

Human Lifetime Entropy in a Historical Perspective (1750-2014)

Patrick Meyer, Grégory Ponthière

► **To cite this version:**

Patrick Meyer, Grégory Ponthière. Human Lifetime Entropy in a Historical Perspective (1750-2014). 2016. halshs-01409679

HAL Id: halshs-01409679

<https://halshs.archives-ouvertes.fr/halshs-01409679>

Submitted on 6 Dec 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



PARIS SCHOOL OF ECONOMICS
ÉCOLE D'ÉCONOMIE DE PARIS

WORKING PAPER N° 2016 – 30

Human Lifetime Entropy in a Historical Perspective (1750-2014)

**Patrick Meyer
Grégory Ponthière**

JEL Codes: D80, J11, N30

Keywords: mortality risk, longevity, age at death, entropy, measurement



PARIS-JOURDAN SCIENCES ÉCONOMIQUES

48, Bd JOURDAN – E.N.S. – 75014 PARIS
TÉL. : 33(0) 1 43 13 63 00 – FAX : 33 (0) 1 43 13 63 10
www.pse.ens.fr

Human Lifetime Entropy in a Historical Perspective (1750-2014)

Patrick Meyer* Gregory Ponthiere†

Abstract

Although it is widely acknowledged that life is risky, it is difficult to provide an intuitive indicator of the riskiness of life, whose metric would have a concrete counterpart for the layman. This paper uses the Shannon entropy index to the base 2 to quantify, in terms of bits (i.e. the amount of information revealed by tossing a fair coin), the risk relative to the age at death in 37 countries from the Human Mortality Database. We identify 5 major stylized facts: (1) over the last two centuries, (period) life entropy at birth exhibits an inverted U shape pattern with a maximum in the first half of the 20th century (at about 6 bits); (2) over the last 150 years, Western countries have converged in terms of (period) life entropy at birth towards levels of 5.6 bits for men and 5.5 bits for women; (3) curves of (period) life entropy at birth for men and women crossed during the 20th century; (4) the entropy age profile shifted from a non-monotonic profile (in the 18th and 19th centuries) to a strictly decreasing profile (in the 20th and 21st centuries); (5) men exhibit a higher life entropy than women below ages 50-55, and a lower one after ages 50-55.

Keywords: mortality risk, longevity, age at death, entropy, measurement.

JEL codes: D80, J11, N30

*IMT Atlantique Bretagne-Pays de la Loire, Institut Mines Télécom, UMR CNRS 6285 Lab-STICC, Technopôle Brest Iroise, CS 83818, 29238 Brest Cedex 3, France

†University Paris East (ERUDITE), Paris School of Economics, and Institut universitaire de France. Address: ENS, bd Jourdan 48, 75014 Paris. E-mail: gregory.ponthiere@ens.fr

1 Introduction

Mors certa, hora incerta. As the philosopher Jankélévitch emphasized in his treatise *La Mort* (1977), uncertainty about the duration of life constitutes a major dimension of human condition. Everyone knows that death is inevitable, and, hence, that one will necessarily die one day. But no one knows *when* precisely (which year, which day, which hour) he or she will die.

Uncertainty about the time of death may, according to Jankélévitch, have significant psychological effects regarding how one perceives one's finite existence. As long as one does not know on which day he or she will die, some form of desperate, irrational, belief may emerge: everyone knows that he or she will die one day or another, but when a new day comes, one believes that death will not arise on that day, but, rather, on *another* day (in the future), and so on and so forth.

Irrational denial of death can emerge since individuals have, in general, no clear idea concerning the size of the risk relative to the age at death. True, individuals know life expectancy statistics, which inform them about the expected duration of life conditionally on the prevailing age-specific mortality rates. But since individuals have little idea of the amount of risk relative to the age at death, and of the speed at which this amount of risk is resolved as they become older, they can believe, even at old ages, that they have a large probability to survive to even higher ages.

Lack of quantification of the risk about the age of death - and the associated denial of death - may partly explain some contemporary puzzles in the economics of risk and insurance.¹ Let us take, for instance, the long term care (LTC) private insurance puzzle (Brown and Finkelstein 2007, 2011). Despite the large probability (about 30 to 50 percents) to enter a nursing home at the old age (Brown and Finkelstein 2009), and despite the large costs of LTC, few people purchase private LTC insurance.² Several causes may explain this paradox, both on the supply side (high loading factors) and on the demand side (crowding out by family care or by social insurance). But the denial of death - and, hence, the denial of dependency preceding death - may also explain the LTC private insurance puzzle.

Ignorance concerning the risk relative to the age at death may seem, at first glance, surprising, since various measures of life riskiness were developed by demographers on the basis of survival curves. More rectangular survival curves are associated with a lower degree of risk about the age at death, the rectangular being the hypothetical case where all members of a cohort would die at the same age (Wilmoth and Horiuchi 1999). Indicators of life riskiness include measures of the verticalization of the survival curve (capturing the concentration of deaths around a normal age at death).³ Other indicators include measures of the horizontalization of survival curves, which capture how long a cohort can live before aging-related deaths significantly reduce the proportion of survivors (Cheung et al 2005).⁴

But indicators of verticalization / horizontalization of survival curves remain hard to understand for non-demographers. Most individuals have some idea of life expectancy, but few of them have an accurate idea of how risky their life is. One possible explanation for this may lie in the *units* used for measurement. Life expectancy is measured in years, which is something most common for individuals, and makes them familiar with those indicators. But years are more difficult to use as measurement units of the degree of life riskiness. For instance, one may have difficulties to understand what a standard deviation of lifespan of x years represents in terms of life riskiness.

The goal of this paper is to develop an intuitive indicator of the riskiness of life, whose metric would have a concrete counterpart for the layman. When considering risk, the most familiar measurement unit is the bit, defined as the amount of information revealed by tossing a fair coin (Pierce 1980). Tossing a coin is the most common randomization device, with which individuals are most familiar. Since at least the Roman empire, individuals are used to flip coins. Hence, if one wants to build a

¹Note that the denial of death can affect various economic decisions. Kopczuk and Slemrod (2005) explore the consequences of the denial of death on savings behavior.

²According to Brown et al (2007), only 9 to 10 percents of the population at risk of LTC have purchased a private insurance in the U.S.

³Such measures include the standard deviations of lifespans, interquantile range, or the coefficient C_{50} developed by Kannisto (2000), which measures the shortest age interval that is necessary to concentrate 50 % of life durations.

⁴One example is the age reached by some high percentile of survivors in a life table (e.g. the age reached by 99 % of survivors).

measure of life riskiness that has a concrete significance for all individuals (including non specialists), it makes sense to measure the risk about the duration of life in terms of bits.

This paper develops a measure of the degree of risk about the duration of life whose metric is the bit, and which can thus be interpreted in terms of the amount of information revealed by flipping a fair coin. We propose to measure the degree of risk about the duration of life by means of Shannon's entropy, defined to the base 2. Relying on Shannon's entropy to the base 2 allows us to make simple comparisons between, on the one hand, the amount of information revealed by a death, and, on the other hand, the amount of information revealed from tossing a coin (i.e. one bit). We are thus able to quantify the risk about the duration of life that remains *unresolved* at a given age in terms of bits, i.e. the quantity of risk unresolved when tossing a fair coin. Then, we use period and cohort lifetables for 37 countries from the Human Mortality Database, to study the size and evolution of life entropy by country, gender and age over the last two centuries.

Anticipating our results, we identify 8 major stylized facts: (1) over the last two centuries, (period) life entropy at birth exhibits an inverted U shape with a maximum in the first half of the 20th century (at a level around 6 bits); (2) the variability of (period) life entropy at birth has declined over time; (3) curves of (period) life entropy at birth for men and women crossed during the 20th century; (4) over the last 150 years, Western countries have been converging in terms of (period) life entropy at birth towards about 5.6 bits for men and 5.5 bits for women; (5) the entropy age profile has shifted from a non-monotonic profile (in the 18th and 19th centuries) to a strictly decreasing profile (in the 20th and 21st centuries); (6) the entropy age profile can be approximated by a piecewise linear function with three distinct segments, coinciding with a respectively, low, medium and high speed of learning about one's duration of life; (7) men exhibit a higher life entropy than women below ages 50-55, and a lower one after ages 50-55; (8) comparing cohorts born in the 19th and 20th centuries, the entropy age profile has tended to shift upwards (with a gain of 0.5 bits at each age during the active life) and to the right (with a gain of about 2 bits beyond age 100).

This paper complements the literature on the measurement of the variance of the age at death using entropy indicators, such as Demetrius (1976), Nagnur (1986), Hill (1993) and Noymer and Coleman (2013). Our contribution with respect to those studies is twofold. First, we use, unlike those studies, Shannon entropy to the base 2, in order to make our measures of life entropy rely on the intuitive bit metric. Second, we apply our measure to the Human Mortality Database, and consider various period/cohort measures of entropy at birth and along the life cycle, in order to quantify how large life riskiness is, and of how life riskiness varies across ages, genders, countries and epochs.

The rest of the paper is organized as follows. Section 2 presents Shannon entropy to the base 2, and explains how one can interpret its values in terms of the amount of information revealed by coin flipping. Section 3 presents the evolution, for men and women, of life entropy around the world (1750-2014) using the Human Mortality Database. Then, Section 4 examines life entropy by age, and examines the speed at which risk about the age at death is resolved as individuals become older. Section 5 examines the issue of unanticipated risk about the age at death, by comparing period life entropy with cohort life entropy. Section 6 concludes.

2 Shannon entropy as a measure of life riskiness

Let us first introduce Shannon entropy to the base 2 and explain why this specific indicator has some intuitive appeal for the purpose of measuring and comparing life riskiness across periods and countries.

Actually, considering the degree of risk about the duration of life in the light of a measure of life entropy amounts to regard the life of an individual as a lottery, whose different possible outcomes correspond to different ages at death. The outcome of that lottery of life - the age at death - is revealed only at the time of the death of the individual. From that perspective, one can measure the degree of risk of the lottery of life by the amount of information that we learn when the outcome of that lottery is realized, that is, when the individual dies (his age at death is then revealed).

The amount of information obtained when the age at death is revealed can be measured by life

entropy. Entropy is a general measure of the information revealed by an event (Pierce 1980). It can be regarded as a measure of the uncertainty resolved. When applied to the uncertainty about the duration of life, this allows us to measure the degree of risk associated to the age at death. The event under study is thus the "death of someone". The information that we learn in case of the event "death at a given age" is given by the measured life entropy.

Various formulas of entropy are available, depending on the logarithm that is used. In the following, we rely on entropy to the base 2, which is standard in information theory. When assuming a maximum duration of life equal to 115 years, the life entropy to the base 2 at age k , denoted by H_k , is:

$$H_k = - \sum_{i=k}^{115} p_{i,k} \log_2 (p_{i,k}) \quad (1)$$

where $p_{i,k}$ is the probability of a (remaining) life of duration i conditional on surviving up to age k , with $k \leq i \leq 115$.

The formula for life entropy at age k can be interpreted as follows. Take, for instance, a person dying at age k . That person could potentially have lived more years, and could have potentially died at many other (higher) ages ($k+1, k+2, \dots$). The event of the death of that person at age k allows us to learn an amount of information equal to H_k . This amount of information measures the quantity of risk about the duration of life that remained unresolved at age k , just before the person died. When the amount of information revealed by the death event is large, this means that the degree of riskiness involved in the lottery of life was also large. On the contrary, when the amount of information revealed by the death event is low, then the degree of life riskiness was also low. The amount of information revealed by the death of someone depends on (1) the age at which death occurs; (2) the distribution of the possible durations of life at the age at death.

Entropy is measured in bits, defined as the amount of information revealed by flipping a fair coin. In our case, where we focus on the entropy of life, the measurement unit consists of "bits per life" (or "bit per death" to be more accurate, since it is the death event that reveals the information). We will thus say that the entropy of life for an individual of age k is equal to H_k bits per life, or, in short, H_k bits. Saying that the life entropy of an individual at age k equals H_k bits means that for a person of age k , the amount of risk about the duration of life that is unresolved at age k equals H_k bits. Hence, if that person were to die at age k , the amount of information that would be learnt would be equal to H_k bits. One can thus interpret H_k bits as a measure of life riskiness at age k .

To illustrate this, Figure 1 considers two extreme lotteries of life. Consider first the case of a degenerate lottery of life, that is, a perfectly rectangular survival curve. In that case, the amount of information learnt when the outcome of the lottery of life is realized is zero. Indeed, when a person dies, nothing is learnt, since it is known in advance that each individual will die at the same age. Hence the age at death that is realized does not bring any amount of information. That degenerate lottery is thus associated to the minimal life entropy at birth (equal to 0). Consider now the extreme case of a triangular survival curve, each age at death being equiprobable (with a probability $\frac{1}{115}$). In that case, the amount of information revealed when an individual dies is much larger. The degree of risk associated to that lottery of life is maximum. The triangular survival curve is associated with the maximal life entropy at birth (equal to 6.905).

The main appeal of the Shannon entropy to the base 2 for the measurement of the degree of risk of life lies in the intuitive meaning of its measurement unit. Clearly, if one keeps in mind that the amount of information revealed by flipping a fair coin is equal to 1 bit, one can easily interpret the amount of information revealed by dying at a particular age under a particular lottery of life. For instance, if we take the triangular survival curve, we have that the amount of information learnt at the death of an individual at age 0, equal to 6.905 bits, is slightly inferior to what is learnt when flipping 7 fair coins (i.e. 7 bits). Thus, in the hypothetical case where each duration of life would be equally likely, the degree of riskiness of a life would be almost equivalent to the degree of risk involved in tossing 7 fair coins. This gives some concrete idea of how risky a life is.

Obviously, the two cases shown on Figure 1 are extreme examples. In real life, survival curves are




Figure 1: Rectangular and triangular survival curves for the same life expectancy (57.5) and the associated life entropy.

neither rectangular, nor triangular. Hence the associated degree of life entropy lies strictly within the interval $]0, 6.905[$ bits. Moreover, we know that, during the last two centuries, survival curves have been subject to a rectangularization process, implying a reduction of the variance of the age at death. But whereas this tendency implied a reduction in the associated life entropy, it is difficult to know to what precise extent life entropy has decreased. The goal of the next sections is to use Shannon entropy to the base 2 to cast some light on the evolution of survival curves over the last centuries.

3 Life entropy in a historical perspective

This section uses lifetables data from 37 countries in the Human Mortality Database to study the evolution of Shannon life entropy to the base 2 across countries, periods and genders.⁵

As a starting point, this section focuses on life entropy *at birth*, based on period lifetables. Life entropy at birth measures the amount of risk about the age at death that is unresolved at age 0 (or, alternatively, the amount of information revealed by the death of someone at age 0). It constitutes thus a measure of the riskiness of life as a whole, conditionally on the prevailing lifetable. Life entropy at birth provides thus a synthetic measure of the degree of risk relative to the age at death when considering the entire potential lifespan.⁶

3.1 Life entropy at birth by country and gender

Let us first take the subsample of countries with the longest time series. For that purpose, Figure 2 shows, over the period 1750-2014, the evolution of life entropy at birth for women and men in France, England and Wales, Iceland and Sweden.

⁵Figure 11 of Appendix A summarizes the availability, in terms of year interval, of period lifetables for each country within the database.

⁶Section 4 will consider the evolution of life entropy by age, that is, along the life cycle, to examine the speed at which the risk about the duration of life is resolved as individuals become older.

On each graph, the y axis measures life entropy at birth, expressed in bits units. Those measures can be interpreted as follows. A life entropy at birth equal to 6 bits for French women in 1850 means that, based on the survival conditions prevailing in 1850 for French women, the lottery of life faced by French women was, at that time, equivalent, in terms of risk, to the flipping of 6 fair coins, in the sense that the amount of information revealed by the death of a French women at age 0 is equal to the amount of information revealed by tossing 6 fair coins.




Figure 2: Life entropy at birth (period) in France (FRATNP), England and Wales (GBRTENW), Iceland (ISL) and Sweden (SWE), for women (f) and men (m)

A first major stylized fact is that life entropy at birth exhibits, over the period considered, an inverted U shape, with a maximum level during the first part of the 20th century. The increasing part of the inverted U pattern is explained by the decline in infant mortality, which, by making more scenarios of life duration possible, increased the risk about the duration of life. The decreasing pattern in the second part of the 20th century reflects the rectangularization of the survival curve, which

reduced the degree of life riskiness (by making the scenarios of death before the old age less likely).

A second point to be stressed consists of the variability of life entropy at birth. Life entropy exhibits very large fluctuations in the 19th century, and much smaller variability after 1950. This fact is well illustrated by the case of Island, which exhibits variations with magnitudes superior to 1.5 bit in the 1840s and 1880s, whereas fluctuations are of much smaller magnitudes in the 20th century.

A third observation concerns the comparison of life entropy at birth across gender. In each country considered, entropy at birth is, at the beginning of the period under study, higher for women than for men, whereas, at the end of the period, life entropy at birth is lower for women than for men. There has been an inversion of the gender rank in terms of life entropy at birth.

It should be stressed that those three stylized facts are observed for all countries for which we have long time series (see the Appendix B for the entropy figures).

Let us now focus on some large countries. Figure 3 shows, for women and men, life entropy at birth in Japan, Russia and the U.S. during the 20th century. Two main observations can be made. First, whereas all countries under comparison exhibit quite similar entropy levels around 1960 (equal to about 6 bits for men and 5.9 bits for women), the pattern of life entropy varies strongly across countries. While life entropy at birth for women in Japan has declined by about 0.4 bit between 1960 et 2010, it has stagnated around 5.8 bits in Russia. Thus the evolution of life riskiness varies strongly across those countries. Second, although series for men and women in the same country are correlated, life entropy is, over the 20th century lower for women than for men, in line with the right part of Figure 2. Moreover, the decline in life entropy differs across genders. In the case of Japan the decline was larger for women.




Figure 3: Life entropy at birth (period) in Japan, Russia and the U.S., for women (left) and men (right)

3.2 Life entropy at birth: convergence or divergence?

Let us now examine whether the different countries converge or diverge in terms of life entropy at birth. To examine that issue, Figure 4 plots, for women (top) and men (bottom), the (period) life entropy at birth in the countries for which we have observations starting in 1855 or before. We also show the associated standard deviation coefficient.




Figure 4: Life entropy at birth (period) in Denmark, England and Wales, Finland, France, Iceland, Italy, Netherlands, Norway, Scotland, Sweden and Switzerland and associated standard deviation by year, for women (f) and men (m)

As shown on Figure 4, the countries under comparison exhibit quite different levels of life entropy at birth in 1855, ranging between 5.3 bits and 6.3 bits for men, and between 5.6 bits and 6.3 bits for women. Moreover, there has been, over the period 1855-2008, a significant convergence of countries in terms of life entropy at birth, with entropy levels converging towards 5.6 bits for men and 5.5 bits for women. The extent of the convergence is measured by the standard deviation coefficient, which has, for men (resp. women), decreased from 0.28 (resp. 0.20) to 0.10 (resp. 0.10) between 1855 and 2008. This convergence between Western countries, which arises for both men and women, constitutes another important stylized fact.

4 Life entropy along the lifecycle

Up to now, we focused only on life entropy measured at birth. That indicator captured the degree of life riskiness while taking the human lifespan as a whole. However, we can also compute life entropy for all other ages, and consider how life entropy evolves along the life cycle.

Focusing on life entropy at all ages of life allows us to examine how quickly the risk about the duration of life is resolved as a person becomes older. The intuition goes as follows. The entropy at a given age provides the amount of risk about the duration of life that remains unresolved at that age. Hence considering how life entropy evolves with the age allows us to learn how the risk about the duration of life is being progressively resolved as the individual becomes older.

Put it differently, individuals of different ages are not facing the same lottery of life, and do not face the same degree of life riskiness. Comparing life entropy at different ages allows us to quantify how different those lotteries are in terms of riskiness.

4.1 Life entropy age profiles across centuries: the case of Sweden

In order to study the evolution of the life entropy age profile over long periods of time, let us first focus on the case of Sweden. Figure 5 shows, for women and men, the life entropy age profile for years 1751, 1839 and 2014. Each curve can be interpreted as follows: it gives, for each age, the quantity of information that one will learn in case of death of the person at that age, or, alternatively, the amount of risk about the duration of life that remains unresolved at that age.

Figure 5 can be interpreted as follows. Take, for instance, the entropy age profile for women in 2014. If one reads the graph from the x axis to the y axis, the entropy age profile tells us that, on the basis of the survival conditions prevailing for women in Sweden in 2014, the amount of risk about the duration of life is, at age 80, equal to about 4.5 bits. Alternatively, one can read the graph from the y axis to the x axis, and look for critical ages at which the remaining risk about the age at death is equivalent to, for instance, the amount of risk associated to tossing a given number of fair coins. For instance, the age at which the remaining risk to be resolved is equivalent to 3 fair coins is 95 years.




Figure 5: Period life entropy by age in Sweden for women (left) and men (right)

Three main observations can be made. First, the overall shape of the life entropy age profile - in particular its monotonicity - has changed over time. In the 18th and 19th centuries, life entropy was,

at the beginning of life, increasing with the age, and then decreasing after infancy. On the contrary, in the 20th century, life entropy decreases monotonically over the entire age interval. That change can be explained by the strong reduction of infant mortality during the last centuries. In the past, surviving the first few years of life was opening many possibilities in terms of life duration, and, hence, increasing the amount of risk about the duration of life. Nowadays, infant mortality is much lower, so that lifetime entropy does no longer increase with the age during early childhood. That first observation can be made for all countries in the sample for which we have a sufficiently long time series. Hence, this shift from a non-monotonous to a monotonous entropy age profile constitutes a robust stylized fact.

A second observation concerns the slope of the life entropy age profile during adulthood, that is, the speed at which risk about the duration of life is reduced as the individual becomes older. The life entropy age profile can be approximated by a piecewise linear function with three segments. The slope of each segment of the function is, in absolute value, increasing with the age on all profiles under comparison. At very high ages, a larger amount of risk about the duration of life is resolved as the person becomes one year older.⁷ The decomposition of the life entropy age profile into those three segments has significantly changed over time. The first segment, which coincides with childhood and young adulthood, has, over the years, become spread on a longer age interval, and has also become more flat, which means that, in the 21st century, a much smaller quantity of risk about the duration of life is resolved as the individual becomes older below age 50, in comparison with what used to prevail in the 18th and 19th centuries. The slopes of the second and the third segments have remained quite similar under the three profiles under comparison, but the increase in the age interval associated to the first segment has pushed those two other segments to the right.

A third observation consists in the tendency of life entropy curves associated to the 19th and 20th century to cross somewhere in the middle of life. In other words, life entropy has tended, over the last centuries, to decrease in the young ages of life, and to increase in the older ages of life. For instance, life entropy at age 20 for men was equal to about 6 bits in 1751, against about 5.5 bits in 2014. On the contrary, life entropy at age 80 was equal to about 3.7 bits in 1751, against about 4.2 bits in 2014. However, due to the intersection, life entropy has changed little around age 50.

4.2 Life entropy age profile across countries

Let us now compare life entropy age profiles across countries and gender. For that purpose, Figure 6 shows, for men and women, life entropy profiles by age for Australia, France, England and Wales, Japan, Russia and the U.S., for the first and the last observations. The comparison of life entropy age profiles across countries, gender and period allows us to identify two major stylized facts.

First, whatever the country, the gender and the period considered (abstracting from the short increasing segment mentioned above), all entropy age profiles exhibit a kind of piecewise linear function with three segments whose slope is, in absolute value, increasing with the age. That feature is a robust stylized fact, which can be observed for each country and gender under study (beyond the 6 countries shown here).

This piecewise linear form reveals the existence of three ages in terms of learning about one's duration of life. The first segment coincides with the young age, during which very little information about the duration of life is learnt when the individual becomes older. The second segment, which starts nowadays around age 55 and lasts around age 100, is the mature age, during which one starts learning, year after year, more about one's duration of life. During that period, the amount of risk about one's age at death falls from about 5 bits to 2 bits. Such a 3-bits reduction of risk is substantial, but is spread on an age interval exceeding 40 years. Finally, the third segment, which starts around age 100, coincides with the very old age, during which there remains little risk about the duration of life (about 2 bits), and during which the learning about the duration of one's life is extremely rapid: the remaining 2 bits of entropy vanish in just a few years.

Another robust stylized fact concerns the comparison of life entropy profiles by gender. As shown in the Figure, the entropy age profiles for men and women intersect around age 50. Clearly, men

⁷At very high ages, risk about death is equivalent to tossing a unique coin.

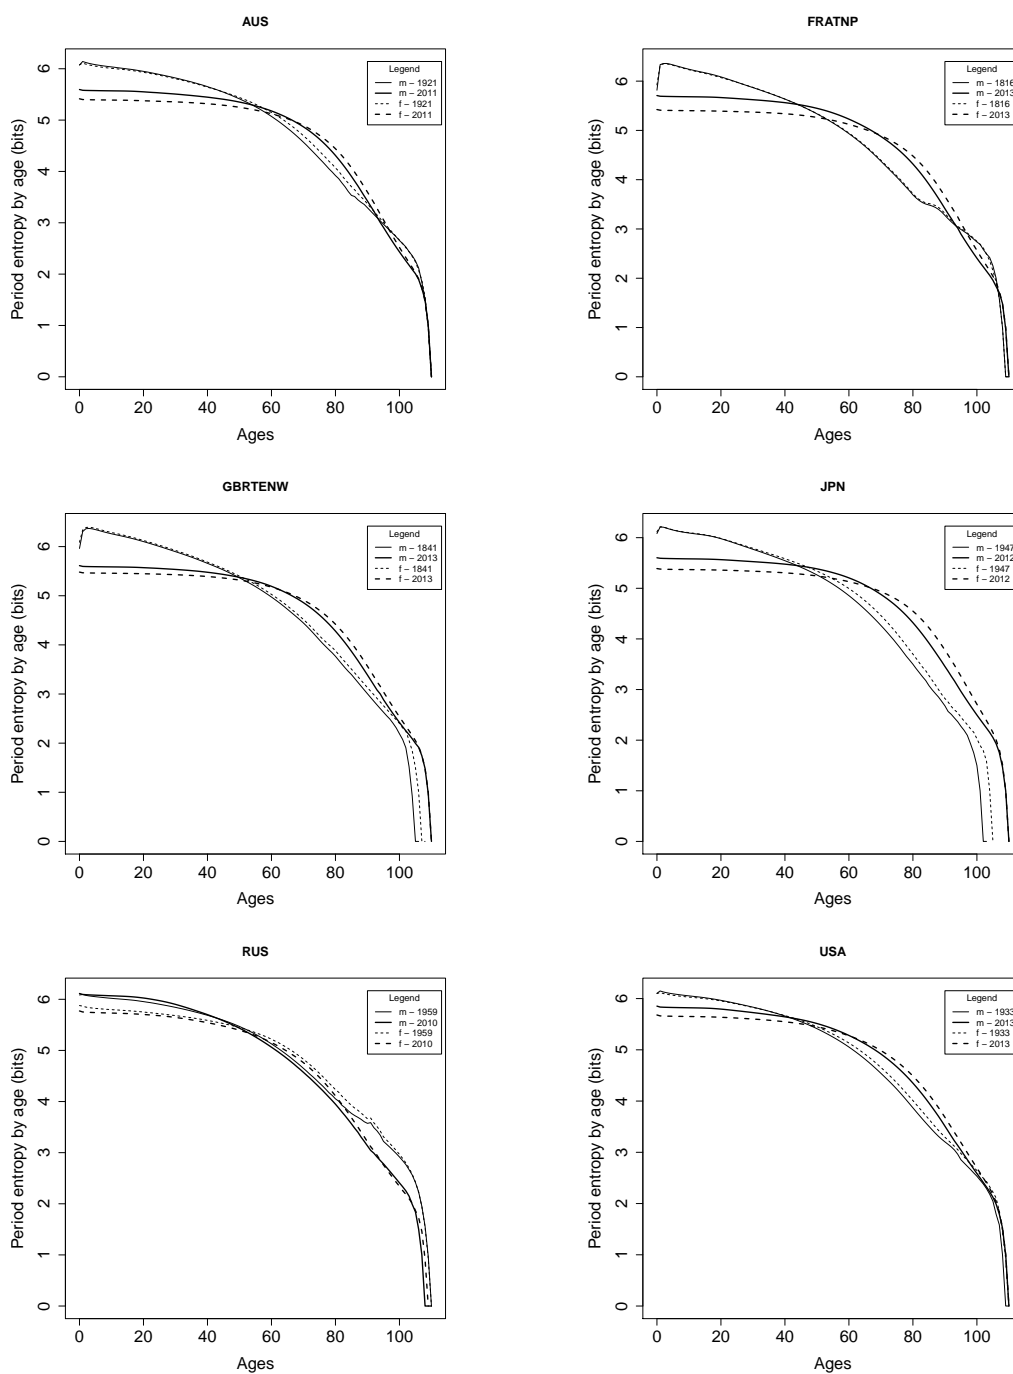


Figure 6: Period life entropy by age in Australia (AUS), France (FRATNP), England and Wales (GBRTENW), Japan (JPN), Russia (RUS) and the US (USA) for women (f) and men (m), first and last available year.

exhibit, in the first half of life, a higher life entropy than women. Then, in the second half of life, this is the opposite, and women exhibit a higher life entropy than men.

5 Cohort life entropy

Up to now, our analysis of human lifetime entropy has focused only on period data, based on mortality rates observed at year t (period lifetables). As a consequence, our entropy figures for a given year did not capture the actual risk about the duration of life, but the risk about the duration of life on the basis of the assumption that currently observed mortality rates will keep on prevailing during the entire lifetime of the individual.

As a complement, this section proposes to study lifetime entropy on a cohort basis (i.e. based on cohort lifetables). Cohort lifetime entropy at birth reflects the actual extent of risk about the duration of life (i.e. conditionally on the actual distribution of the age at death for cohort members), rather than the expected extent of risk about the duration of life (i.e. conditionally on the expected distribution of the age at death for cohort members).

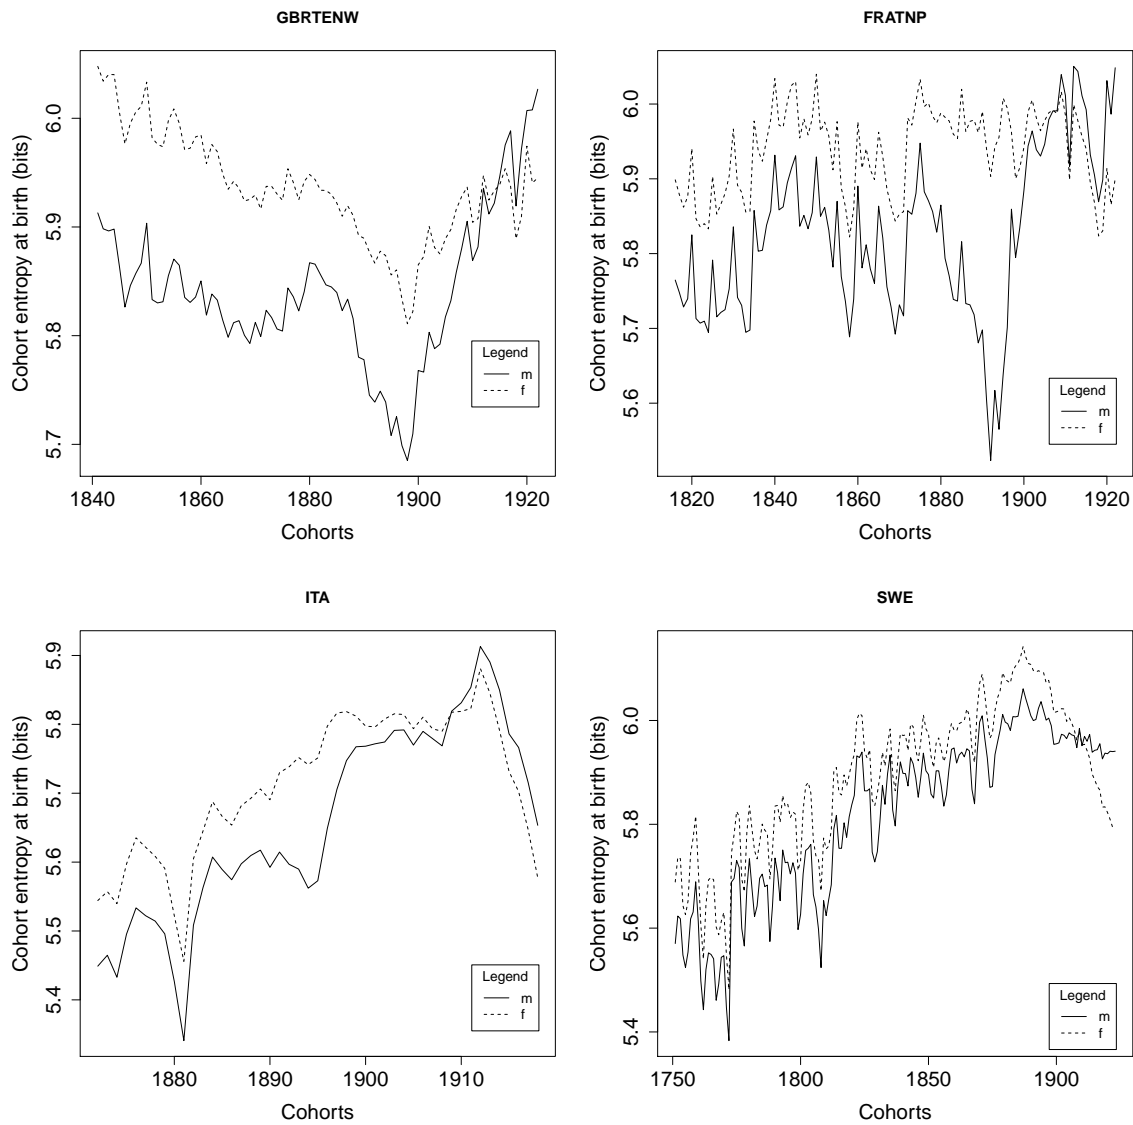


Figure 7: Life entropy at birth (cohort) in England and Wales, France, Italy and Sweden, for women (f) and men (m)

5.1 Cohort lifetime entropy at birth in a historical perspective

Figure 7 shows cohort life entropy at birth for men and women in England and Wales, France, Italy and Sweden. A first observation concerns the general pattern of cohort life entropy at birth for cohorts born during the 19th century. That pattern varies strongly across countries. Whereas cohort life entropy at birth tends to increase strongly in Sweden until cohort 1880 and then declines, the trend in cohort life entropy at birth remains stagnant in France during the 19th century (with strong fluctuations across the trend), and is decreasing in England and Wales until cohort 1895.⁸

Another observation concerns the comparison of cohort life entropy at birth across gender. Actually, Figure 7 shows that, although women exhibit a higher cohort entropy than men for all cohorts born in the 19th century, the two curves intersect for the cohorts born in late 19th century, and men exhibit a higher cohort lifetime entropy at birth for last few cohorts under study, that is, those born in the early 20th century. This observation also holds for the other countries for which cohort data are available.

5.2 Cohort life entropy: convergence or divergence?

Is there a convergence or a divergence across countries in terms of cohort life entropy at birth? To address that issue, Figure 8 shows, for women and men, the evolution of cohort life entropy at birth over the cohorts 1855-1920, as well as the evolution of the standard deviation across countries.

Figure 8 shows that there has been an unambiguous convergence in terms of cohort life entropy at birth for cohorts born between 1855 and 1920. Note that this convergence of life entropy across countries is achieved not by variations across countries in terms of the growth rate of life entropy, but by the mere fact that, while some countries with a higher initial entropy have exhibited a decrease in life entropy, other countries, with lower initial entropy levels, have exhibited a rise in life entropy over the cohorts considered.

5.3 Period versus cohort life entropy at birth

Are cohort and period life entropy patterns similar? To answer that question, Figure 9 compares, for France, England and Wales, Iceland and Sweden, the cohort and period life entropies at birth.⁹

Figure 9 shows that the relation between period and cohort life entropies at birth varies strongly across periods and countries, reflecting the different timing of shifts in survival curves across time and space. For instance, we can see that, for the cohorts of men born in England and Wales between 1845 and 1900, period life entropy at birth has systematically exceeded cohort life entropy at birth. This means that the persons in those cohorts faced, during their life, a lower life entropy than the one that could have been expected based on the yearly age-specific probabilities of death. On the contrary, for cohorts of men from Iceland born during the same period (i.e. between 1845 and 1900), cohort life entropy is not systematically below period life entropy.

The comparison of period and cohort life entropies can also be used to better quantify the impact of World War One (1914-1918) on the riskiness of life. Period life entropy at birth for French males fell strongly at the beginning of the war, making life entropy at birth drop from about 6 bits to 5.4 bits. Turning now to cohort analysis, Figure 9 shows that the impact of the War on the riskiness of life is not concentrated on just a few cohorts, but, rather, is spread on many cohorts, born between 1875 and 1895. Cohorts of French males born in the 4th quarter of the 19th century are characterized by entropy levels that fall strongly across cohorts, from 5.9 bits for men born around 1875 to about 5.5 bits for men born around 1895. Thus Figure 9 shows how different were the lives of successive cohorts born in just two decades, depending on their age at the beginning of the War.

⁸Note that this finding is not in contradiction with the results shown in Section 3 on the basis of period data. Actually, period data were studied on a much longer time period, during which a clear inverted U shape pattern could be identified. See Section 5.3 on the period/cohort comparison.

⁹A similar comparison is carried out for women on Figure 14 in the Appendix C.




Figure 8: Life entropy at birth (cohort) in Denmark, England and Wales, Finland, France, Iceland, Italy, Netherlands, Norway, Scotland, Sweden and Switzerland and associated standard deviation by year, for women (f) and men (m).

5.4 Cohort life entropy along the life cycle

Whereas the previous section focused on cohort life entropy at birth, it is also possible to examine how cohort life entropy has evolved with the age of cohort members, that is, how the risk about the duration of life is progressively resolved as one becomes older. Figure 10 shows cohort life entropy along the lifecycle for men and women in France, England and Wales, Iceland and Sweden, for the first and the last cohorts under study in each country.

On the basis of Figure 10, three main observations can be made. First, in each country, we can see that a major change in the form of entropy age profiles concerns the very beginning of life. For cohorts born in the 18th or 19th century, infant mortality was high, and being able to survive the first few

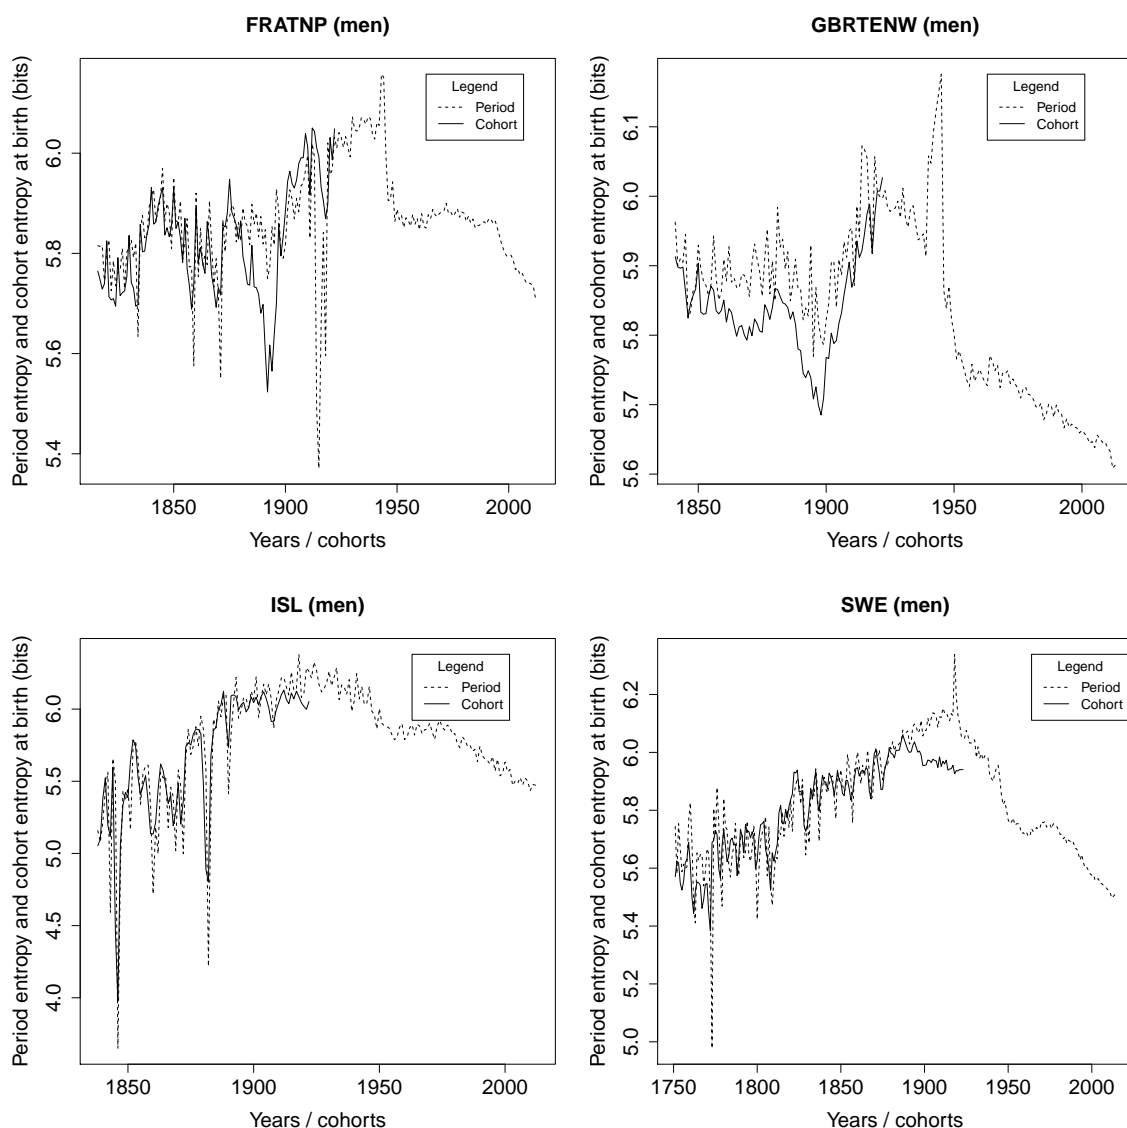


Figure 9: Period and cohort life entropy at birth in France (FRATNP), England and Wales (GBRTENW), Iceland (ISL) and Sweden (SWE) for men

years of life was opening new possibilities in terms of life duration, and, hence, increasing life entropy significantly. This rise in life entropy at low ages is no longer observed among cohorts born in the 20th century (or, when observed, it leads to a much lower rise in life entropy).

A second observation concerns the overall shape of the entropy age profiles by cohorts between ages 40 and 100. For those ages, cohort entropy age profiles have tended to shift upwards over time. The overall shape of the entropy age profile between ages 40 and 100 looks pretty much the same across cohorts, but the level of cohort entropy has changed. We can see that, above age 40, cohorts born in the early 20th century faced a higher life entropy at all ages in comparison to cohorts born in the 18th or 19th century. The gap varies with the age, but is equal on average to about 0.5 bits. Thus the improvement of survival conditions has led to a rise in the riskiness of life after age 40.

A third observation concerns the very high ages of life (i.e. beyond age 100). For cohorts born

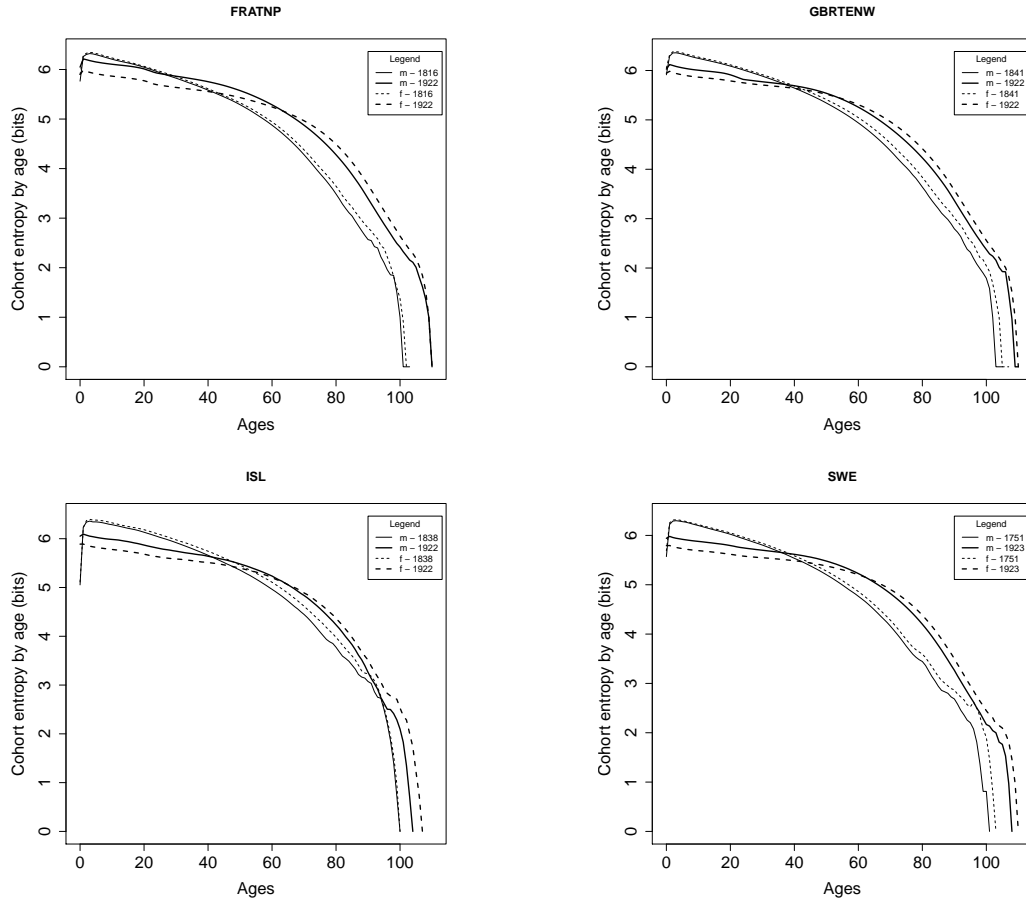


Figure 10: Cohort life entropy by age in France (FRATNP), England and Wales (GBRTENW), Iceland (ISL) and Sweden (SWE) for women (f) and men (m), first and last available years.

in the 18th or 19th century, entropy beyond age 100 was pretty low, and falling strongly to zero bits (risk fully resolved).¹⁰ However, for cohorts born in the early 20th century, entropy at age 100 remains equal to about 2 bits. Thus, for high ages, the entropy age profile has tended to shift to the right, allowing cohorts to face a significant risk about the duration of their life at ages when, in the past, little risk remained. This major change can be explained by the shift of survival curves to the right (i.e. the rise in limit longevity), which opened new perspectives at very high ages, and, also increased risk about the duration of life at those ages.

6 Conclusions

Life is inherently risky: no one knows how long one will live. In this paper, we proposed to develop a simple, intuitive measure of the quantity of risk about the duration of life: Shannon's entropy index to the base 2. The main intuitive support for the use of that index comes from the associated metric.

¹⁰This can be explained by the fact that cohort lifetables for those cohorts exhibited probabilities to reach ages beyond 95 years that were often close to zero or equal to zero when the lifetable was closed at a maximum age equal to 95. Note that the fact that lifetables were closed at lower ages is not a pure measurement artifact: it reflects that surviving beyond those ages was, at those distant epochs, extremely rare.

Indeed, this index quantifies the degree of life riskiness in terms of bits, that is, in terms of the amount of information revealed by tossing a fair coin.

We identified 8 stylized facts. The first four stylized facts were obtained by considering how life entropy at birth has evolved over the last three centuries in the countries under study. We showed that: (1) over the last two centuries, (period) life entropy at birth exhibits an inverted U shape pattern with a maximum in the first half of the 20th century (at about 6 bits); (2) the variability of (period) life entropy at birth has declined over time; (3) curves of period life entropy at birth for men and women crossed during the 20th century; (4) over the last 150 years, Western countries have converged in terms of (period) life entropy at birth towards 5.6 bits for men and 5.5 bits for women.

Those stylized facts suggest that, when interpreting the evolution of survival conditions, focusing only on the life expectancy indicator may leave aside a big part of the picture. True, over the last two centuries, life expectancy at birth has more than doubled in the countries under study. However, our calculations suggest that, over the same period, the quantity of risk associated to the duration of life has first increased and then decreased, to reach about 5.5 bits. This change, although significant, still leaves today's life entropy at a level that is quite high. Hence the main message is that, although survival prospects have been strongly improved over the last centuries, life remains quite risky, in the sense that the death of someone still reveals nowadays a large amount of information.

Four other major stylized facts were obtained by computing life entropy at all ages of life. We showed that: (5) the entropy age profile has shifted from a non-monotonic profile (in the 18th and 19th centuries) to a strictly decreasing profile (in the 20th and 21st centuries); (6) men exhibit higher life entropy than women below ages 50-55, and exhibit lower life entropy than women after ages 50-55; (7) the entropy age profile can be approximated by a piecewise linear function with three distinct segments, coinciding with a respectively, low, medium and high speed of learning about one's duration of life; (8) comparing cohorts born in the 19th and 20th centuries, the entropy age profile has tended to shift upwards (with a gain of 0.5 bits at each age during the active life) and to the right (with a gain of about 2 bits beyond age 100).

Those findings point to major changes in the speed at which one learns about the duration of one's life as one gets older. Before the major advances against infant mortality, the quantity of risk associated to a life was growing with the age at low ages, and then decreasing with the age. This is no longer the case nowadays, where a person, when becoming older and older, necessarily learns more about his/her duration of life. Another interesting interpretation of those findings is that, although the entropy age profile has changed over time and across cohorts, it always keeps a kind of piecewise linear form, which suggests that there exists, from the perspective of learning about one's age at death, a robust division of life into three ages of life.

All in all, using the Shannon entropy index to the base 2 allows us to cast some original light on the evolution of survival conditions over time, by allowing us to summarize in a single number the degree of risk about the duration of life. Moreover, in comparison with other measures of life riskiness, the Shannon entropy index to the base 2 has the advantage of relying on a metric that has a clear, intuitive counterpart for the laymen used to toss fair coins.

This leads us to a final suggestion: when presenting life expectancy statistics in reports or newspapers, providing, for instance in parentheses, the associated amount of bits could bring an important additional information. Allowing humans to have a more concrete idea of what the lottery of life is could definitely help them in making (more) rational decisions in the context of a risky lifetime.

7 References

Brown, J., Coe, N., Finkelstein, A. (2007). Medicaid crowd-out of private long-term care insurance demand: evidence from the Health and Retirement Survey. NBER Chapters, in: *Tax Policy and the Economy*, volume 21, 1-34, National Bureau of Economic Research.

Brown, J., Finkelstein, A. (2007). Why is the market for long-term care insurance so small? *Journal of Public Economics*, 91 (10), 1967-1991.

- Brown, J. and Finkelstein, A. (2009). The private market for long term care in the U.S.. A review of the evidence. *Journal of Risk and Insurance*, 76(1): 5-29.
- Brown, J., Finkelstein, A. (2011). Insuring long term care in the United States. *Journal of Economic Perspectives*, 25 (4), 119-142.
- Cheung, S.L., Robine, J.M., Jow-Ching Tu, E., Caselli, G. (2005). Three dimensions of the survival curve: horizontalization, verticalization and longevity extension. *Demography*, 42 (2), 243-258.
- Hill, G. (1993). The entropy of the survival curve: an alternative measure, *Canadian Studies in Population*, 20 (1), 43-57.
- Human Mortality Database. University of California, Berkeley (USA), and Max Planck Institute for Demographic Research (Germany). Available at www.mortality.org or www.humanmortality.de
- Jankélévitch, W. (1977). *La mort*. Flammarion, Paris.
- Kannisto, V. (2000). Measuring the compression of mortality. *Demographic Research*, 3 (6).
- Kopczuk, W., Slemrod, J. (2005). Denial of death and economic behavior, *The B.E. Journal of Theoretical Economics - Advances*, 5 (1), 1-26.
- Nagnur, D. (1986). Rectangularization of the survival curve and entropy: the canadian experience, 1921-1981. *Canadian Studies in Population*, 13 (1), 83-102.
- Noyer, A., Coleman, C. (2013). A universal pattern of the evolution of life table entropy and life expectancy, mimeo.
- Pierce, J. (1980). *An introduction to information theory. Symbols, signals and noise*. Dover, London.
- Wilmoth, J., Horiuchi, S. (1999). Rectangularization revisited: variability of age at death within human populations. *Demography*, 36 (4), 475-495.

A Data ranges




Figure 11: Ranges of period lifetables (left) and cohort lifetables (right) per country in the Human Mortality Database

B Period life entropy at birth : long periods




Figure 12: Life entropy at birth (period) in Belgium (BEL), Switzerland (CHE), Denmark (DNK) and Finland (FIN), for women (f) and men (m)

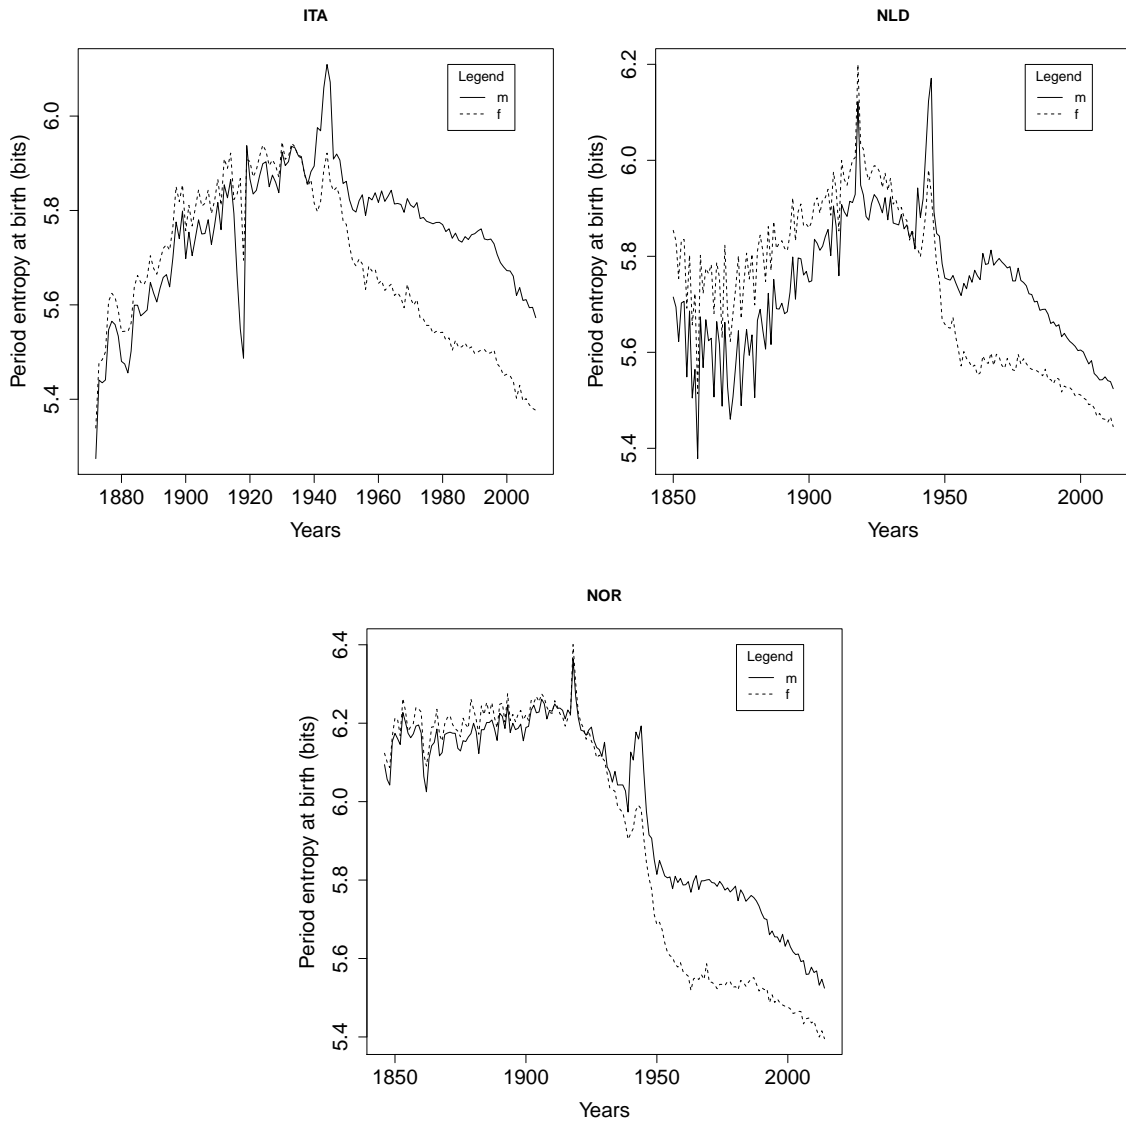


Figure 13: Life entropy at birth (period) in Italy (ITA), Netherlands (NLD) and Norway (NOR), for women (f) and men (m)

C Period vs cohort entropies at birth for women




Figure 14: Period and cohort life entropy at birth in France (FRATNP), England and Wales (GBRTENW), Iceland (ISL) and Sweden (SWE) for women

D Period vs cohort entropies by age for men

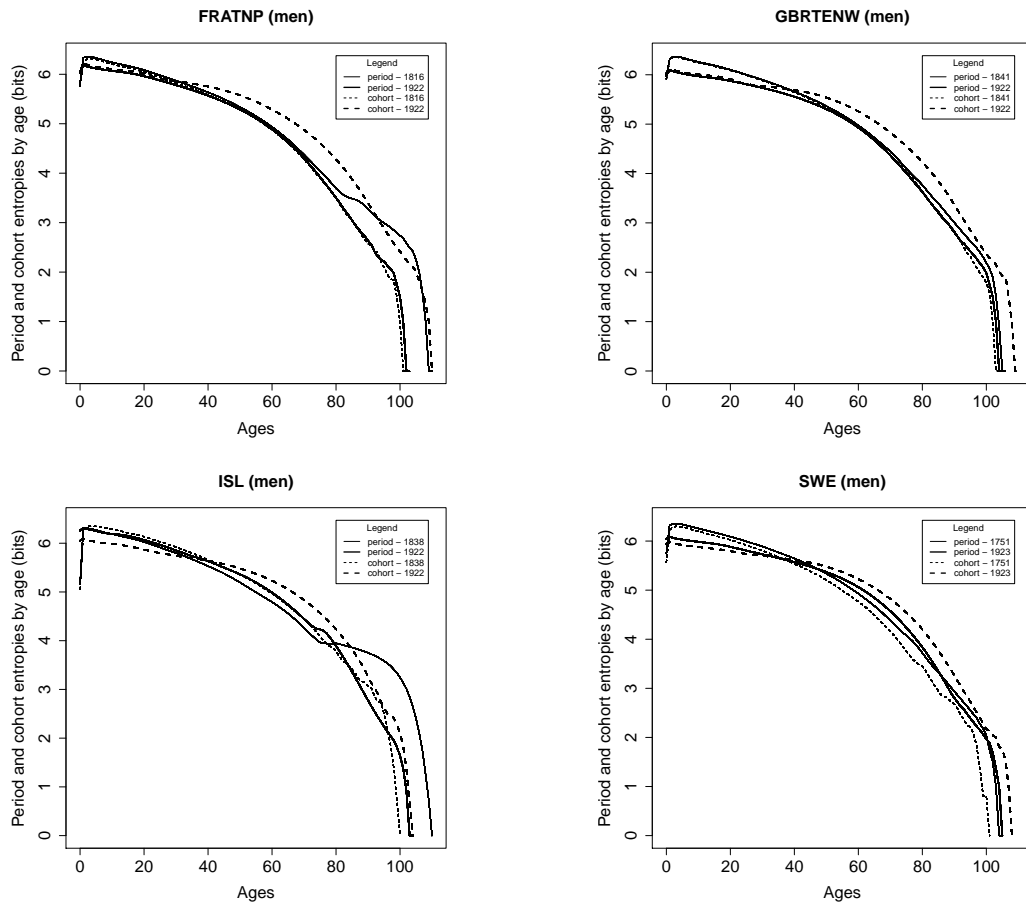


Figure 15: Cohort vs period life entropy along the life cycle in France (FRATNP), England and Wales (GBRTENW), Iceland (ISL) and Sweden (SWE) for men, first and last available years.