

Optimal Aging and Death: Understanding the Preston Curve*

Carl-Johan Dalgaard[†] and Holger Strulik^{**}

First Version: July 2010. This Version: March 2012.

Abstract. Does prosperity lead to greater longevity? If so, what is the strength of the income channel? To address these questions we develop a model of optimal lifetime consumption, in which households are subject to physiological aging. In modeling aging we draw on recent research in the fields of biology and medicine. The speed of the aging process, and thus the time of death, are endogenously determined by optimal health investments. A calibrated version of the model accounts well for the observed non-linear cross-country link between longevity and income, also known as “the Preston Curve”.

Keywords: Aging, Longevity, Health investments, Savings, Preston Curve.

JEL: D91, J17, J26, I12.

*We would like to thank David Canning, James Carey, Francesco Caselli, Karl Dietrich, Leonid Gavrilov, Natalia Gavrilova, Michael Grimm, Alexia Fürnkranz-Prskawetz, Oded Galor, Chad Jones, Soeren Leth-Petersen, Arnold Mitnitski, Fidel Perez-Sebastian and seminar participants at Brown University, Harvard University, the German council for population economics meeting in Mannheim, and at the LEPAS workshops in Vienna and Alicante for helpful comment. This research was funded by the European Commission within the project “Long-Run Economic Perspectives of an Aging Society” (LEPAS) in the Seventh Framework Programme under the Socio-economic Sciences and Humanities theme (Grant Agreement: SSH7-CT-2009-217275)

[†]Department of Economics, University of Copenhagen, Oester Farimagsgade 5, building 26, DK-1353 Copenhagen, Denmark; email: carl.johan.dalgaard@econ.ku.dk.

^{**}University of Hannover, Wirtschaftswissenschaftliche Fakultät, Königsworther Platz 1, 30167 Hannover, Germany; email: strulik@vwl.uni-hannover.de.

1. INTRODUCTION

In a seminal contribution Samuel Preston (1975) documented a striking empirical fact: In a cross-section of countries higher levels of income per capita are associated with greater life expectancy; the curve that fits the data best – dubbed “the Preston curve” by later research – is concave. This discovery served to promote the idea that prosperity leads to greater longevity, and that income inequality works to lower average longevity; the latter being a consequence of the observed concavity of the mapping between income and life expectancy.¹

These are powerful ideas, which continue to be influential. As Bloom and Canning (2007, p. 498) observe: “Samuel H. Prestons classic paper, ‘The Changing Relation between Mortality and Level of Economic Development’, published in 1975, remains a cornerstone of both global public health policy and academic discussion of public health.” Despite its prominence, however, the interpretation of the Preston curve continues to be shrouded in mystery. Does the slope of the curve reflect a causal impact of income *per se*? Or is income perhaps a stand-in for other underlying correlates with longevity and income, such as efficiency differences in health care? Deaton (2003, p. 152) formulates the key challenge for research in this area succinctly: “If income is indeed directly protective, we need to know whether the effect is really nonlinear [...] because it is this [...] that determines whether and by how much income redistribution can improve population health”. The objective of the present research is to offer some progress in this regard.

We employ a theory-driven approach in order to elucidate the mechanisms underlying the shape and position of the Preston curve. In order to understand human longevity we need to understand the determinants of the process which *culminates* in death: the aging process (Arking, 2006). We therefore develop a life-cycle model for a representative consumer who is subject to physiological aging. In modeling the aging process we draw on recent research in the fields of biology and medicine. This approach has the advantage that estimates for key

¹A second observation made in the paper was that the curve shifts upwards over time, implying greater longevity at all levels of income per capita. Preston hypothesized that these shifts represented improvements in health technology (broadly defined), and noted that the shifts accounted for the lions share of global improvements in longevity over time.

physiological parameters influencing the aging process are available to us; this is what facilitates a credible calibration of life expectancy. In our model the individual consumes, saves, and makes deliberate investments in slowing down the aging process, thereby postponing the “date of expiry”. This model allows us to study the impact from, for example, changes in income and health care efficiency on changes in longevity.

With the model in place we subsequently calibrate it to US data, and proceed to examine the calibrated model’s ability to account for the *cross-country* income gradient in the year 2000, as reflected in the Preston curve. Accordingly, using income data for a cross section of countries as “input” we ask the model to predict the associated levels of life expectancy for those countries, which we then compare to the estimated income gradient: the Preston curve. Since the model is calibrated for the US such that income matters for health and life-expectancy, it is unsurprising that it predicts a positive correlation between income and life-expectancy across countries. What is striking, however, is that the model does a good job at predicting the position and concave slope of the Preston curve; about 80% of variation nested in the Preston curve can be accounted for by the model. That is, the pure income effect on longevity, mediated by health investments, goes a long way towards explaining the Preston curve.

What should one make of the left-over residual (i.e., the remaining 20% of the variation captured by the Preston curve)? When we predict longevity using our model we are keeping the relative price of health goods (to non-health goods) in efficiency units fixed at the US level. Empirically, the (efficiency adjusted) price of health is likely correlated with income across countries. Such variation can, according to our model, account for the residual. While the model does not allow us to decompose this left-over residual, it undoubtedly captures a range of factors such as cross-country variation in relative prices of health goods; climate (influences disease prevalence, and thus attainable longevity at given health investment rate); health care institutions (influences how many years a given dollar amount of investment can buy) and more.² The main message from the analysis, however, is that the Preston curve

²The Preston curve may also be capturing - in part at least - an impact of longevity on income. Since our model deliberately abstracts from this causal mechanism, instead distilling the power of the income-to-longevity mechanism, this sort of influence would also go into the left-over residual. At the same time it is

indeed seems to capture a strong - nonlinear - causal effect of income on longevity. To put it in popular terms: “The Preston curve is 80% about the impact of income on longevity”.

The proposed theory is centered around a novel approach to the modeling of aging. Based on research from the natural sciences, we conceptualize the process of aging as one whereby the human organism gradually loses redundancy and thus becomes more *fragile*. This leads us to a law of motion for human frailty, which depends on physiological parameters and health expenditures. The process of increasing frailty is relentless and accelerating over time, but it may be slowed by health investments. The incentive to slow down aging is a longer life, which facilitates greater consumption and thus utility. The costs of doing so is foregone utility from consumption of goods that do not serve to prolong life. In a setting where individuals are maximizing lifetime utility from consumption we examine the optimal intertemporal choice with respect to savings and health investments. Subject to the physiological constraints faced by humans, and the standard budget constraint, this allows us to characterize the optimal speed of aging and thereby optimal longevity. The aging process (and time of death) is thus endogenously determined by optimal health investments. The physiological foundation of the aging process, as implemented in the analysis, implies that key physiological parameters all have been estimated with great precision in the medical science literature. This is in large part what enables us to calibrate the model in a meaningful way, and quantitatively confront the foundation of the Preston curve.

Our analysis is primarily related to the seminal work of Grossman (1972) on the demand for health and to Ehrlich and Chuma’s (1990) important work on optimal longevity.³ A key theoretical difference between our work and previous contributions lies in the law of motion which governs the aging process; the counterpart to the law of motion for health capital in the existing literature. We elaborate on the broader differences between the standard approach

worth noting that there is an ongoing debate about the impact of life expectancy on growth. See e.g. the interesting exchange between Acemoglu and Johnson (2007), Hazan and Zoabi (2006) and Hazan (2009), who suggest a modest (if any) effect of life expectancy on growth, and Cervalletti and Sunde (2010, 2011) who argue for a positive impact of longevity on growth.

³Other notable contributions in this tradition are Reid (1998), Eisinger (1999) and Foster (2001). Also related are contributions investigating the evolution of life expectancy at the aggregate level; see Cervellati and Sunde (2005) and Galor and Moav (2007).

to health capital accumulation and our proposed approach below. Another difference is that we employ our model to examine the underlying forces that shape the Preston curve from a quantitative angle; this has not been attempted in the existing literature.

A few remarks on the setup are warranted. First, one could imagine that health investments also stimulate productivity, which inevitably influences the income—longevity correlation in the cross-country data. In the present context, however, we are interested in understanding the potential strength of the causal mechanism running from income to longevity; this objective is compromised if we admit the reverse chain of causality to be present in the model, for which reason it is suppressed. Second, in our model health matters for the time horizon, but does not enter the per period utility function. Here we follow Becker (2007) and use a standard utility function for consumption goods implying that any explanation for the health gradient results from the budget constraint. Finally, the accumulation of health deficits and death itself are clearly stochastic events from the individual viewpoint. In order to get some nicely interpretable analytical results we follow the conventional literature of Grossman (1972), Ehrlich and Chuma (1990), and Hall and Jones (2007) and focus on a deterministic framework. At the expense of analytical tractability, uncertain lifetime, as well as utility from state of health, can easily be integrated into the framework developed below without substantial implications for the key results (see Strulik, 2011, 2012).

The paper proceeds as follows. In the next section we derive the key law of motion that characterizes the aging process, and we discuss how this law of motion differs from the standard health accumulation equation. Section 3 contains the main analysis of optimal aging and death. In Section 4 we calibrate the model to US data and examine its predictions as well as its dynamical properties in the presence of shocks to income, technology and more. Then, in Section 5, we apply the model to the study of the Preston curve. Finally, Section 6 is reserved for concluding remarks.

2. INTRODUCING BIOLOGICAL AGING

2.1. Modeling Human Aging as a Process of Deficit Accumulation. Aging is defined as the intrinsic, cumulative, progressive, and deleterious loss of function that eventually culminates in death (Arking, 2006). At the individual level the aging process exhibits great heterogeneity, and is only imperfectly captured by chronological age; some 60 year-olds are as fit as some 40 year-olds and vice versa. Indeed, biologists and gerontologists stress that individual aging should be viewed as an event-dependent process, rather than as a time-dependent process. There is no such thing as a “biological clock”, which determines the speed of individual aging and timing of death.

At the population level, however, age is a better predictor of aging and death. The so-called Gompertz-Makeham law of mortality implies that the death rate rises exponentially with age.⁴ The fact that the age-specific mortality rate increases is obviously a manifestation of the aging process. But then why is the death rate (for the average individual) increasing with age?

A reductionist approach might suggest that the organism ages because of aging organs, which in turn is caused by aging tissue, brought on by the aging of cells and so on. Such an approach, however, is a dead end. The reason is that, eventually, a level of “disaggregation” is reached, which consists of non-aging entities: atoms, for example. The aging process can not be understood in this manner. But how can a system constructed from non-aging components age in the manner suggested by the Gompertz-Makeham law of mortality?

A promising strand of literature in biology has sought an answer by drawing on reliability theory from engineering.⁵ To understand the basic idea, consider the following model (Gavrilov and Gavrilova, 1991). Suppose we view the organism as a whole as consisting of a fixed number of individual parts, which we will refer to as “blocks”. Each block does not age. That is, the failure rate of a block is constant. This assumption captures that the human

⁴More precisely, the Gompertz-Makeham law states that the age specific mortality rate, $\pi(t)$, evolves with age, t , in accordance with the formula $\pi(t) = a + b e^{\alpha t}$ where a, b and α are parameters. The fit of the Gompertz-Makeham law, for the population aged 20 and above (children and teen-ages are “special”) is extremely good, always featuring an R^2 in excess of 0.95. See e.g., Arking (2006).

⁵Reliability theory is used in engineering to understand the failure rate of mechanical devices. Gavrilov and Gavrilova (1991) were the first to introduce reliability theory into biology.

organism is ultimately constructed from non-aging components, as noted above. Next, assume that the blocks are connected in parallel, and that the system as a whole is assumed to survive as long as there is one functioning block remaining. This assumption is motivated by the physiological fact that the human organism is characterized by a great deal of redundancy; as young adults the functional capacity of our organs is estimated to be tenfold higher than needed for mere survival (Fries, 1980). Though each block does not age the passing of time will reduce redundancy in the system as a whole, which leads to an increasing failure rate of the system. Hence, the simple model successfully reproduces an exponentially rising death rate with age. Many extensions of this basic model have been developed, which have led to new insights into the aging process.⁶ But for present purposes the key point of reliability theory is conveyed by the simple model: senescence can be conceptualized as the gradual loss of redundancy, ultimately leading to organism collapse.

Following the underlying reasoning of reliability theory one may therefore think of aging as being characterized by increasing *frailty*. That is, as the redundancy of the human organism shrinks we become more fragile. An empirical measure of human frailty has been developed by Mitnitski and Rockwood and various coauthors in a series of articles (e.g., Mitnitski et al, 2002a,b; 2005; Rockwood and Mitnitski, 2006).

As humans age we develop an increasing number of disorders, which Mitnitski et al. (2002a) refer to as “deficits”. Some of these deficits may be viewed as rather mild nuisances (e.g., reduced vision) while others are more serious in nature (e.g., strokes). Nevertheless, the notion is that when the number of deficits rises the body becomes more frail. A frailty index can then be estimated as the proportion of the total potential deficits that an individual has, at a given age.⁷

⁶For example, if one extends the framework we have just sketched by assuming that the body consists of a number of *essential* blocks (so that if one fails the organism fails), each of which consists of non-aging elements of which some initially are defect, the model reproduces the Gompertz-Makeham law; see e.g., Gavrilov and Gavrilova (1991).

⁷Some methodological notes. To be in the index the deficit must have been demonstrated to be a group and individual indicator of health, and an important correlate of survival (Mitnitski et al., 2002a). “The total number of potential deficits” are in practise determined by the survey at hand. This may seem arbitrary. But according to Rockwood and Mitnitski (2007) the exact choice of deficits is usually not crucial. Provided sufficiently many indicators (40 or more) are present in the index, results tend to be relatively unaffected. Note that the construction of the index does not preclude that its value declines on occasion.

Mitnitski et al. (2002a) show that the following equation fits data on the proportion of deficits, $D(t)$, of the representative individual at age t very well.

$$D(t) = E + Be^{\mu t}.$$

This “law of increasing frailty” explains around 95% of the variation in the data, and its parameters are estimated with great precision. The parameter E turns out to be common for men and women; using a data set encompassing 66,589 Canadians, aged 15 to 79, Mitnitski et al. (2002a) estimate E to 0.02, with a standard error of 0.001. The parameters B and μ are gender specific. For Canadian men (women) $\log(B)$ is -5.77 ± 0.06 (-4.63 ± 0.06), while μ is 0.043 ± 0.001 (0.031 ± 0.001). Interestingly, very similar estimates for B and μ are obtained on data for Australia, USA and Sweden (Rockwood and Mitnitski, 2007). Hence, in these four developed countries (in spite of differences in samples, the precise contents of the frailty index etc.) the average individual accumulates 3-4% more deficits from one birthday to the next.

We can restate the law of increasing frailty in flow form by differentiating with respect to age:

$$\dot{D}(t) = \mu(D(t) - E), \tag{1}$$

where E works to slow down the speed of deficit accumulation. In order to see that the influence of E in (1) is consistent with Mitnitski and Rockwood’s equation for the stock of deficits, integrate (1) and insert the initial condition $D(0) = D_0$ to get the solution $D(t) = (D_0 - E)e^{\mu t} + E = D_0e^{\mu t} - E(e^{\mu t} - 1)$. Since $e^{\mu t} > 1$ for all $t > 0$, a larger autonomous component E implies less deficits for any given age t . Note also that the compound parameter $(D_0 - E)$ corresponds to Mitnitski et al.’s estimate of B . In the natural science literature the parameter E is interpreted as capturing the impact of non-biological factors on deficit accumulation (Mitnitski et al., 2002a). Accordingly, we will assume that E is amendable to change by way of deliberate investment. This is the way in which the individual may attempt to slow down aging in the model below.

Specifically, we propose the following parsimonious refinement of the process of deficit accumulation:

$$\dot{D}(t) = \mu (D(t) - a - Ah(t)^\gamma), \quad (2)$$

where $D(0)$ is given. The parameter a captures environmental influence on aging beyond the control of the individual (less pollution, say, implying a higher value for a), the parameters $A > 0$ and $0 < \gamma < 1$ reflect the state of the health technology, and h is health investment. While A refers to the general power of health expenditure in maintenance and repair of the human body, the parameter γ specifies the degree of decreasing returns of health expenditure. The larger γ the larger the relative productivity of cost-intensive high-technology medicine in maintaining and repairing highly deteriorated human bodies.⁸

By way of contrast to E , the parameter μ – impressed by its empirical constancy across developed countries – is considered to be a physiological parameter. In the remaining we will refer to this physiological parameter as *the force of aging*, as it drives the inherent and inevitable process of human aging. In a science-fiction version of the model we could perhaps also address how health expenditure, for example through epigenetic regulation or hormone replacement therapy, affects μ . In retrospect, however, there is so far very little evidence that “standard” medical treatments have substantially modified the *rate* at which our bodies decay (Gavrilov and Gavrilova, 1991). Standard treatments like, for example a bypass operation or a liver transplant, are effectively delaying death (by removing one or several health deficits) but they are not manipulating the intrinsic rate of deficit accumulation. Here we confine the analysis to standard treatments, which are appropriately captured in $E(t)$.

In order to capture death, we need to invoke an upper boundary to deficit accumulation, \bar{D} . In the analysis below the representative individual remains alive as long as $D(t) < \bar{D}$. Direct evidence on the existence of an upper boundary for D is found in Rockwood and Mitnitski (2006). Observe that equation (2), along with the restriction that $D(t) < \bar{D}$, provides a complete description of aging until death. In this process, chronological age does *not* play

⁸If individuals’ investment in health influences E , as we suggest, then one may wonder if the “frailty law” should still work empirically, as E then is expected to exhibit individual-level variation. It should; but the cross-section estimate for E should be interpreted as the average level in the sample in question (see Zellner, 1969).

a role in itself. While the model developed below concerns optimal aging and death of a representative agent of a cohort, it is nevertheless worth noting that this formulation is in concordance with a central point made by biologists and gerontologists: individual aging is *not* time-dependent. This follows since $\dot{D}(t)$, by (2), is only influenced by current investments and accumulated deficits; t (chronological age) plays no independent role.

This completes the development of the central equation which governs aging. But before we turn to the analysis of optimal aging and death, we briefly compare our law of motion for deficits with the familiar law of motion from health, which dates back to the work of Grossman (1972).

2.2. Deficit Accumulation vs Health Accumulation. Usually, health is introduced as a state variable similar to human capital. In its most basic form, this would involve an equation such as $\dot{H}(t) = I(t) - \delta H(t)$, where $I(t)$ is investment in health and $H(t)$ is the stock of health capital (see e.g. Ehrlich and Chuma, 1990).

If this equation is employed as a description of the aging process it contains a rather unfortunate implication: It predicts that health depreciation is greater when the stock of health is large (δH), which, of course, usually would mean when individuals are relatively young. In reality, the process of aging is a process where the rate of decline in health status accelerates during life; as explained above, both health deficits and the mortality rate rises exponentially with age, implying slow aging early in life and rapid deterioration in latter stages. In practise, therefore, health losses are greater in states where the health index is *low*, which usually means when one reaches a more advanced age.⁹

Naturally, this problem has not gone unnoticed. Since the work of Grossman (1972), the standard “fix” has been to introduce an age-dependent rate of health depreciation, $\delta(t)$, with $\dot{\delta}(t) > 0$. While the rate of change in the rate of depreciation may seem intuitively appealing the question remains how *exactly* the function is to be specified. Insofar as one does not get

⁹Of course, in other applications beyond aging the law of motion may be perfectly reasonable. As a macro representation, for instance, the depreciation term may reflect lost health capital from diseased individuals. Our critique of the equation is confined to its relevance in the context of modeling the aging process

this choice “right” the resulting model will represent an inaccurate description of the aging process, and, by extension, its predictive power will be diminished.¹⁰

Our approach, by way of contrast, employs a physiological foundation for health depreciation. In order to see this advantage clearly we convert the equation for deficit accumulation into one for health accumulation by assuming that actual health is defined as best attainable health minus accumulated frailties: $H = \bar{H} - D$ where \bar{H} is “maximum health”; the state of health of a normal 15 year old, say. This implies $\dot{H} = -\dot{D}$. Inserting equation (1) into \dot{H} and substituting D for $\bar{H} - H$ provides $\dot{H} = -\mu(\bar{H} - H - E)$, or,

$$\dot{H}(t) = \mu E - \mu(\bar{H} - H(t)). \quad (3)$$

Following the approach above, we may associate E with investments. Finally, using the terminology of Grossman (1972) and Ehrlich and Chuma (1990) we may define the minimum level of health below which life is infeasible as $H_{min} \equiv \bar{H} - \bar{D}$, where \bar{D} is maximum deficits. We are thus left with a simple linear differential equation for health, which differs in one crucial respect from the one adopted in the economics literature: Consistent with the facts, the equation predicts that health loss is small at a good state of health and increasing losses are predicted when the health stock deteriorates. Hence, a meaningful description of how health deteriorates with age (as the health index erodes) is implicit in our frailty equation. Notice that according to equation (3) an improvement in health status works to slow down aging; i.e., it stimulates health accumulation.

The advantage of our modeling approach is that various parameters have clear empirical counterparts. This is also the case for the frailty index itself, which has already been developed and tested in the natural sciences. These aspects are very useful in the context of calibration and will be valuable when the model is taken to the data and tested for its economic implications. Obviously, in empirical studies it is commonplace to fit some nonlinear

¹⁰Making $\delta(t)$ an *explicit* function of time also implies that chronological age is supposed to matter to aging *per se*. This is a problematic assumption from a biological perspective, as explained above. This position of biologists is supported by work of health economists. For example, Zweifel et al. (1999) demonstrate that among the elderly health expenditure is not predicted by chronological age once “time remaining until death” is controlled for. This suggests that health status (e.g., frailty), and not the year on the birth certificate, is what matters to health investments.

function $\delta(t)$ to capture health depreciation. We do not deny that this is possible, and in specific applications very useful. But for present purposes it is an undeniable advantage that the mechanism of health deficit accumulation, the force of aging μ , is pinned down by an invariant parameter, which is explained by gerontological theory and estimated from actual data with great precision. This makes us confident that the model cannot only be calibrated to some particular sample, but that it can also be used for out-of-sample predictions. That is, for an attempt to explain the position and slope of the Preston curve.

3. A THEORY OF OPTIMAL AGING AND DEATH

3.1. The optimization problem. Consider an adult maximizing utility from consumption $c(t)$ over his or her life. We are considering a representative member of a cohort, for which reason the maximization problem can be viewed as deterministic. The initial age is for convenience normalized to zero. Longevity T is finite and endogenous. Let $\rho \geq 0$ denote the rate of pure time preference; the rate of time preference need not be strictly positive in order for the problem to be well-defined. Summarizing, intertemporal utility is given by

$$\int_0^T e^{-\rho t} u(c(t)) dt \tag{4}$$

with $u(c) = (c^{1-\sigma} - 1)/(1 - \sigma) + b$ for $\sigma \neq 1$ and $u(c) = \log(c) + b$ for $\sigma = 1$. The case of $\sigma > 1$ might lead to perverse results. If, additionally, consumption is below unity such that $c^{1-\sigma} > 1$, flow utility $u(t)$ becomes negative and the individual may prefer a shorter life. In order to exclude this degenerate case we introduce, inspired by Hall and Jones (2007), a constant $b \geq 0$ which is assumed to be large enough to guarantee positive utility. For our numerical results discussed below, it turns out that c is never below unity which allows us to assume $b = 0$. In this case flow utility converges smoothly to $\log(c)$ for $\sigma \rightarrow 1$. In the calibration, moreover, it turns out that $b = 0$ generates the best fit for the value of life.

It is perhaps tempting to allow frailty to enter the per period utility function as well. Nevertheless, as will become clear, the simple model matches the cross-country income gradient

remarkably well. In only allowing health investments to affect utility via longevity, our approach is similar to that of Becker (2007). At the same time, the present framework can be extended to allow utility from state of health as well as uncertain lifetime without substantive implications for the result (Strulik, 2011, 2012).

The individual receives a wage income w . We assume the wage rate is constant during life; a simplifying assumption which serves to highlight the central workings of the model. Income can be spend on consumption goods c or on health goods h . The relative price of health goods is p . While consumption goods are directly utility enhancing, health goods (e.g. a hip replacement, a weekend at the spa) are instrumental in repairing or delaying bodily decay and, ultimately, in prolonging the period during which consumption goods can be enjoyed. Accordingly, the only value of “good health” in this model is via its impact on longevity.¹¹ Besides spending income on final goods, the individual may invest in capital k and receive a net interest rate r . The individual takes all prices as given. The law of motion for individual wealth is thus given by (5).

$$\dot{k}(t) = w + rk(t) - c(t) - ph(t). \quad (5)$$

The individual is assumed to inherit wealth $k(0) = k_0$, and to leave a bequest $k(T) = \bar{k}$ (which both could be zero).

The problem is to maximize (4) subject to the accumulation equations (2) and (5), the initial conditions $D(0) = D_0$, $k(0) = k_0$, and the terminal conditions $k(T) = \bar{k}$, $D(T) = \bar{D}$. The problem can be solved by employing optimal control theory; the state variables are $k(t)$ and $D(t)$ and the control variables are consumption $c(t)$ and health investments $h(t)$. From now on time indices are suppressed in the interest of brevity.

¹¹At first sight it may not be obvious that a hip replacement delays death. Bypass surgery or organ transplants are probably more immediately intuitive examples for life-extending health expenditure. But gerontological theory as well as analysis of the frailty index suggest that health deficit accumulation is Markovian: the probability to get yet another health deficits next period depends positively on the number of already accumulated deficits (Mitnitski et al., 2006). Together with the evidence for the phenomenon that death can be expected at a given value of the frailty index, it is clear that any untreated health deficit is conducive to bodily decay and death.

3.2. **Optimal Aging.** From the first order conditions we obtain the Euler equations

$$g_c \equiv \frac{\dot{c}}{c} = \frac{r - \rho}{\sigma} \quad (6)$$

$$g_h \equiv \frac{\dot{h}}{h} = \frac{r - \mu}{1 - \gamma}. \quad (7)$$

While equation (6) is the standard Consumption Euler equation, equation (7) provides a novel Euler equation for health expenditures: the “Health Euler equation”. As in the context of non-health expenditure a higher intertemporal marginal rate of transformation (i.e., a higher real rate of interest) calls for growing health expenditures. At the same time, growth in health expenditures is tempered by the force of aging, μ . The intuition is that if μ is high, deficits will accumulate very fast at the end of life, making late-in-life health investments a relatively ineffective way of prolonging life. Instead, the optimal strategy is to invest more heavily early in life. Conversely, if the force of aging is low (i.e., μ is “small”) late-in-life health expenses are more effective in prolonging life, for which reason it can be optimal to allow h to grow over time. Finally, growth of health expenditure is also influenced by γ , which captures the curvature of the health investment function: a larger γ increases growth in health expenditures, *ceteris paribus*. Intuitively, if γ is “small”, diminishing returns set in rapidly, which makes it optimal to smooth health expenditure to make the deficit-reducing effect as large as possible. Thus, in this setting where health does not enter the per period utility function it is the state of technology that determines the extent of expenditure smoothing across life.

It is interesting to note that it is not *necessarily* optimal to plan for increasing health expenditures during life. In societies where the force of mortality is sufficiently strong it may be the optimal strategy to prioritize early-in-life health spending, and thus allow health spending to *decline* over time. One way to think about this result is to associate early-in-life investments in health with “preventive measures” and late-in-life health expenses with “treatment measures”. If so, then the Health Euler simply says that in societies where individuals are aging at a rapid pace ($r < \mu$) it is optimal to focus resources on prevention, rather than on treatment.

Even if it is optimal to allow health expenditure to rise over time, the expenditure *share* for health, $\epsilon_h \equiv h/(h + c)$, may nevertheless fall, if pure consumption is growing sufficiently rapidly. Hence, the expenditure share for health is increasing over time if $g_h > g_c$, or (using (6) and (7)) if

$$r - \frac{1 - \gamma}{\sigma} \cdot (r - \rho) > \mu, \quad (8)$$

If the ratio $\frac{1-\gamma}{\sigma}$ increases it becomes less likely that the condition is met. The interpretation is that a small σ leaves little incentive to smooth consumption, which implies (*ceteris paribus*) faster growth in consumption. Meanwhile, γ parameterizes the incentive to smooth health investments; the smaller γ is the greater the incentive to smooth health investments. Consequently, when $\frac{1-\gamma}{\sigma}$ increases it implies a greater desire, on the part of the consumer, to smooth health investments relative to consumption, which suggests a declining share of health during the life cycle as long as $r > \rho$.

A higher rate of interest will increase the rate of growth of both health investments and consumption. But whether a higher r makes it more likely that the condition is fulfilled depends on the incentive to smooth consumption and health investments, respectively. If consumers have a greater desire to smooth health investments ($\frac{1-\gamma}{\sigma} > 1$) a higher r reduces the health expenditure share.

This leaves the impact of ρ and μ . Health spending is independent of ρ since h prolongs life, but does not affect per period utility, while a higher ρ (as usual) works to slow down consumption growth.¹² Hence, it becomes more likely that the health share is rising if individuals are highly impatient. At the same time, in the presence of a greater force of aging people will act more impatiently vis-a-vis health investments, inducing them to invest early in life. If the force of aging is sufficiently strong, the expenditure share for health will therefore be declining over the life cycle.

The bottom line is that the life cycle path for the health share is ambiguous. In particular, it is not a given that the expenditure share should be rising with age. In theory one may

¹²Appendix A offers a more detailed discussion of why g_h is independent of ρ .

therefore expect differences in the health share across countries that differ in terms of culture, technology and physiology.

3.3. Optimal Death. In contrast to the pioneering contributions on the topic of optimal health investments (Grossman, 1972; Ehrlich and Chuma, 1990), the present model is sufficiently simple to allow for an explicit solution of the involved differential equations, from which we can then infer optimal longevity T . In order to see this, first note that the boundary value problem with variable terminal value T requires that the boundary conditions $D(0) = D_0$, $k(0) = k_0$, $k(T) = \bar{k}$, $D(T) = \bar{D}$ and $h(T) = 0$ are fulfilled. Then integrate (5) in order to solve for $k(T)$ and integrate (2) in order to solve for $D(T)$. Finally solve the associated Hamiltonian for $H(T) = 0$. This provides the following set of equations (see Appendix A for details):

$$\bar{D} = D_0 \exp(\mu T) - \frac{\mu A h(0)^\gamma \exp(\mu T)}{g_D} [\exp(g_D T) - 1] - a [\exp(\mu T) - 1], \quad (9a)$$

$$\bar{k} \exp(-rT) = k_0 - \frac{w}{r} [\exp((r)T) - 1] - \frac{c(0)}{g_c - r} [\exp((g_c - r)T) - 1] - \frac{p h(0)}{g_D} [\exp(g_D T) - 1] \quad (9b)$$

$$0 = u_T - \frac{\exp(-\sigma g_c T)}{c(0)^\sigma} \times \quad (9c)$$

$$\left\{ \frac{(\bar{D} - a)}{\gamma A} p h(0)^{1-\gamma} \exp((1 - \gamma)g_h T) - \frac{1 - \gamma}{\gamma} p h(0) \exp(g_h T) - w - r\bar{k} + c(0) \exp(g_c T) \right\}$$

where $g_D \equiv (\gamma r - \mu)/(1 - \gamma)$ and $u_T \equiv \log(c(0)) + g_c T + b$ in the case of log-utility and $u_T \equiv [c(0) \exp(g_c T) - 1]^{1-\sigma}/(1 - \sigma) + b$ otherwise. These three equations can be solved for the three unknowns: $c(0)$, $h(0)$, and T . Having found the optimal initial values and the optimal terminal time, the four-dimensional dynamic system (2) and (5) – (7) is fully specified and it can be solved for the optimal life-cycle trajectories of c , h , k and D .

4. COMPARATIVE DYNAMICS

4.1. Calibration and Model Predictions for the US. Before we turn to the experiments, we calibrate the model to USA data, though we have to rely on data from Canada when it

comes to pinning down the force of aging. Accordingly, we take from Mitnitski et al. (2002a) the estimate of $\mu = 0.043$ for Canadian men.¹³

In order to calibrate γ we turn to data on growth in health expenditures over the life cycle. The model predicts that countries that are technologically similar, and are inhabited by genetically similar populations should exhibit similar investment patterns across the life cycle. Is there any evidence to support this prediction? Following, for example, Hall and Jones (2007) and Fonseca et al. (2009), we may try to gauge life-cycle developments by examining per capita spending on health across age groups. Figure 1 illustrates such data for four “Western Offshoots”: Australia, Canada, New Zealand and United States.

The immediately visually arresting theme is that the four trajectories appear to be more or less parallel, suggesting similar per capita spending growth. Indeed, if one calculates the slope of the trend, the average annual growth rates across the life cycle are nearly identical: 2.0% (Australia), 2.1% (Canada), 1.9% (New Zealand) and 2.0 % (USA). Hence, while the four countries differ in terms of the level of health spending per capita, the rise in health spending per capita across life is nevertheless strikingly similar. These patterns are consistent with the predictions of the health euler derived above.¹⁴


Next we turn to the Health Euler (7) in order to calibrate the curvature of the health production function, i.e. γ . We put $r = 0.06$ (e.g., Barro et al., 1995) and using data from Figure 1 we put $g_h = 0.021$, to capture the growth of health spending across the life cycle in Canada from which we obtained the value for μ . This produces $\gamma = 1 - (r - \mu)/g_h = 0.19$, which squares well with the independent estimates obtained by Hall and Jones (2007).¹⁵

¹³As explained in Section 2: the force of aging within the US and Canada are similar (Rockwood and Mitnitski, 2007). Thus, using the estimate from the Canadian sample should be a good approximation. While Rockwood and Mitnitski (2007) stress the similarity of their results for US and Canadian populations they do not report the detailed results for their US analysis, for which reason we are forced to rely on the results from the Canadian sample.

¹⁴That the countries otherwise seem to behave differently in the context of health is nicely illustrated by their aggregate health shares. In 2006 the share of total health spending in GDP was 8.7% in Australia, 10.0% in Canada, 9.4% in New Zealand and 15.3 % in the US. Notice that the observation of an approximately constant growth of health expenditure across ages in a particular period of observation (around the year 2000 in our case) is compatible with the fact that life cycle health expenditure changed over time (as observed for the US by Meara et al., 2004). In our theory this would be attributable to change in μ , ρ or, presumably most likely, in γ , the marginal productivity of health expenditure.

¹⁵Hall and Jones allow the curvature of the health production function to be age dependent. The average value is close to 0.2.

Figure 1: Health expenditures per capita across age groups: Selected countries 1999-2001



The figure shows log health expenditures per capita by age-group in Australia, Canada, New Zealand and the USA. *Sources:* US from Keehan, Lazenby, Zezza, and Catlin (2004); Canada from Health Canada (2001); Australia from Australian Institute of Health and Welfare (2004); for New Zealand from New Zealand Ministry of Health (2004). *Notes:* (i) In order to consolidate the age intervals for the purpose of illustration simple averages across age intervals and gender have been invoked. (ii) For the USA the first age group concerns individuals aged 19-44. (iii) The year of data collection varies slightly, as indicated by the legend to the figure: The data for the US is from 1999; data for Australia and Canada is from 2000 and data for New Zealand concerns 2001.

Mitnitski et al’s regressions do not involve children. Much like the Gompertz-Makeham law for mortality, individuals below roughly the age of 20 are presumably not well described by the law of frailty (in stark contrast to the group above 20). Hence, when calibrating the model we assume individuals are “born” at the age of 20.

With this in mind we do the following in order to parameterize the deficit accumulation equation. From Mitnitski et al.’s (2002a) regression analysis we can back out $D(0) = D(20) = 0.0274$ as the relevant initial value for a 20 year old and $\bar{D} = 0.1005$ 55.6 years later; the life-expectancy of a 20 year old US American in the year 2000 was 75.6. In order to identify a ,

we assume that before 1900 the role of technology in the repair of health deficits of adults was virtually zero. In 1900 the life expectancy of a 20 year old U.S. American was 42 years (Fries, 1980). From countries for which extended historical time series are available we know that life expectancy for adults was about the same in the 18th, 19th, and early 20th century (See England and Sweden at www.mortality.org). Setting $A = 0$ and thus shutting down the health investment channel, we obtain $a = 0.013$ such that the model predicts a life-expectancy at age 20 of 42 years.


Since consumption tends to be essentially constant across the life cycle, once family size has been taken into account (Browning and Ejrnaes, 2009), we put $\rho = r$. The most natural specification of the intertemporal elasticity of substitution is probably unity (Chetty, 2006). So we set utility to be logarithmic for the benchmark case. In order to focus on health expenditure as a motive for savings we assume $k(0) = k(T) = 0$ for the baseline simulation. Hall and Jones (2007) emphasize that the presence of autonomous utility $b > 0$ is essential in their theory to produce the result that the health expenditure share is increasing with age. By way of contrast, the present theory does not rely on $b > 0$, the model predicts an empirically plausible association of health expenditure and age, through a different, novel channel: the law of increasing frailty. To emphasize this point we set $b = 0$ in our benchmark calibration. Moreover, the case of $b = 0$ provides the best fit of the model to independent estimates of the value of life, as detailed below.

Finally, in order to identify the role of health technology in preventing death, we take GDP per worker in the US in the year 2000 (PPP\$ 77,003) and assume a capital share of $1/3$, which implies an annual labor income (in international dollars) of \$ 46,669. We normalize the relative price of health in the year 2000 to unity, and estimate A such that the individual dies with deficits \bar{D} at age 75.6, which was the life-expectancy of a 20 year-old US American male in the year 2000. This provides the estimate $A = 0.0014$. Note that this identification leaves room for a very broad interpretation of health promoting technology. It does not only include advancements in medical science but also advancements in general knowledge about nutrition and hygiene, i.e. all knowledge that makes health expenditure more effective.

In sum, the model is calibrated to match initial deficits, end-of-life deficits, and longevity in 2000; $D(0)$, \bar{D} and T , respectively. It also matches the growth of health spending across age groups exactly. Yet the *path* of deficits (between age 20 and death), as well as the evolution of expenditure *shares* across the life cycle, are left unrestricted. This provides the opportunity for an informative consistency check of our approach: does the path of the health expenditure share and cohort frailty match the data?

Figure 2 shows the basic run of the model; including the association between model predictions and actual data. In the lower right panel of Figure 2 stars indicate actual health expenditure shares by age-group inferred from Mazzocco and Szemely (2010), which can be compared to the solid line representing the prediction from the calibrated model. Admittedly, our calibration predicts a too flat slope of the age trajectory for the health expenditure share. We conjecture this failure is mainly caused by the fact that our model does not allow wages to decline with age. Nevertheless, the model does well on average.

FIGURE 2: OPTIMAL AGING: BASIC RUN



Solid lines: basic run. $T = 55.6$. Parameters: $a = 0.013$, $A = 0.0014$, $\mu = 0.043$, $\rho = r = 0.06$, $D_0 = 0.0274$, $\bar{D} = 0.10$, $p = 1$. Dashed line: no health expenditure ($A = 0$) stars: data, red dot: US mean male age and age-structure-weighted expenditure share.

The upper left panel shows the accumulation of deficits for the average US (male) citizen. The dotted line reflects the scenario where there are no health investments occurring; the solid line involves optimal health investments as predicted by the model. In the figure we also illustrate the law of frailty, as estimated by Mitnitski et al. (2002a) (represented by stars). As can be seen, the model’s fit is rather good in that the path of deficits is fairly close to the one found in the data.

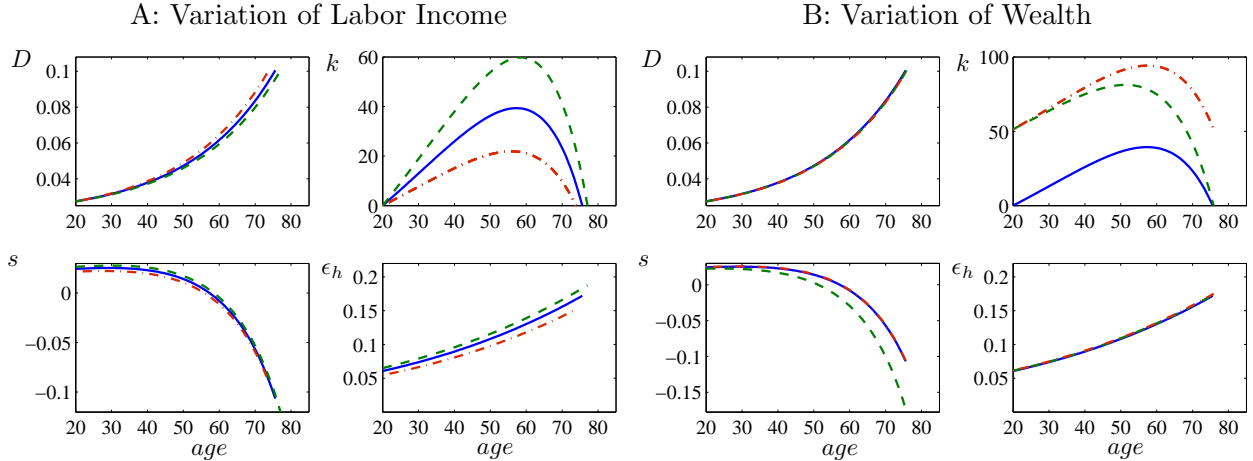
As another check, we can employ the model to calculate the value of life (VOL) of a US citizen at different age intervals, so as to compare the present frameworks predictions with previous estimates. The VOL provides a monetary expression of aggregate utility experienced during life until its end, that is period utility is converted by the unit value of an “util”, $u'(c)$. Applying the formula $\tilde{V}(t) = \int_t^T e^{-\rho(\tau-t)} u(c)/u'(c) d\tau$ we obtain the VOL at age t . Our benchmark calibration predicts a VOL of about \$ 7.9 million for a 30 year-old, \$ 7.1 million for a 45 year-old, \$ 5.1 million for a 60 year old. In order of magnitude these values correspond well with Murphy and Topel’s estimate of \$ 6.3 million for the value of a statistical life. Observe that our estimates for VOL are obtained under the assumption $b = 0$. Clearly, there are no alternative values of $b \geq 0$ that will allow us to match existing VOL estimates better, since the predicted VOL is already at the upper end of the existing estimates. This suggests $b = 0$ is a sensible benchmark.¹⁶

Finally, the model also holds predictions about the path of the savings rate (s) and wealth across the life cycle. Contingent on the imposed parameter values health expenditures are rising during life, as we just saw, which requires early-in-life savings so as to finance late-in-life investments. The wealth of the representative individual therefore follows a “hump shaped” trajectory; as seen from the upper right hand side window, supported by positive savings early in life and dissaving late in life (lower left hand side window). These patterns are qualitatively consistent with the standard life cycle theory of consumption. We have not examined whether the model matches wealth and savings quantitatively, as it seems unlikely

¹⁶In the Appendix we nevertheless discuss an alternative specification with positive b (with and without adjustments for σ).

that we can predict wealth during working years without a careful discussion bequests and inheritances.¹⁷

FIGURE 3: HEALTH AND WEALTH OVER THE LIFE CYCLE



Green (dashed): wage income increases by 1/3. $\Delta T = 1.48$, implied elasticity 0.08). Red (dotted): wage income decreases by 1/3 ($\Delta T = -1.79$, implied elasticity -0.10).

Green (dashed lines): $k(0) = w$, $\bar{k} = 0$ ($\Delta T = 0.27$). Red (dotted): $k(0) = w$ and $\bar{k} = w$ ($\Delta T = 0.29$).

4.2. Experiment 1: Income. Panel A of Figure 3 shows how the agent reacts if his income, w , is perturbed. The green (dashed) line is associated with an increase in w of 1/3, the red (dashed-dotted) line depicts the reaction to a reduction of w by 1/3 (in all the experiments below, “green” is associated with increases, and “red” with reductions in the parameter of interest). As can be seen from the figure, the consequence of higher income is an increase in longevity, peak wealth and the share of health spending. In regard to the latter, note that we are keeping r, ρ, μ, γ constant. Hence the increase in the health share solely reflects a “level effect” through $h(0)/c(0)$. The intuition for why income changes entail a larger change in health spending than regular consumption is that the incentive to smooth the latter is stronger due to diminishing - per period - marginal utility. Higher income therefore leads to a larger adjustment in $h(0)$ than $c(0)$.

The issue of main interest, however, is the quantitative impact on longevity. As seen from the top left hand side corner of panel A, the impact is modest though not inconsequential.

¹⁷The relative importance of life cycle motive vs. the bequest motives for saving is still in debate (e.g., Dynan et al., 2002).

An increase of income of $1/3$ (achievable in a generation with a growth rate of about 1% per year) translates into an increase in longevity of 1.5 years; the reduction involves a fall in life span of 1.8 years. If we convert the impact into an elasticity - the elasticity of longevity with respect to income - we find it to be around 0.09. These effects are not outlandish in comparison to econometric estimates.¹⁸

4.3. Experiment 2: Wealth and Bequest. In the baseline model we did not allow for wealth transfers. Hence, an interesting question is how longevity changes if the individual receives an inheritance, and is forced to pass on a bequest. The dashed line in Figure 3, panel B, examines the impact from offering the individual a transfer comparable to her annual wage, while she is not obliged to pass on any bequest. The dotted line then examines the effect from forcing the individual to pass on a bequest equal to the initial transfer.

When agents receive an inheritance they are naturally able to invest more in health than otherwise, as it increases their life time income. It is interesting to observe, however, that what really matters is the initial inheritance; if bequests are passed on the simulated increase in longevity is about a quarter of a year (0.27 years), whereas the increase is only marginally higher without the bequest requirement. Why is the bequest requirement so relatively inconsequential? The intuition is that what dominates over a lifetime is the accumulated interest on the initial bequest. A 6% annual interest will double the value of the initial inheritance almost every decade. In this light it is perhaps unsurprising that the paying of a one year's wage in bequest has a relatively minor impact on the results.

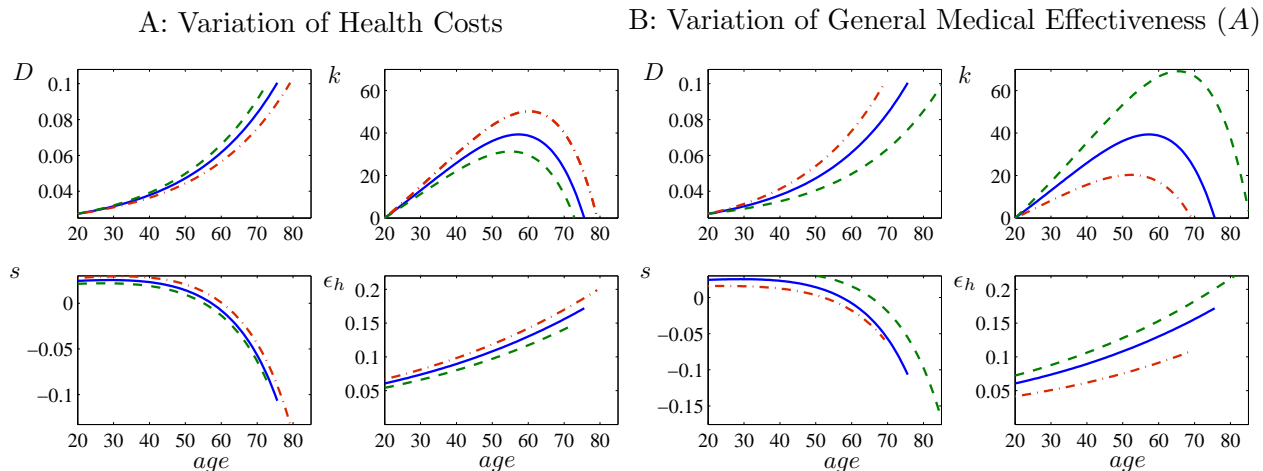
4.4. Experiment 3: Health costs. Our next experiment concerns health costs; the relative price of h . Here we consider a doubling of the relative price of health goods. Rising relative health prices is a realistic scenario; the CPI of medical care has risen faster than the GDP deflator in the US (Cutler et al., 1998).¹⁹

¹⁸See, for example, Pritchett and Summers, 1996. Since the dynamic system is non-linear, the results from our experiments are unlikely to be symmetrical. Life-expectancy can be read off the figure as the age coordinate of the right endpoint of the lines. When $k(T)=0$, which is always the case but in Fig. 3B, life-expectancy can be read off most conveniently where the k curve hits the age axis.

¹⁹However, according to Cutler et al., productivity of health care has also risen. Indeed, the authors argue that relative medical prices, appropriately quality adjusted, might have been declining over time. Below we examine the influence from changes in health productivity.

As is clear from Panel A of Figure 5, when the relative price of health increases, individuals substitute towards regular consumption. As a result, the health share declines. With less health investments, savings decline as well. The end result of a doubling of the relative price for longevity is a reduction by 2.5 years. This amounts to a longevity-price elasticity of 0.09.

FIGURE 5: HEALTH AND WEALTH OVER THE LIFE CYCLE



Green (dashed): doubling p , $\Delta T = -2.7$, implied elasticity -0.09. Red (dotted): halving p ($\Delta T = 3.6$, implied elasticity 0.13).

Green (dashed): 33 percent increase of A ($\Delta T = 9.8$, implied elasticity 0.53). Red (dotted): 33 percent decrease of A ($\Delta T = -6.4$, implied elasticity -0.35).

4.5. Experiment 4: Health Technology. In the experiments above, the impact from the parameter of interest were indirect. For instance, an increase in income translates into both higher health spending and higher consumption. Medical technologies (or the productivity of health investments more broadly), however, have a direct impact on the evolution of deficits and therefore on longevity. A larger impact is therefore to be expected.

In Figure 5, panel B, we depict the impact of health technologies on longevity. We examine the impact from increasing A by 1/3; an increase of a similar magnitude to that which we analyzed in Section 4.2, in terms of income. If health productivity rises by 33% the consequence is an increase in longevity of nearly a decade. The implied elasticity is about 1/2. This is a very large effect, suggesting that the impact from improvements in health productivity easily may have towered that of rising income.

5. THE PRESTON CURVE

In his classic paper Preston (1975) documented that there is a positive and concave cross-country relationship between GDP per capita and life expectancy at birth. Subsequent research has found that the Preston curve still holds (e.g., Deaton, 2003) and that the relationship between income and longevity seems to have emerged sometime late in the 19th century (Bloom and Canning, 2001).

While the Preston curve thus continues to hold, questions remain as for their interpretation. Preston (1975) himself argued that an impact from income on longevity could occur in various ways, including via diet, access to clean water and sanitation, as well as medical treatment. At the same time, it is hard to exclude that the same link could be attained with causality running from life expectancy to income. Or, perhaps income is simply correlated with something else that matters to longevity, like human capital (Bloom and Canning, 2007).

Our model establishes a causal link between income (in the form of wages) and life expectancy (in the form of longevity, T). The mediating factor is health investments, h , which would include basic investments like access to clean water and sanitation, as well as more sophisticated investments associated with medical treatment. But are these mechanisms sufficient to create the empirical income gradient?

In order to confront this question we begin by estimating the Preston curve for 2000. The fitted values from this regression constitute the income gradient to be explained. Naturally, the results from the regression most likely confounds three separate channels: The impact of income on longevity; the impact of longevity on income (reverse causality), and the influence from factors correlated with both income and longevity (omitted variables bias). Our model, however, allows us to isolate the pure effect of income on longevity. Hence, as a second step, we feed the income levels of the countries in our sample through our model, *keeping the parameters we calibrated for the US fixed*, and record the implied levels of life expectancy. As a third step we then compare the longevity predictions of our model to those associated

with the observed income gradient; this comparison allows us to quantify how important the income-to-longevity channel is in accounting for the Preston curve.²⁰

The prediction of the model should not be expected to coincide with the Preston Curve since it is hard to believe that the income-to-longevity channel is the *only* channel captured by the Preston curve. But the deviation between points on the Preston Curve and the prediction of our model will have a sensible interpretation: they reflect either the impact from factors that shift the Preston curve and are correlated with income, or, they reflect the opposite chain of causality (i.e., longevity influencing income). In theory, candidate omitted factors that are correlated with income empirically are captured by the relative price of health goods (p) and health efficiency (A); both p and A will shift the Preston curve if perturbed.²¹ Our analysis does not allow us to pin down the source of the deviation between the model’s prediction and points on the Preston curve (omitted variables vs. reverse causality). But it does allow us to gauge how important they are in accounting for the Preston curve in its entirety: Their influence can be quantified as the fraction of the variation nested in the observed income gradient that our model *cannot* account for.

5.1. Estimating the Preston Curve. While the original Preston curve concerns GDP per capita and life expectancy at birth, we obviously need to modify the “input data” slightly. Our model does not involve children; life expectancy at birth is therefore not the optimal empirical counterpart to T . Instead, we collected data on male life expectancy at the age of 20, to retain consistency with our calibration for the US.²² Unfortunately, life expectancy at age 20 is not recorded consistently each year in all countries; so our actual data involves observations circa 2000. Our income measure is only slightly different from the one employed by Preston (1975) in that we opt for GDP per worker rather than GDP per capita. GDP

²⁰Naturally the actions of private agents and governments may not be as optimal as our model suggests. Deviations from optimality will be reflected in deviations between the model and the observe Preston curve. That is, suboptimal behavior will appear in the left-over residual.

²¹The list of factors that may influence “ A ” is long but probably includes, among others, technology (if health technology does not diffuse immediately), health care institutions (how efficient is the health care system in translating spending into results?), climate (since it influences disease gradients and thereby longevity for income given) and human capital (if education means that a dollar’s worth of health spending is used more effectively).

²²All data sources are found in Appendix C.

per worker is arguably a more appropriate empirical proxy for wages w , the natural income measure in our model.

The country sample was restricted in two ways (beyond via availability of the basic data): following Preston (1975) we ignore countries with a population below 2 million citizens, and in addition we omitted the OPEC countries; GDP per worker is probably a poor guide to average wages in these countries. This leaves us with a sample of 65 countries in 2000.

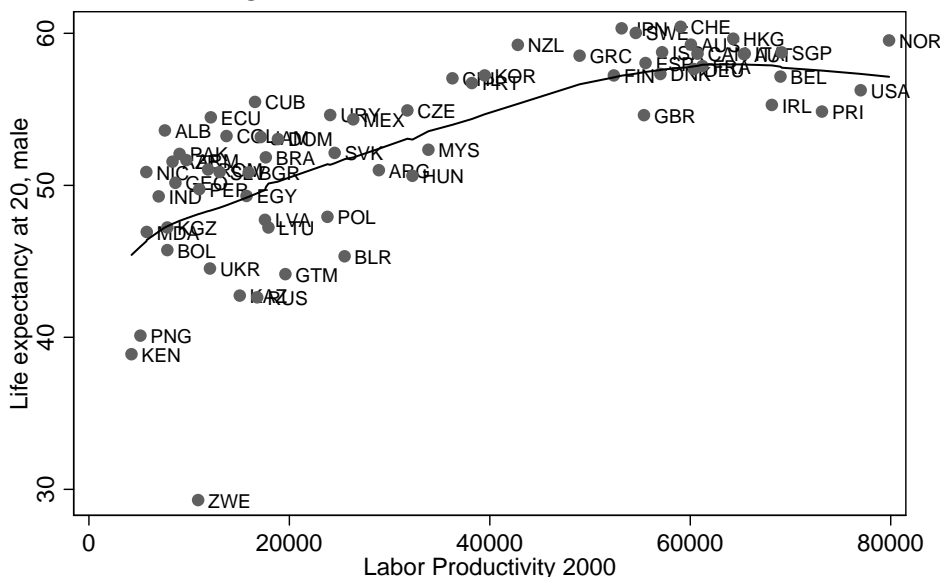
With this data in place we proceed by estimating the income gradient semi-parametrically. That is, we make no a priori functional form assumptions about the income gradient, but add a linear control for the year of observation for longevity. Specifically, we estimate the equation $y_i = f(z_i) + x_i\beta + \epsilon_i$, where the dependent variable is life expectancy at the age of 20, z_i is GDP per worker, x_i is the observation year for life expectancy which adjusts for the fact that life expectancy is not measured in *exactly* the same year, and ϵ_i is a noise term. The function f is assumed to be a smooth, single-valued function with a bounded first derivative, but data decides the exact relation between GDP per worker and life expectancy.²³ The result is depicted in Figure 6.

In comparison to the original Preston curve the estimated function is somewhat more linear to behold. This result is likely caused in part by the fact that we are ignoring child mortality, and in part by the fact that data is missing in many of the poorest countries. Nevertheless, we do see a flattening of the income gradient at high levels of GDP per worker; the income gradient is non-linear and concave. The estimated income gradient, depicted in Figure 6, is the target we compare our model's predictions to.

5.2. Understanding the Preston Curve. In order to examine the importance of the pure income channel, in accounting for the Preston Curve, we feed the income levels (PPP GDP per worker in 2000) for the 65 countries through the model, keeping the parameters we calibrated to US data fixed. Figure 7 allows for a visual comparison between the model's predictions

²³See Lokshin (2006) for details on the estimation algorithm that we employ.

Figure 6: The Modified Preston Curve



The Figure shows the cross-country link between life expectancy at age 20 and GDP per worker in 2000 (65 countries). *Notes:* The line is estimated semi-parametrically, with year of data collection for life expectancy being the linear control. Labor productivity is significant (p-value of 0.000). See Lokshin (2006) for details on the estimation algorithm. See Appendix C for data sources.

regarding life expectancy in 2000, and the empirically estimated income gradient from the cross section in 2000, whereas Table 1 provides some summary statistics.²⁴

TABLE 1: SUMMARY STATISTICS, MODEL VS. DATA


	Data	Model
LE(min)	45.4	47.4
LE(max)	58.0	55.8
LE(mean)	52.4	51.6
LE(median)	51.5	51.2
R^2		0.78

The R^2 is the coefficient of determination: the estimated income gradient constitute data, whereas the model’s prediction constitutes “fitted values”. The reported number thus reflects the fraction of the variation in life expectancy along the income gradient that the model can account for.

As can be seen from the figure, the model (solid line) does a pretty good job at matching the Preston curve (dotted curve). As seen from the table, the model underestimates the range of life expectancies spanned by the Preston curve somewhat, but at average or median income

²⁴Note that the empirical Preston curve – in contrast to the predicted Preston curve – does not go through the US data point (the empirical curve overestimates US health given wealth). Consequently, the two curves do not intersect at the US data point.

FIGURE 7: THE MODEL VS. THE PRESTON CURVE



The figure compares the empirically estimated Preston curve (dotted) to the Preston curve predicted by the model (solid).

the actual income gradient and the one predicted by the model essentially coincide. The table also provide a summary measure of the goodness-of-fit of the calibrated model in the form of the coefficient of determination: the R^2 of the income-to-longevity link comes to 0.78, which means that 78% of the variation in life expectancy along the Preston curve is accounted for by our model.

While the bulk of the variation in the observed link between life expectancy and income thus is accounted for, there are systematic deviations between the Preston curve and predicted life expectancy, which accounts for the remaining 20% of the variation. There are two complementary explanations for the systematic nature of the deviations; the first is based on the assumption that “omitted variable bias” is influencing the empirically estimated income gradient, whereas the second would pertain to the case where reverse causality is thought to be important.

Consider the latter case first. If reverse causality is an issue one may think of the estimated income gradient as the outcome from the interaction of two separate underlying schedules: an income-to-longevity schedule and a longevity-to-income schedule. The former is in theory captured by our model, which by construction does not admit the reverse line of causality. Now, if the longevity-to-income schedule has a *steeper* slope than the income-to-longevity

schedule in the (income, life expectancy) space, then the estimated Preston curve will feature a slope that is strictly larger than the income-to-longevity schedule. As a result, our model (capturing only the income-to-longevity mechanism) should overestimate life expectancy at the bottom of the income distribution and underestimate it at the top.

Alternatively, suppose reverse causality is not an issue. If so, then we would interpret deviations between the Preston curve and the model's prediction as the result of omitted variables; factors that are correlated with both life expectancy and income. Theoretically, such factors could map into A (efficiency of health investments), p (relative price of health goods), or both. Our results then suggest that the price level of health in efficiency units (p/A), is higher in most of the poorest countries relative to US. But by the same token p/A must then be larger in the US compared to many of the richest countries, suggesting that the US health care system (at least in 2000) was less efficient than that of many other rich nations. Whether this is true or not is hard to say. But it remains an observable fact that the US constitutes an "outlier" in health expenditures, but not in terms of life expectancy. In any case, *if* omitted variables is the only channel affecting the Preston curve beyond the income-to-longevity channel, the observed difference between the Preston curve and the model's prediction would have to mean that the price of health in efficiency units is higher in many of the poorest places, yet lower in the richest places, relative to its level in the US.

In practise, of course, we have no way of knowing which of these two "stories" is more important in accounting for the left-over residual. What we do know is that they both influence the Preston curve, and that they together account for some 20% of the variation along the income gradient.

In sum, the analysis suggests that the Preston curve largely, but not exclusively, is due to the causal influence from income on longevity; 78% of the variation in life expectancy along the income gradient is accounted for by our model. This insight leads to an important conclusion. People in the poorest countries are dying earlier than citizen's in rich nations. Our analysis suggests that, to a first approximation, the sad reality seems to be that poor people spend less on health because they are poor and live shorter lives because of it. These conclusions

mimic Preston's (1975) own conjecture regarding the underlying forces that generate his curve rather well. That is, the nonlinear link between income and life expectancy is to a large extent caused by lower health investments in several dimensions, ranging from clean water to medical treatment. Changes in relative prices and health technologies do matter. But in order to understand the income gradient, they do not seem to be the main culprits.

6. CONCLUDING REMARKS

In the present paper we have proposed a theory of optimal aging and death. Individuals maximize lifetime utility subject to the usual budget constraint, but also taking their physiology into account. The physiological constraint concerns the gradual emergence of health deficits that constitute the aging process. While aging and death are inevitable, individuals can invest in their health which serves to slow down aging and prolongs life. Contingent on preferred health investments the aging process as well as the time of death are determined.

The model holds strong predictions regarding optimal health spending across the life cycle, as well as for the optimal evolution of health expenditure shares. Interestingly, it is not necessarily optimal for health spending, or its expenditure share, to rise during life. The analysis has worked out under what circumstances it is optimal for spending to be rising, constant, or declining during life.

The calibrated model is able to predict spending shares and frailty across age groups, in the US, fairly well. Encouraged by these findings we used the calibrated model to elicit information about the dynamic impact on frailty and longevity from shocks to technology, income and more. One interesting outcome of the simulations were that changes in relative health prices and income have relatively modest effects on longevity. By way of contrast, the impact on longevity from changes in health efficiency is much larger. This suggests that governments aiming to improve health outcomes might be better off focusing on e.g. health technology, and on the institutional set-up of the health sector, than focusing on (say) subsidizing prices on health investments.

In a cross-country setting the calibrated model is able to account well for the celebrated Preston curve. This suggests that most of the observed link between income and longevity

across countries can be attributed to variations in health investments. The model does not fully account for the Preston curve; deviations reflect either reverse causality or omitted factors that are correlated with both income and longevity. This residual, however, only accounts for some 20 % of the variation in longevity along the income gradient; it would appear that the Preston curve mostly reflect the impact of income on longevity.

APPENDIX A: DERIVATIONS

Derivation of (6) and (7). The Hamiltonian associated with the problem of maximizing (4) subject to (2) and (5) reads

$$H = \frac{c^{1-\sigma} - 1}{1-\sigma} + \lambda\mu(D - a - Ah^\gamma) + \phi(rk + w - c - ph).$$

For $\sigma = 1$ the first term is replaced by $\log(c)$. The first order conditions wrt. c and h and the co-state equations are

$$c^{-\sigma} - \phi = 0 \quad \Rightarrow \quad c^{-\sigma} = \phi \quad \Rightarrow \quad \sigma\dot{c}/c = -\dot{\phi}/\phi \quad (\text{A.1})$$

$$-\lambda\mu A\gamma h^{\gamma-1} - p\phi = 0 \quad (\text{A.2})$$

$$\lambda\mu = \lambda\rho - \dot{\lambda} \quad \Rightarrow \quad \mu - \rho = -\dot{\lambda}/\lambda \quad (\text{A.3})$$

$$\phi r = \phi\rho - \dot{\phi} \quad \Rightarrow \quad r - \rho = -\dot{\phi}/\phi. \quad (\text{A.4})$$

Equation (A.4) is the well known Euler equation requiring that the shadow price of consumption (ϕ) grows at the rate of the interest rate less the time preference rate. Analogously, the Euler equation (A.3) requires that the shadow price of health grows at the rate of health deterioration (μ) less the time preference rate.

Log-differentiate (A.2) wrt. time and insert (A.3) and (A.4) to obtain optimal growth of health expenditure:

$$\frac{\dot{\lambda}}{\lambda} - \frac{\dot{\phi}}{\phi} = (1 - \gamma)\frac{\dot{h}}{h} \quad \Rightarrow \quad -\mu + \rho + r - \rho = (1 - \gamma)\frac{\dot{h}}{h}.$$

Observe that ρ cancels out. Intuitively, the growth rate of health expenditure depends positively on the growth rate of the the shadow price differential, i.e. the growth rate of λ/ϕ . If the shadow price of health (λ) grows at higher rate than the shadow price of consumption (ϕ), it indicates that the future contribution of health to utility is more important than the future contribution of consumption (both measured relative to current contribution) and thus it is optimal that health expenditure increases over time (\dot{h}/h is positive). Since the time preference enters both equations symmetrically, it has no significance for the *growth* of health expenditure. Of course it will affect the level of health expenditure, see main text. Since r enters only the growth rate of the shadow price of consumption and μ enters only the growth rate of the shadow price of health, they do not drop out but affect growth of health expenditure with opposite sign. Specifically solving for the growth rate of expenditure we get the ‘‘Health Euler’’, i.e. equation (7) of the main text. As usual, (6) is obtained by inserting (A.3) into (A.1). Note also, from (A.2), that the shadow price of health ϕ is negative because the associated stock variable D is a ‘‘bad’’ rather than a good.

Derivation of (9a)-(9c). Begin with noting that, because g_h is optimally constant according to (7), the differential equation (2) can be rewritten as $\dot{D} = \mu(D - a - Ah(0)^\gamma \exp(\gamma g_h t))$.

Given $D(0) = D_0$ the solution at time T is:

$$D(T) = D_0 \exp(\mu T) - \mu A h(0)^\gamma \exp(\mu T) \int_0^T \exp(\gamma g_h - \mu) dt + \mu a \exp(\mu T) \int_0^T \exp(-\mu t) dt.$$

At the time of expiry the boundary condition requires $D(T) = \bar{D}$. Solve the integrals in the above equation to get (9a) in the text. Next, integrate (5) and insert $k(0) = k_0$ and $k(T) = \bar{k}$ to obtain

$$\begin{aligned} \bar{k} &= k_0 \exp(rT) + w \exp(rT) \int_0^T \exp(-rt) dt \\ &\quad - c(0) \exp(rT) \int_0^T \exp[(g_c - r)t] dt - ph(0) \exp(rT) \int_0^T \exp[(g_h - r)t] dt. \end{aligned}$$

Divide by $\exp(rT)$. Note that $g_h - r = (\gamma r - \mu)/(1 - \gamma) \equiv g_D$ and solve the integrals to obtain (9b) in the text.

Finally note that optimal expiry requires that the Hamiltonian assumes the value of zero at T . Otherwise, it would have been optimal to live longer or die earlier. Also, at expiry $D(T) = \bar{D}$ and $k(T) = \bar{k}$. Thus the Hamiltonian reads

$$0 = H(T) = u(c(T)) + \lambda(T) \mu [\bar{D} - a - Ah(T)^\gamma] + \phi(T) [r\bar{k} + w - c(T) - ph(T)].$$

Insert $\lambda(T)$ and $\phi(T)$ from (A.1) and (A.2) to get

$$0 = u(c(T)) - \frac{p}{c(T)^\sigma} \left[\frac{(\bar{D} - a)h(T)^{1-\gamma}}{\gamma A} - \frac{h(T)}{\gamma} - \frac{w + r\bar{k}}{p} + c(T) + h(T) \right]$$

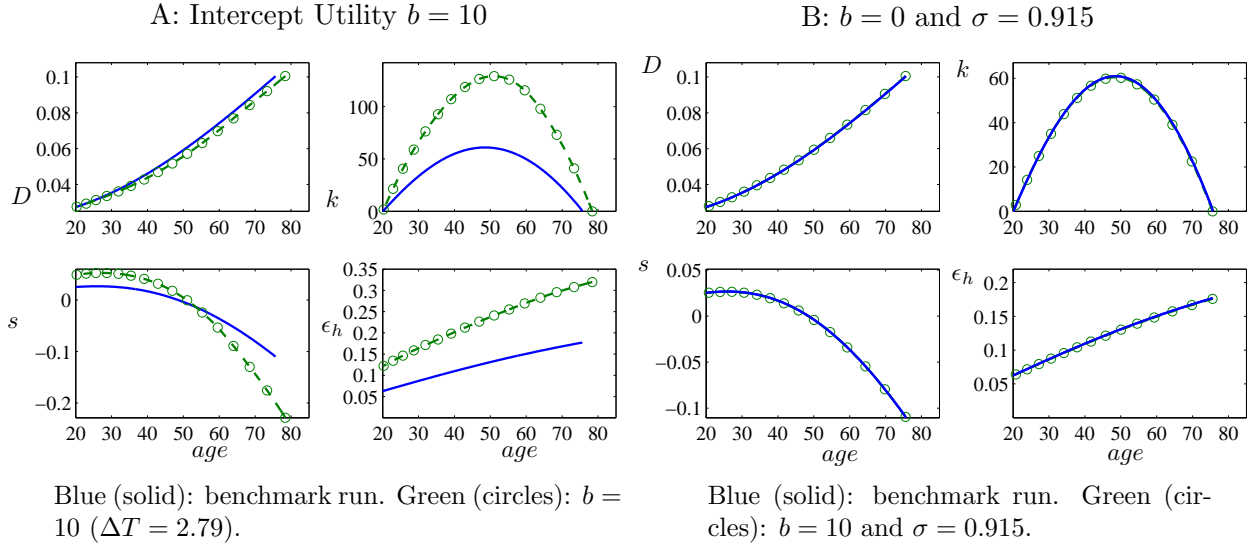
where $u_T \equiv \log c(T) + b$ in the case of log-utility and $u_T \equiv [c(T) - 1]^{1-\sigma}/(1 - \sigma)$ otherwise. Noting that $c(T) = c(0) \exp(g_c T)$ and $h(T) = h(0) \exp(g_h T)$ this provides (9c) in the text.

APPENDIX B: PREFERENCES

Panel A of Figure A.1 compares the benchmark run with the solution when health is a “superior good”, as in Hall and Jones (2007), reflected by green circled lines showing the solution for $b = 10$. The introduction of an autonomous component in utility raises health expenditure and the slope of the health expenditure curve, as in Hall and Jones. In our case, this implies that the model predicts too much and too steeply increasing health expenditure.

The counterfactually high demand for health can be controlled for by increasing the intertemporal elasticity of substitution. A higher IRS (lower σ) implies that the individual has less preference for consumption smoothing than in the log case. Since we maintain the benchmark assumption $r = \rho$, i.e. constant level of consumption throughout life, the change in preferences can only manifest itself in a higher initial level of consumption. This also brings lower savings, a lower level of health investments and thus less longevity. As Panel B demonstrates reducing σ to 0.915 exactly counterbalances the “superior good” assumption and the

FIGURE A.1: HEALTH AS A LUXURY GOOD



predicted life cycle trajectories (green circled lines) coincide with those of the benchmark model (blue solid lines).

APPENDIX C: DATA SOURCES

Data on life expectancy at age 20. Our samples for life expectancy pertains to males, and for circa 2000. Specifically, our sample involves observations covering the period 1997–2006. The data is available in the Demographic yearbook for 2006 . Alternatively, the data can be obtained online at the web address:

<http://unstats.un.org/unsd/demographic/sconcerns/mortality/mort2.htm>

In our sample, used for estimation, the median year of observation is 2003 with a standard deviation of 2.73 years.

Data on Labor productivity. In the regressions we employ GDP per worker (RGDPW) for 2000, from Penn World Tables, Mark 6.3.

REFERENCES

- Acemoglu, D. and S. Johnson, 2007, Disease and development: the effect of life expectancy on economic growth, *Journal of Political Economy* 115, 925-985.
- Arking, R., 2006, *The Biology of Aging: Observations and Principles*, Oxford University Press, Oxford.
- Australian Institute of Health and Welfare, 2004. *Australia's health 2004*. Australia's health no. 9.
- Barro, R., G. Mankiw and X. Sala-i-Martin, 1995. Capital mobility in neoclassical models of growth. *American Economic Review* 85, 103-115
- Becker, G., 2007. Health as human capital: synthesis and extensions. *Oxford Economic Papers* 59, 379-410.
- Bloom DE, Canning D., 2001. Cumulative causality, economic growth and the demographic transition. In: Birdsall N, Kelley AC and Sinding SW (eds). *Population Matters: Demographic Change, Economic Growth, and Poverty in the Developing World*. New York: Oxford University Press, pp. 165-97.
- Bloom, D. and D. Canning, 2007. Commentary: The Preston Curve 30 years on: still sparking fires. *International Journal of Epidemiology* 36, 498-499
- Browning, M. and Ejrnæs, M., 2009. Consumption and children. *Review of Economics and Statistics* 91, 93-111
- Cervellati, M. and U. Sunde, 2005. Human capital formation, life expectancy, and the process of development, *American Economic Review* 95, 1653-1672.
- Cervellati, M. and U. Sunde, 2010, Longevity and lifetime labor supply: evidence and Implications revisited, Working Paper (St. Gallen).
- Cervellati, M. and U. Sunde, 2011, Life expectancy and economic growth: the role of the demographic transition, *Journal of Economic Growth*, 16, 99-133.
- Chetty, R., 2006, A new method of estimating risk aversion, *American Economic Review* 96, 1821-1834.
- Cuttler, D. Mark B. McClellan, Joseph P. Newhouse, and Dahlia Remler, 1998. Are Medical Prices Declining? Evidence from Heart Attack Treatments. *Quarterly Journal of Economics* 113, 991-1024.
- Deaton, A., 2003. Health, inequality, and economic development. *Journal of Economic Literature* 41, 113-158.

- Dynan, K.E. and Skinner, J. and Zeldes, S.P., 2002. The importance of bequests and life-cycle saving in capital accumulation: A new answer. *American Economic Review* 92, 274–278.
- Ehrlich, I. and Chuma, H., 1990. A Model of the Demand for Longevity and the Value of Life Extension. *Journal of Political Economy* 98, 761-82
- Eisenring, C., 1999. Comparative dynamics in a health investment model. *Journal of Health Economics* 18, pp. 655-660
- Fonseca, R., P-C. Michaud, T. Galama, A. Kapteyn., 2009. On the Rise of Health Spending and Longevity. IZA discussion paper No. 4622
- Fries, J.F., 1980. Aging, natural death, and the compression of morbidity. *New England Journal of Medicine* 303, 130-135.
- Foster, M., 2001. The meaning of death: some simulations of a model of healthy and unhealthy consumption. *Journal of Health Economics* 20, 613-638.
- Galor O. and O. Moav, 2007. The Neolithic Origins of Contemporary Variation in Life Expectancy. Working Paper (Brown University)
- Gavrilov, L.A. and Gavrilova, N.S., 1991, *The Biology of Human Life Span: A Quantitative Approach*, Harwood Academic Publishers, London.
- Grossman, M., 1972. On the concept of health capital and the demand for health. *Journal of Political Economy* 80, 223-255
- Hall, R.E. and Jones, C.I., 2007. The Value of Life and the Rise in Health Spending. *The Quarterly Journal of Economics* 122, 39–72
- Hazan, M., 2009, Longevity and lifetime labor supply: evidence and implications, *Econometrica* 77, 1829-1863.
- Hazan, M. and H. Zoabi, 2006, Does longevity cause growth? A theoretical critique, *Journal of Economic Growth* 11, 363-76.
- Health Canada, 2001. *Health Expenditures in Canada by Age and Sex, 1980-81 to 2000-01*. Ottawa.
- Keehan, S. P, H. Lazenby, M. Zezza, and A.C. Catlin., 2004. Age Estimates in the National Health Accounts. *Health Care Financing Review* 1, 1–16.
- Lokshin, M., 2006. Difference-based semiparametric estimation of partial linear regression models. *The Stata Journal* 6, 377-83
- Mazzocco, M. and Személy, B., 2010, The decline of the US saving rate and the rise of health Expenditure, Working Paper, California Center for Population Research.

- Meara, E., C. White, and D. M. Cutler, 2004. Trends in Medical Spending by Age, 1963-2000. *Health Affairs* 23, 176-183.
- Mitnitski, A. and Bao, L. and Rockwood, K., 2006, Going from bad to worse: a stochastic model of transitions in deficit accumulation, in relation to mortality, *Mechanisms of ageing and development* 127, 490-493.
- Mitnitski, A.B. and Mogilner, A.J. and MacKnight, C. and Rockwood, K., 2002a. The accumulation of deficits with age and possible invariants of aging. *Scientific World* 2, 18161822.
- Mitnitski, A.B. and Mogilner, A.J. and MacKnight, C. and Rockwood, K., 2002b. The mortality rate as a function of accumulated deficits in a frailty index. *Mechanisms of ageing and development* 123, 1457-1460.
- Mitnitski, A. and Song, X. and Skoog, I. and Broe, GA and Cox, J.L. and Grunfeld, E. and Rockwood, K., 2005. Relative fitness and frailty of elderly men and women in developed countries and their relationship with mortality. *Journal of the American Geriatrics Society* 53, 2184-2189.
- New Zealand Ministry of Health, 2004. *Population Ageing and Health Expenditure: New Zealand 2002-2051*. Public Health Intelligence Occasional Bulletin No 22
- Preston, S.H., 1975. The changing relation between mortality and level of economic development. *Population Studies* 29, 231-248
- Pritchett, L. and Summers, L.H., 1996. Wealthier is healthier. *Journal of Human Resources* 31, 841-868
- Ried, W., 1998. Comparative dynamic analysis of the Grossman model. *Journal Health Economics* 17, pp. 383-425.
- Rockwood, K. and A.B. Mitnitski, A.B., 2006, Limits to deficit accumulation in elderly people. *Mechanisms of ageing and development* 127, 494-496.
- Rockwood, K. and Mitnitski, A.B., 2007, Frailty in relation to the accumulation of deficits. *Journals of Gerontology Series A: Biological and Medical Sciences* 62, 722
- Strulik, H., 2011, Health and education: understanding the gradient, Discussion paper, University of Hannover.
- Strulik, H., 2012, Optimal aging and uncertain death, Discussion paper, University of Hannover.

- Zellner, A., 1969, On the aggregation problem: A new approach to a troublesome problem. In: M. Beckmann and H.P. Künzi (eds.) : *Economic Models, Estimation and Risk Programming: Essays in honor of Gerhard Tintner*, Springer Verlag, Berlin.
- Zweifel, P. and Felder, S. and Meiers, M., 1999. Ageing of population and health care expenditure: a red herring? *Health Economics* 8, 485–496.