



Long-Run Economic Perspectives of an Ageing Society

Understanding Differences in Longevity across the EU Member States

Carl-Johan Dalgaard

Holger Strulik

Lepas Report No. D 4.2

April 2011

Available online at: <http://www.lepas-fp7.de/lepasrep-d42.pdf>

The project LEPAS – Long-Run Economic Perspectives of an Ageing Society – is a joint research initiative of the Universities of Alicante, Copenhagen, Hannover and the Vienna Institute of Demography. It is funded through the seventh framework programme by the European Community, Grant Agreement: SSH-2007-3.1.01- 217275. Webpage: <http://www.lepas-fp7.de/>

Understanding Differences in Longevity across the EU Member States

Lepas Report D4.2

Carl-Johan Dalgaard and Holger Strulik*

April 2011

Executive Summary. Longevity varies considerably across the post-enlargement EU member states. Evaluated around the year 2000 life expectancy at age 20 varied by roughly a decade while income per worker varied by about a factor of eight. With respect to income, convergence within the members of the European Union is to be expected in the years to come. Naturally, the question arises whether and how much the convergence of income will help to close gap in longevity. Moreover, beyond income, what are the most effective ways at raising longevity within the EU? In order to address these questions the Lepas Project has developed a physiologically-founded economic model of health demand over the life-cycle, drawing on recent advances in the modeling of aging and longevity from the natural sciences. The main conclusions summarized in this report are:

- Convergence of *income* per capita (labor productivity) across the EU member states can be expected to close a considerable fraction the life expectancy gap. But productivity advances are not the most powerful determinant of longevity.
- The most powerful determinant of increases in longevity is improvement in health *efficiency* (medical technological progress). Health efficiency should be the main policy target.
- Targeting *prices* (subsidies on wages or prices in the health sector) is a relatively ineffective way of increasing longevity.

A background paper to this report is available as Dalgaard and Strulik (2011).

*University of Copenhagen, Oester Farimagsgade 5, bygning 26,1353 Koebenhavn K, Carl.Johan.Dalgaard@econ.ku.dk and University of Hannover, Wirtschaftswissenschaftliche Fakultät, Königsworther Platz 1, 30167 Hannover, Germany; email: strulik@vwl.uni-hannover.de. 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).


1. THE FACTS

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. A 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. 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.”

Figure 1A depicts a recent – modified – version of the Preston curve for a cross-section of countries, focusing on labor productivity (GDP per worker in PPP\$) rather than GDP per capita, and life expectancy at age 20 rather than at birth. Despite these slight differences in the variables involved, the regularity is essentially that recorded by Preston: citizen’s of countries that are wealthier tends to be healthier. Moreover, the association between productivity and adult life expectancy is positive and non-linear, featuring a “flattening” at the top.


FIGURE 1: LIFE EXPECTANCY AND LABOR PRODUCTIVITY

A: The modified Preston curve



The figure shows the cross-country link between life expectancy at age 20 and GDP per worker cirka the year 2000. *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 Dalgaard and Strulik (2011) for details on the estimation algorithm. See Appendix for data sources.

B: Highlighting EU



The figure shows male life expectancy at age 20 vs. GDP per worker in PPP\$, inside and outside EU, cirka the year 2000: 85 countries. *Notes:* Life expectancy at age 20 is not recorded for all countries in the world, for which reason the sample only includes 25 out of the 27 current EU member States: Cyprus and Netherlands are missing.

As can be discerned from Figure 1B, there is considerable variation in the data within the European Union (EU); both in terms of labor productivity levels and life expectancy. In 2000 GDP per worker varied by a factor of eight (from about 12,000 PPP\$ in Romania to roughly 100,000 in Luxembourg), and life expectancy (recorded around the same time) varied by a remarkable 13 years (from ca. 46 years in Lithuania to a little in excess of 59 in Sweden).

The strong income gradient, visually obvious from Figure 1, suggests that higher income enables more direct investments in health, which slows down the aging process and prolongs life. It is thus tempting to infer that convergence in income levels across the EU member states would entail convergence in longevity. However, reflecting on the identity of the countries at the extremes of the two distributions leads one to the conclusion that there may be more to it than that: longevity is neither the lowest in Romania nor the greatest in Luxembourg. Hence, one might suspect that other factors could be responsible for the apparent link between longevity and income: Health technology or human capital for instance. To make matters worse, it seems impossible to rule out *a priori* that the positive correlation could be attributed (fully or partially) to an impact from longevity on income, i.e., the reverse line of causality.

Accordingly, in order to make progress it is necessary to try to unravel the forces that have shaped the Preston curve. For that purpose we suggest a theory-driven approach; by modeling the aging process, and the decision to invest in life prolonging initiatives, it is possible to examine the extent to which the income-to-health channel can account for the Preston curve by way of model simulation. Subsequently we then attempt to gauge the likely impact from income convergence on the distribution of health within the EU.

2. MODELING AGING

The analysis conducted in Dalgaard and Strulik (2011) aims to distil the logic of the income-to-longevity channel. The theory builds on two elements. First, a sound modeling of the aging process for a representative member of society, drawn from the natural science literature. The second element is the standard economics modeling of intertemporal choice of consumption for the representative consumer, yet extended by the additional option to invest in health, which serves to slow down aging. We describe these two elements in turn.

3. MODELING AGING

Following the underlying reasoning of reliability theory we think of aging as being characterized by increasing *frailty* (Gavrilov and Gavrilova, 1991; Lepas Report D2.1). 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). As humans age they 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.

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, however, 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:

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

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 *level* 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 at 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.

Specifically, the following parsimonious refinement of the process of deficit accumulation is employed in Dalgaard and Strulik (2011):

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

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 efficiency by which investments in health initiatives are converted into tangible results vis-a-vis reductions in speed of aging. Accordingly, A may naturally reflect the state of health technology, but also health institutions which influences how much of the health spending that is channeled into health initiatives rather than, say, administration and the like. 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.

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. While the force of aging is exogenous, it is unlikely to be universally constant.

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 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; chronological age t plays no independent role.

4. INTERTEMPORAL CHOICE

The modeling of the representative individual's economic choices follows the standard approach in dynamic economics. In each period of life the individual receives a wage income and, if wealth is positive, an interest income. This income can either be spend on non-health consumption, it can be saved for future purposes, or it can be invested in health, which, as explained above, serves to reduce the speed of aging and thus delays the time of death.

The representative individual derives utility from taking initiatives that lower the speed of aging because it delays the date of expiry, which is ultimately determined by investments in h . At the same time individuals enjoy other forms of consumption which are not life prolonging. As a result, the individual faces a trade-off between consuming today or making health investments, which allow for greater consumption in the future. Hence, the solution to the optimization problem provides a program for optimal consumption over the life cycle, along with optimal health investments.

The economic model of consumption demand captures two fundamental characteristics of human behavior, *consumption smoothing* and *time discounting*. Consumption smoothing reflects the law of diminishing marginal utility. Consuming a further unit now (e.g. another car or television set) provides more utility but less additional utility than the last consumption unit. Facing such a concave mapping from consumption into utility at every instant in time, consumers essentially prefer a smooth path of consumption expenditure over time. Consumption smoothing provides an incentive for health expenditure. Instead of using additional income to consume more now (and running into diminishing marginal utility) some income can be invested into health in order to prolong life and thus adding more periods over which consumption expenditure can be stretched (e.g. another year of driving pleasure). Time preference, on the other hand, captures human impatience. Individuals prefer to consume a given unit of goods now more than later on. For obvious reasons we expect the rate of time preference to be close to the interest rate, i.e. the price of present consumption in terms of future consumption.

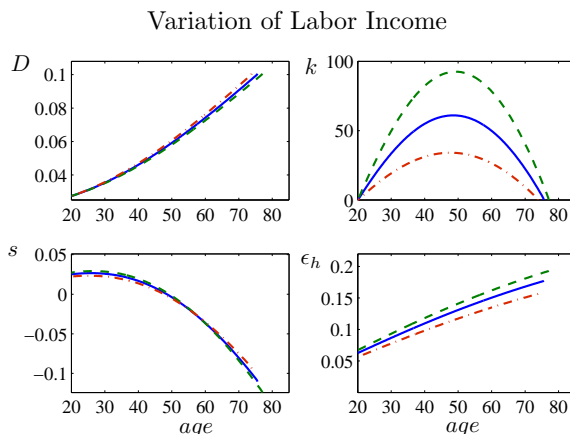
Given these fundamental trade-offs in human behavior it is a priori unclear how individuals would optimally like to allocate their life-time income on consumption, savings, and health investments now and in the future and how long they thus want to live given their life-time budget constraint and facing the force of aging. In order to answer this question we calibrated

the model of consumer behavior to US data. That is, the parameters of the model are chosen such that the model matches behavior of the average US male citizen exactly in several dimensions such as in terms of labor income, life expectancy, evolution of frailty over life and more. With the calibrated model in hand it is possible to conduct experiments, whereby parameters are changed after which the consequences for aging and longevity can be assessed.

5. POLICY EXPERIMENT 1: INCOME

Figure 2 shows how the representative individual reacts if his income is perturbed. The blue line shows life-trajectories for the representative US individual for frailty (deficit accumulation) D (top left panel), capital accumulation k (top right panel), the savings rate s (bottom left), and the health expenditure share ϵ_h (bottom right) from the age of 20 to death. The green (dashed) line is associated with an increase of labor income by 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.

FIGURE 2: HEALTH AND WEALTH OVER THE LIFE CYCLE



Green (dashed): wage income increases by 1/3. $\Delta T = 1.47$, implied elasticity 0.07). Red (dotted): wage income decreases by 1/3 ($\Delta T = -1.78$, implied elasticity -0.09).

The intuition for why higher income leads to longer life is found in the desire to smooth consumption. With increasing income, health spending rises more than regular consumption. This occurs as the incentive to smooth the latter is relatively stronger due to diminishing per period marginal utility. Higher income therefore leads to a larger adjustment in the level of health spending compared to non-health consumption. The end result is a slower speed of aging, and a longer life.

The issue of main interest, however, is the quantitative impact on longevity. As seen from the top left hand side corner of the figure, the impact is modest though not inconsequential. An increase of income of 1/3 (achievable in a generation with an income growth rate of about 1%

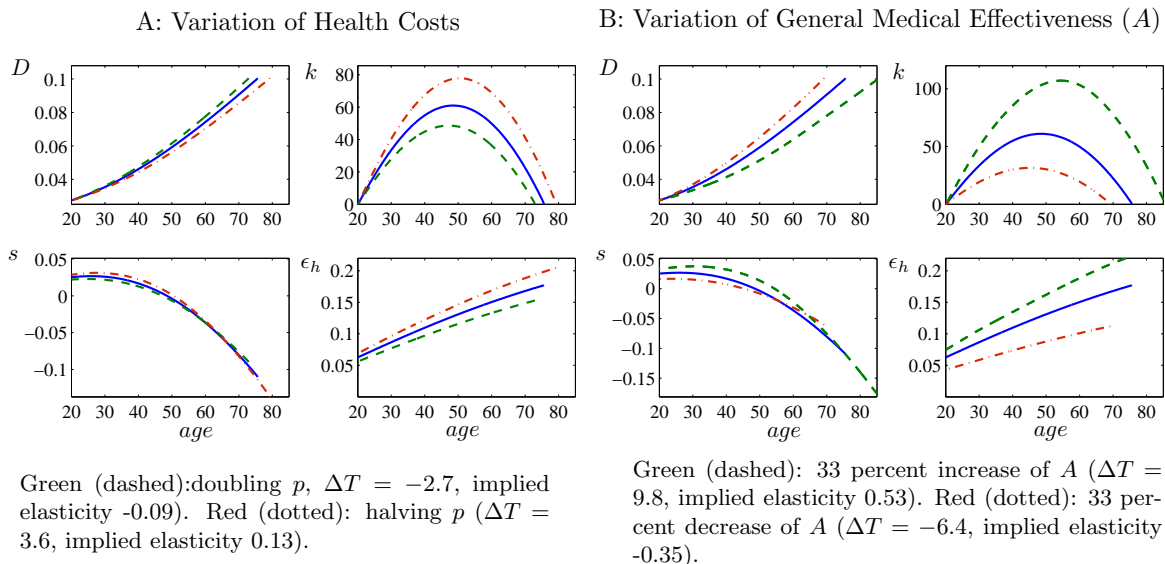
per year) translates into an increase in longevity of 1.7 years; the reduction involves a fall in longevity of 1.6 years. If we convert the impact into an elasticity - the elasticity of longevity with respect to income - we find it to be around 0.08, implying that a 100 percent increase of income would result in an 8 percent increase of life-expectancy. This effect is close to econometric estimates (e.g., Pritchett and Summers, 1996). Below we return to the issue of whether the magnitude of this effect is sizable enough to account for the Preston curve.

6. POLICY EXPERIMENT 2: HEALTH COSTS

The 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 price index of medical care has risen faster than the price index of GDP (Cutler et al., 1998).

As is clear from Panel A of Figure 3, when the relative price of health increases, individuals substitute towards regular consumption. As a result, the health share declines. With less health investments, savings (s) decline as well. The end result of a doubling of the relative health price for longevity is a reduction by 2.5 years. This amounts to a longevity-price elasticity of 0.09. While this elasticity is of roughly the same numerical size as the income elasticity it is well to bear in mind that the potential for variation is far smaller; income levels vary across the world easily by a factor of 30 or more, far towering any cross-country variation in price levels.

FIGURE 3: HEALTH COSTS AND MEDICAL EFFECTIVENESS



7. POLICY EXPERIMENT 3: HEALTH TECHNOLOGY

In Figure 3, panel B, we depict the impact of health productivity, i.e. general medical effectiveness (parameter A in equation 2) on longevity. 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 Panel B we examine the impact from increasing A by $1/3$; an increase by the same factor as that which we analyzed 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. It also suggests that improvements in health efficiency have a large potential to improve life expectancy.

8. THE PRESTON CURVE: IS INCOME ALL THAT?

As discussed above, a key issue is whether the income gradient, captured by the so-called Preston curve, reflects a causal impact of income on longevity mediated by health investments. 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 individual countries through the model, keeping the parameters we calibrated to US data fixed. Figure 5 allows for a visual comparison between the model's predictions regarding life expectancy in 2000, and the empirically estimated income gradient from the cross section in 2000, whereas Table 1 provides some summary statistics. Accordingly, the solid line represents the best fit of the income gradient to the data displayed in Figure 1; it therefore captured fitted values from estimating the model $l = f(y) + \epsilon$, where l is longevity, y is income and the function $f()$ is determined by the data. See Dalgaard and Strulik (2011) for details.

FIGURE 5: THE MODEL VS. THE PRESTON CURVE



The figure compares the empirically estimated Preston curve (dotted) to the Preston curve predicted by the model (solid). Source: Dalgaard and Strulik (2011).

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 does underestimate the range of life expectancies spanned by the Preston curve somewhat, but at average or median income the actual income gradient and the one predicted by the model essentially coincide. The table

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	51.3
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.

also provides 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 is thus accounted for, there are systematic deviations between the Preston curve and predicted life expectancy, which account for the remaining 22% of the variation. There are two complementary explanations for the systematic nature of the deviation; 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. Reverse causality arises if changes in life expectancy instigates changes in income. For instance, suppose that people who live longer also generally are healthier at any point in time during their lives, and therefore are capable of exerting more labor market effort generating higher labor productivity. If so, then the data depicted in figure 1 may not reflect an impact of income on longevity at all, but rather, the influence of longevity on labor productivity. Now, 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 the two explanations is more important in accounting for the left-over residual. What we do know is that they both may 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. 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.

What does this imply for the prospects of convergence in life expectancy within the EU? It is a well known feature of the growth process that relative income gaps tend to narrow over time, between countries sharing structural characteristics like investment rates, population growth, institutions etc. This has led to convergence within Europe in the past (e.g., Barro and Sala-i-Martin, 1991), and this process will presumably continue in the future. So, suppose the poorest EU countries gradually manage to catch-up with the richest EU countries: how would this affect the distribution of life expectancy?

The analysis above allows us to answer this question employing the calibrated “income elasticity” (see above). If income convergence is taking place it implies that income dispersion across EU drops. But how big of a reduction in inequality in life expectancy would this instigate? Our analysis would suggest that the answer is: Not a lot. More specifically, for each percentage point reduction in the standard deviation in (log) income per worker we would expect to see a reduction of 0.01 percent in the standard deviation of (log) life expectancy.¹ Or, to put it differently: if inequality in income is cut in half, inequality of life expectancy is reduced by 0.6%.

¹The calculation works as follows. Assume that life expectancy is related to income in the following fashion: $\log(l) = \lambda \cdot \log(y)$, which implies that $std(\log(l)) = \lambda \cdot std(\log(y))$, where $std()$ denotes the standard deviation across EU regions. The elasticity of inequality in longevity with respect to inequality in income is then $\lambda \cdot std(\log(l)) \cdot std(\log(y))^{-1}$. Around 2000 we have $std(\log(l)) = 0.08$, $std(\log(y)) = 0.6$; inserting $\lambda = -0.09$ into the formula leads to the stated result.

The calculation illustrates that convergence in income will enable convergence in life expectancy in the years to come within the EU; but not in a big way. It is also worth bearing in mind, that the convergence process is believed to be relatively slow. A typical estimate for the speed of income convergence across regions would be that roughly 2 % of the income gap is eliminated per year, implying that it will take about 35 years to cut income inequality in half within the EU, and thus attain the above mentioned reduction in inequality in longevity. This serves to underscore the insight from the policy experiments that that health efficiency should be a major policy theme within the EU.

9. CONCLUDING REMARKS

The present report has summarized progress in the modeling of the economics of aging. In particular, the theory allows us to examine the impact from various policy initiative (such as initiatives leading to a reduction of the health inefficiencies, or health prices) on longevity. It also allows us to gauge the impact of income convergence on inequality in longevity at age 20. Our analysis reveals that income convergence will lead to reductions in “health inequality” across EU; but not by much. The key focus area for policy should be health efficiency.

The present research program will be further developed in various ways. First, the analysis shows the need to understand certain physiological characteristics of the European populations better. Specifically, there is a need to understand the speed of aging, which is an issue that can be addressed empirically. Second, the framework presented above can be extended to a scenario where the representative consumer also chooses optimal labor supply. This will allow for an analysis of optimal retirement, which will shed light on an issue of key public finance interest: what is the impact of changes in income, health prices and efficiency for the length of time between (optimal) retirement and (optimal) death? These will be areas of focus for the LEPAS consortium.

REFERENCES

- Barro, R. and X. Sala-i-Martin, 1991. Convergence across states and regions. *Brookings Papers on Economic Activity*, 107-82.
- Bloom, D. and D. Canning, 2007. Commentary: The Preston Curve 30 years on: still sparking fires. *International Journal of Epidemiology* 36, 498-499.
- Cutler, 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.
- Dalgaard, C-J and H. Strulik, 2011. Optimal aging and death: Understanding the Preston curve. Discussion Paper 11-09 (University of Copenhagen).
- Gavrilov, L.A. and Gavrilova, N.S., 1991, *The Biology of Human Life Span: A Quantitative Approach*, Harwood Academic Publishers, London.
- 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.
- 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.

APPENDIX: DATA SOURCES

Data on life expectancy at age 20. We use a samples of data on life expectancy at the age of 20 for males in “2000”. Specifically, the sample involves observations covering the period 1997–2006. The data is available in the Demographic yearbook for 2006 (2000 sample). Alternatively, the data can be obtained online at the web address:

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

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