and H. Strulik (2020). Negligible senescence: An economic life cycle model for the future. *Journal of Economic Behavior & Organization* 171(C), 264–285.
- Edwards, R. D. and S. Tuljapourkar (2005). Inequality in Life Spans and a New Perspective on Mortality Convergence Across Industrialized Countries. *Population and Development Review* 31(4), 645–674.
- Ehrlich, I. (2000). Uncertain lifetime, life protection, and the value of life saving. *Journal of Health Economics* 19(3), 341–367.
- Ehrlich, I. and G. S. Becker (1972). Market Insurance, Self-Insurance, and Self-Protection. *Journal of Political Economy* 80(4), 623–648.
- Ewbank, D. C. and S. H. Preston (1990). Personal health behaviour and the decline in infant and child mortality: The United States, 1900–1930. In J. Caldwell, S. Findley, P. Caldwell, G. Santow, W. Cosford, J. Braid, and D. Broers-Freeman (Eds.), *What We Know About Health Transition: The Culture, Social and Behavioural Determinants of Health*, pp. 116–149. Canberra: Health Transition Center, Australian National University.
- Ferrie, J. P. and W. Troesken (2008). Water and Chicago's mortality transition, 1850–1925. *Explorations in Economic History* 45(1), 1–16.
- Fonseca, R., P.-C. Michaud, T. Galama, and A. Kapteyn (2020). Accounting for the Rise of Health Spending and Longevity. *Journal of the European Economic Association* 19(1), 536–579.
- Frank, R. G. and R. J. Zeckhauser (2009). Health Insurance Exchanges—Making the Markets Work. *New England Journal of Medicine* 361(12), 1135–1137.
- Frankovic, I. and M. Kuhn (2019). Access to health care, medical progress and the emergence of the longevity gap: A general equilibrium analysis. *The Journal of the Economics of Ageing* 14(C).
- Frankovic, I. and M. Kuhn (2023). Health insurance, endogenous medical progress, health expenditure growth, and welfare. *Journal of Health Economics* 87(C).
- Fries, J. F. (1980). Aging, natural death, and the compression of morbidity. *New England Journal of Medicine* 303(3), 130–135.
- Fuchs, V. R. (2010). New priorities for future biomedical innovations. *New England Journal of Medicine* 363(8), 704–706.
- García Sanchez, P., L. Marchiori, and O. Pierrard (2025). Long-term care expenditures and investment decisions under uncertainty. *Journal of Population Economics* 38(1), 4.
- Gertler, P. and J. Gruber (2002). Insuring consumption against illness. *American Economic Review* 92(1), 51–70.

- Grossman, M. (1972). On the Concept of Health Capital and the Demand for Health. *Journal of Political Economy* 80(2), 223–255.
- Hall, R. E. and C. I. Jones (2007). The Value of Life and the Rise in Health Spending. *The Quarterly Journal of Economics* 122(1), 39–72.
- Harris, K. M. and M. P. Keane (1998). A model of health plan choice: Inferring preferences and perceptions from a combination of revealed preference and attitudinal data. *Journal of Econometrics* 89(1), 131–157.
- Harris-Kojetin, L., L. McCormack, E. Jael, J. A. Sangl, and S. Garfinkel (2001). Creating more effective health plan quality reports for consumers: Lessons from a synthesis of qualitative testing. *Health Services Research* 36(3), 447–476.
- Hood, C. M., K. P. Gennuso, G. R. Swain, and B. B. Catlin (2016). County health rankings: Relationships between determinant factors and health outcomes. *American Journal of Preventive Medicine* 50(2), 129–135.
- Hosseini, R., K. A. Kopecky, and K. Zhao (2021). *How Important Is Health Inequality for Lifetime Earnings Inequality?* FRB Atlanta Working Paper 2021-1, Federal Reserve Bank of Atlanta.
- Hu, T.-W., H.-Y. Sung, and T. E. Keeler (1995). Reducing cigarette consumption in California: Tobacco taxes vs an anti-smoking media campaign. *American Journal of Public Health* 85(9), 1218–1222.
- Hutchings, M. I., A. W. Truman, and B. Wilkinson (2019). Antibiotics: past, present and future. *Current Opinion in Microbiology* 51, 72–80. Antimicrobials.
- Hwang, J. and S. Kim (2023). Unexpected longevity, intergenerational policies, and fertility. *Journal of Population Economics* 36, 1607–1640.
- Jayachandran, S., A. Lleras-Muney, and K. V. Smith (2010). Modern medicine and the twentieth century decline in mortality: Evidence on the impact of sulfa drugs. *American Economic Journal: Applied Economics* 2(2), 118–146.
- Jones, C. I. and P. J. Klenow (2016). Beyond gdp? welfare across countries and time. *American Economic Review* 106(9), 2426–57.
- Kannisto, V. (2000). Measuring the compression of mortality. *Demographic Research* 3(6).
- Keane, M. P. and S. Thorp (2016). Complex decision making: The roles of cognitive limitations, cognitive decline, and aging. Volume 1 of *Handbook of the Economics of Population Aging*, pp. 661–709. North-Holland.
- Kenkel, D. (1990). Consumer health information and the demand for medical care. *The Review of Economics and Statistics* 72(4), 587–595.
- Lichtenberg, F. R. (2018). The Impact of New Drug Launch on Life-Years Lost in 2015 from 19 Types of Cancer in 36 Countries. *JODE - Journal of Demographic Economics* 84(3), 309–354.
- Lucas, R. (1987). *Models of Business Cycles*. Basil Blackwell.
- Marmot, M. and J. J. Allen (2014). Social determinants of health equity. *American Journal of Public Health* 104(Suppl 4), S517–S519.
- McCormack, L. A., S. A. Garfinkel, J. H. Hibbard, E. C. Norton, and U. J. Bayen (2001). Health plan decision making with new medicare information materials. *Health Services Research* 36(3), 531–554.
- McFadden, D. (2006). Free markets and fettered consumers. *The American Economic Review* 96(1), 3–29.
- Miller, B. and S. Nyce (2023). Healthcare USA: The Big Paycheck Squeeze. *Insider* 33(5).
- Mishan, E. J. (1971). Evaluation of life and limb: A theoretical approach. *Journal of Political Economy* 79(4), 687–705.
- Mitnitski, A., A. Mogilner, C. MacKnight, and K. Rockwood (2002). The accumulation of deficits with age and possible invariants of aging. *The Scientific World Journal* 2002(2), 1816–1822.
- Murphy, K. and R. Topel (2006). The value of health and longevity. *Journal of Political Economy* 114(5), 871–904.
- Newhouse, J. P. (1992). Medical Care Costs: How Much Welfare Loss? *Journal of Economic Perspectives* 6(3), 3–21.
- Nichols, R. L. (2004). Preventing surgical site infections. *Clinical Medicine and Research* 2(2), 115–118.
- Nielsen, S., M. Khalek, C. Benn, P. Aaby, and S. Hanifi (2021). National immunisation campaigns with oral polio vaccine may reduce all-cause mortality: Analysis of 2004–2019 demographic surveillance data in rural bangladesh. *EclinicalMedicine* 36, 100886.
- Oeppen, J. and J. W. Vaupel (2002). Broken limits to life expectancy. *Science* 296(5570), 1029–1031.
- Ozkan, S. (2025). Income differences and health disparities: Roles of preventive vs. curative medicine. *Journal of Monetary Economics* 150, 103698. *Labor Market Dynamics: Drivers and Consequences*.

- Pashchenko, S. and P. Porapakkarm (2017). Work incentives of medicaid beneficiaries and the role of asset testing. *International Economic Review* 58(4), 1117–1154.
- Pestieau, P. and G. Ponthiere (2016). Longevity Variations and the Welfare State. *JODE - Journal of Demographic Economics* 82(2), 207–239.
- Riley, J. C. (2001). *Rising Life Expectancy: A Global History*. Cambridge University Press.
- Rockwood, K. and A. Mitnitski (2007). Frailty in relation to the accumulation of deficits. *The Journals of Gerontology: Series A, Biological Sciences and Medical Sciences* 62(7), 722–727.
- Schelling, T. C. (1968). The life you save may be your own. In S. B. Chase (Ed.), *Problems in Public Expenditure Analysis*, pp. 127–162. Washington, D.C.: Brookings Institute.
- Scott, A. J. (2023). The economics of longevity—an introduction. *The Journal of the Economics of Ageing* 24, 100439.
- Seierstad, A. (2009). *Stochastic Control in Discrete and Continuous Time*. Boston, MA: Springer.
- Shampine, L. F., I. Gladwell, and S. Thompson (2003). *Solving ODEs with MATLAB*. Cambridge University Press.
- Shattock, A. J., H. C. Johnson, S. Y. Sim, A. Carter, P. Lambach, R. C. W. Hutubessy, K. M. Thompson, K. Badizadegan, B. Lambert, M. J. Ferrari, M. Jit, H. Fu, S. P. Silal, R. A. Hounsell, R. G. White, J. F. Mosser, K. A. M. Gaythorpe, C. L. Trotter, A. Lindstrand, K. L. O'Brien, and N. Bar-Zeev (2024). Contribution of Vaccination to Improved Survival and Health: Modelling 50 Years of the Expanded Programme on Immunisation. *The Lancet* 403(10441), 2307–2316.
- Soares, R. R. (2015). Gary Becker's Contribution in Health Economics. *JODE - Journal of Demographic Economics* 81(1), 51–57.
- Stansbury, A. and L. H. Summers (2020). The Declining Worker Power Hypothesis: An Explanation for the Recent Evolution of the American Economy. *Brookings Papers on Economic Activity* 51(1), 1–96.
- Strulik, H. (2015). Frailty, mortality, and the demand for medical care. *The Journal of the Economics of Ageing* 6, 5–12.
- Strulik, H. (2022). Medical progress and life cycle choices. *The Journal of the Economics of Ageing* 23(C).
- Taksler, G. B., E. R. Pfoh, K. A. Martinez, M. M. Sheehan, N. M. Gupta, and M. B. Rothberg (2022). Comparison of National Data Sources to Assess Preventive Care in the US Population. *Journal of General Internal Medicine* 37(2), 318–326.
- Vaupel, J. W., Z. Zhang, and A. A. van Raalte (2011). Life expectancy and disparity: an international comparison of life table data. *BMJ Open* 1(1).
- Wagstaff, A. (2007). The economic consequences of health shocks: Evidence from Vietnam. *Journal of Health Economics* 26(1), 82–100.
- Wang, H.-C., J. C. Yue, T.-C. Chang, and T.-C. Chang (2023). Morbidity compression and cancer insurance. *JODE - Journal of Demographic Economics* 89(3), 465–482.
- Yaari, M. E. (1965). Uncertain Lifetime, Life Insurance, and the Theory of the Consumer. *Review of Economic Studies* 32(2), 137–150.
- Zhang, J., J. Zhang, and R. Lee (2001). Mortality decline and long-run economic growth. *Journal of Public Economics* 80(3), 485–507.