

PAPER

Mortality in the US by education level

Cristian Redondo Lourés and Andrew J. G. Cairns*

Maaxwell Institute for Mathematical Sciences and Department of Actuarial Mathematics and Statistics, School of Mathematical and Computer Sciences, Heriot-Watt University, Edinburgh EH14 4AS, UK

*Corresponding author. Tel: 0131 451 3245. Email: A.J.G.Cairns@hw.ac.uk
(www.macs.hw.ac.uk/~andrewc/ARCSources/)

(Received 25 March 2020; accepted 26 March 2020; first published online 28 April 2020)

Abstract

Different mortality rates for different socio-economic groups within a population have been consistently reported throughout the years. In this study, we aim to exploit data from multiple public sources, including highly detailed cause-of-death data from the United States Centers for Disease Control and Prevention, to explore the mortality gap between the better and worse off in the US during the period 1989–2015, using education as a proxy.

Keywords: US mortality; Education level; Socio-economic status; Mortality inequalities

1. Introduction

It is a well established fact that different socio-economic groups within a single population experience different levels of mortality, with the lower socio-economic groups experiencing consistently higher mortality rates than their better off counterparts. Education has been used as a proxy for socio-economic status in several mortality studies in the past because information on educational attainment of both decedents and the general population is more readily available than other socio-economic indicators (e.g. income, wealth, occupation). In particular, differences in mortality between different education categories in the US have been reported in the past, see for example, Jemal *et al.* (2008), Olshansky *et al.* (2012), or Case & Deaton (2015).

There are two reasons why we choose to analyse mortality by education level in the US. First, the size of the population is such that, even when we focus on relatively rare causes of death (CoDs), absolute death counts will be high enough that reliable rates can be computed separately for single years of age and calendar years, for males and females, and different education groups. The second reason is that the data needed for this kind of analysis is freely available online, and both a comprehensive list of all deaths in the US and certain characteristics of the resident population can be downloaded from the web pages of official US agencies.

To the best of our knowledge, there are no other studies that analyse such a complete and detailed range of CoDs as the one we can construct with this data. This will help uncover which CoDs have influenced most heavily the recent trends in mortality for different education groups and which are the key drivers of the observed gap in all-cause mortality. Having a comprehensive list of CoDs, each of them potentially caused by very different underlying factors (smoking, poor diet, genetic makeup), will provide important insight on why mortality differences arise and what could drive their future evolution.

We also propose a method to smooth out the noisy estimates of educational attainment in the US population obtained from survey data, as long as reliable estimates of whole population and

death numbers are provided. This method is fairly general and easy to extend to other types of data beyond education, which would help obtain underlying population estimates for any characteristic as long as survey data and death counts are provided. It also allows us to extrapolate subpopulation data to ages beyond the maximum provided by the survey, again conditional on knowing the death counts and total exposures. It is thanks to this smoothing of the sampling error of survey data that we are able to accurately report death rates by individual year of age and calendar year for different education groups, and not only average quantities such as standardised death rates and life expectancies.

In this paper, we will discuss how we construct the full dataset, which will be made publicly available, consisting of deaths and exposures by single year of age, single calendar year, and different education groups for males and females. We will also discuss the resulting empirical death rates, both for all-cause mortality and selected CoDs. The outline of the present work is as follows: section 2 introduces the data sources and explains what they provide us with, as well as some of their shortcomings. Section 3 deals with how we can correct some of the issues that appear when separating our death counts and exposures by education level, explaining the procedure followed to build our dataset. Section 4 contains the resulting estimates for the proportion of people in each subpopulation and uses a series of tests to prove the strength of our results. The extrapolation procedure is explained in section 5, in which we also describe the final dataset characteristics. Section 6 contains the resulting death rates for different CoDs and a discussion on their importance. Section 7 summarises our findings and discusses the potential for further research.

2. Data

2.1 Death counts

Individual deaths data come from the National Vital Statistics System, a program managed by the National Centre for Health Statistics (NCHS) – part of the US Centers for Disease Control and Prevention (CDC). The NCHS collects vital data that has been acquired by the state authorities and makes it publicly available through their web page, see NCHS (2016). Published data are transcribed directly from individual death certificates. These are anonymised to comply with data protection laws and published annually. The death certificate is partly filled in by the funeral director, see Rostron *et al.* (2010). This can cause problems with the reliability of certain items and will be discussed in more detail later. Although death certificates might not provide an exact count of all deaths it can be assumed that it is accurate for our purposes. The NCHS estimates that the coverage is above 99%, see US Department of Health and Human Services (1995).

The information that we are interested in is gender, age, calendar year, educational attainment and CoD. The first three are quite straightforward and are recorded in almost every death (e.g. only a negligible number, smaller than 0.02% over the whole period analysed, lacks information on age). We do not have any reasons to doubt the accuracy of any of these elements. The recording of education attainment on death certificates is less straightforward. This item started being recorded only in 1989, and even then not all states introduced it simultaneously, resulting in many entries with missing education in the earlier years. The death certificate was changed again in 2003, including a revised classification of educational attainment. Again, some states adopted this new certificate immediately, while some others used the old certificate for the whole period analysed. Due to this, we have two different ways in which education is recorded that coexist during the 2003–2015 period. Some background on the treatment of educational attainment in death certificates can be found in Hoyert *et al.* (2006). We explain later how this is handled.

Regarding the quality of the data, there are several issues that are well known in the literature. For example, as already mentioned, some parts of the death certificate have to be filled by the funeral director. These might not have personally known the deceased or have otherwise access to all the information required to fill all fields completely. This is especially true for the educational attainment item, which is known to be biased towards the “finished high school” level, that is, people who attended (but never graduated) high school, as well as people who attended college

but never got a degree, tend to be recorded as high school graduates who never attended university (see Rostron *et al.*, 2010). Some other times the funeral director, unable to find what the educational attainment of a person was, reports it as unknown, although this is relatively rare (the certificates with missing education are only about 3% in the last years included in our analysis).

However, many other entries are missing education because different states started recording it at different points in time. Entries with unknown education make almost 30% of the deaths in the earlier years, slowly declining to slightly below 10% in 1996. Fortunately, some states adopted the 1989 standard death certificate immediately and have consistently reported education in more than 90% of their death certificates every year since 1989. These can be used to impute an education to all entries missing it using the ideas from Sasson (2016). In a later section, we will explain how we implement this in our case.

CoDs are recorded using the International Classification of Diseases (ICD) codes. Up to 1998, the ICD-9 standard was used, with later years using ICD-10. This change affects not only the way certain illnesses are coded, but also how the underlying CoD is determined. After the chain of events that led to death by a practitioner has been listed, an automated system determines the underlying CoD following some certain rules. These rules change with each revision of the ICD, which means that discontinuities are introduced in the temporal trends of CoD mortality. The effect of this can be quite big for certain CoDs. As is common practice when ICD rules change, deaths for a certain year were classified using both standards, thus allowing the estimation of some comparability ratios (published in Anderson *et al.*, 2001, which also contains more detail on how the coding and determination of underlying cause work). With this, we can estimate the size of the discontinuity in the number of deaths attributed to certain causes when the ICD9–ICD10 change was implemented. In principle, applying these comparability ratios could make the death rates evolve smoothly, but changes in the trend can still occur that are caused by the change in practice rather than an actual change in CoDs. If we are, however, interested in the ratios between the mortalities experienced by different subpopulations, these would be unaffected by all these changes. In this paper, we deal exclusively with the construction of the database and the crude death rates produced, so correcting these discontinuities is not crucial. We only need to be aware of their potential impact when interpreting the temporal trends observed for a specific CoD.

2.2 Exposures

In line with the deaths data, we seek to estimate exposures (average population) by education level as well as gender, year and age. To achieve this, we exploit data from three sources: the Human Mortality Database (HMD), Current Population Survey (CPS) and American Community Survey (ACS).

- Exposures without subdivision by education have been extracted from the HMD. These have been then filtered using the methods proposed by Cairns *et al.* (2016) to adjust for potential anomalies in the data.
- For each group by gender, year and age, we need estimates of the proportion of the group that have attained different levels of education. This comes from two potential sources: the CPS and the ACS. The CPS data cover the full period 1989–2015, providing data by individual year of age up to 79 years. The ACS data surveys a larger number of people, but only covers 2004 to the present. As a consequence, we use the CPS data for our main calculations and use the larger ACS dataset for independent verification of aspects of our results.

The CPS is a monthly survey conducted by the Census Bureau. In their own words (US Census Bureau, 2000):

The CPS is administered by the Census Bureau using a probability selected sample of about 60,000 occupied households. The fieldwork is conducted during the calendar week that includes the 19th of the month. The questions refer to activities during the prior week; that is,

the week that includes the 12th of the month. Households from all 50 states and the District of Columbia are in the survey for 4 consecutive months, out for 8, and then return for another 4 months before leaving the sample permanently. This design ensures a high degree of continuity from one month to the next (as well as over the year). The 4-8-4 sampling scheme has the added benefit of allowing the constant replenishment of the sample without excessive burden to respondents.

To be eligible to participate in the CPS, individuals must be 15 years of age or over and not in the Armed Forces. People in institutions, such as prisons, long-term care hospitals, and nursing homes are ineligible to be interviewed in the CPS. [. . .] No upper age limit is used, and full-time students are treated the same as nonstudents. One person generally responds for all eligible members of the household. The person who responds is called the ‘reference person’ and usually is the person who either owns or rents the housing unit. If the reference person is not knowledgeable about the employment status of the others in the household, attempts are made to contact those individuals directly.

The CPS is designed to be a survey about the characteristics of the workforce, and therefore the sample is designed to obtain an error for the unemployment rate below a certain threshold. Once the participating households have been selected, their occupants are surveyed for four straight months and then dropped for 8 months, before returning to the sample exactly 1 year after their initial inclusion. In this second round, people are interviewed for another 4 months and then dropped permanently. Each month, 1/8 of the households are dropped permanently and another 1/8 enters the “freezing” period, with 1/8 of the sample for the following month being comprised of completely new households entering the survey, and another 1/8 of “defrozen” households. This way there is a 75% overlap between the households surveyed in two consecutive calendar months and a 50% overlap between the same month of two consecutive calendar years. More details on sample selection and rotation can be found in Chapter 3 of US Census Bureau (2006), particularly in pages 1 and 10–14.

The rotation pattern means that the samples for the same calendar month of two consecutive calendar years are not independent. Half of the households interviewed in a certain year will be the same ones that were interviewed in the previous one, which links the two data points. This is done in order to smooth the yearly drift of the quantities of interest, but could be a problem for us if we want to treat the different points in our time series as independent random variables. However, because we are not interested in the total population estimates given by the CPS, but only in the ratio of respondents that reported to have a certain education, we can restrict the data we retrieve to be the one reported by households that had not taken part in the survey the previous year. We do this at the expense of losing approximately half of the counts for each of our groups, which will increase the sampling error by a $\sqrt{2}$ factor.

The way the sample is designed still makes the not-in-sample-previous-year subsample of households representative of the whole population. Details on sampling can be found, as explained before, in Chapter 3 of US Census Bureau (2006), but the basic idea is as follows. The whole territory of the US is divided into small, relatively homogeneous primary sampling units (PSUs). These are grouped together in strata which are expected to have similar workforce characteristics, and one PSU is selected to be surveyed, with its results representing all other similar PSUs. In each PSU to be surveyed, a list of addresses is taken from the latest Census and complemented with the addresses of recently built housing obtained from building permits. From these, a sample of addresses is chosen to be surveyed and split in several subsamples to be used in the rotation scheme. The way the rotations groups are set up is such that all subsamples are balanced at the levels that we are interested in, which are the state and national level.

Figure 1 shows how this rotation works. Each row represents a sample of the US population. In a given month, for example, July 2012, 16 samples (labelled S1–S16 here) are part of the CPS sample. Four of them (S1–S4) are being interviewed that month, but they were also interviewed




Figure 1. Plot showing how the 28 samples (labelled S1–S28) leave and enter the rotation group during a calendar year. Green: data are collected and not correlated with the one for the same month 1 year before. Orange: data are collected but correlated with the one collected the previous year and are discarded. Grey: the subsample is in the waiting period. White: the subsample is not in sample (NIS).

exactly 1 year ago. Eight others (S5–S12) are in the 8-month waiting period and are not interviewed. Finally, the last group of four (S13–S16) is being interviewed in July 2012, but was not in July 2011. We can see that as months pass, some of the subsamples are being dropped and others take their place, and that in exactly 1 year’s time (July 2013) the group S13–S16 is now being interviewed for the second time in a month of July, while the group S25–S28 is interviewed in July 2013 but was not in July 2012. Choosing only the uncorrelated data is equivalent to discarding the data coming from the orange part of the plot.

3. Educational Attainment

The first thing we need to study educational attainment is a clear definition of the different levels that we are going to use. We need to find education groups that will produce informative results and that can be easily and consistently defined with the available data. Recent studies of mortality in the US that use education as a covariate (including Case & Deaton, 2015) classify people in three groups: people who never attended college (whether or not they graduated high school), people who attended university but never got a bachelor’s degree and people who obtained a bachelor’s or higher qualification. We aim to use a similar grouping in our work. However, our data sources do not provide information on educational attainment in a way that can be straightforwardly converted into the groups defined above, so care must be taken to allocate our death counts and exposures into one of the three groups consistently. A degree-based classification is used in both

Table 1. Cross-tabulation of old code (years of schooling) and new code (highest grade completed or degree received).

Old code	New code			Total
	0–HS diploma	Some college/associate	Bachelor’s or higher	
0–12	39,219	2,832	147	42,198
13–15	508	13,683	350	14,541
16+	70	881	16,735	17,686
Total	39,797	17,396	17,232	74,425

Source: Park (1999) – abbreviated.

the CPS and the 2003 death certificate (which, as already stated, some states had not yet adopted by 2015), but in the earlier years another classification based on years of schooling was used. This forces us to try to make an equivalence between the two systems to have consistent data throughout all the years in our analysis.

For the exposures data coming from the CPS, we have in fact three different education reporting methods: between 1989 and 1991 a “years of schooling model” was used, in which the number of years spent in elementary education (or college if applicable) was recorded; from 1994 on there is a “grade obtained” approach, in which people are classified according to the highest degree they have earned (with a special “some college, no degree” category); and in 1992–1993 a hybrid approach is used, in which years of schooling are used for people who did not graduate high school, and highest degree obtained is used as in the most modern standard for people with higher education. However, in death certificates, there was no intermediate stage, with the 1989 model recording education as “years of schooling” and the 2003 revision recording “highest degree obtained.” As already mentioned, not all states adopted any of the two death certificate models immediately after they were introduced, which results in a high number of entries with missing education information in the early years after 1989, and a mixture of both standards in data for all years starting in 2003. Moreover, between 1992 and 2003, CPS population data and death counts data were recorded using completely different standards, which forces us to try and find a way to group these different education categories in a manner that is consistent throughout the years and regardless of the educational attainment recording standard used.

A complete, category-by-category basis correspondence between the two standards is extremely difficult, if not impossible, to create. However, it is possible to do some grouping that will roughly correspond to the three education groups we want to work with. In the month of February 1990, people responding to the CPS survey question about their educational attainment were asked to provide two answers, one using the old (years based) and one using the new (degree based) standard. The results obtained are shown in Table 1 adapted from Park (1999) to summarise their findings for the groups we are interested in only. We find that data recorded in both standards can be merged in a way in which the membership of each of the three groups of interest is fairly consistent:

- Lower educated group. The group of all people who have never been to college. It corresponds to categories 0 to 12 years of schooling in the old standard (row labels in the attached table), and to categories from no formal schooling to HS Diploma in the new standard (0–HS Diploma in the attached table column headings).
- Medium educated group. These are people who have been to college but have not achieved a bachelor’s degree. It corresponds to categories 13–15 years of schooling in the old standard and “Some College/Associate degree” in the new one.
- Higher educated group. This group is formed by people who hold a bachelor’s degree or higher. It is formed by people in the categories 16 or more years of schooling in the old standard, or in the categories bachelor’s or further in the new recording.

In principle, this grouping should solve most of the problems arising from the differences in education recording. The memberships of each group for February 1990 are as follows:

- Lower educated group: 42,198 (old standard) and 39,797 (new standard).
- Medium educated group: 14,541 (old) and 17,396 (new).
- Higher educated group: 17,686 (old) and 17,232 (new).

The main source of error is the misclassification of 2,538 (out of the 2,832) people who said they had been schooled for 12 years in the old standard (i.e. until the end of high school) but claimed to have some college education when asked the degree-based question (probably because they attended college for less than a year). For the lower educated group, having a big membership, absorbing these extra members does not make a very big difference, but the medium educated group death rates could potentially be affected. Specifically, death rates for group 2 might be underestimated (death counts use old standard → underestimation of number of deaths compared to exposures using the new standard) in the period 1992–2003. The consistency of the membership of group 3 make any misclassification issues less of a problem for people in that group. Also note that we only studied the equivalence of the two educational attainment recording methods for a very specific case (CPS survey, February 1990) and aggregated on age for people aged 25–64 years. The results found there might not be representative of the equivalences between groups of education for all the length of our study or for people outside that age range. Moreover, and even though our grouping should solve most of the problems coming from education being recorded in two different ways, any data quality problems (e.g. misreporting) still remain.

3.1 Education in the death records

Now that we have defined our education groups we can tackle the problem of missing information in the death counts. In general, missing education in a death certificate happens due to one of two reasons: individual entries for which education was not known (which is a small percentage every year), and state-wide non-reporting due to the late adoption of the 1989 death certificate by some states (which is the reason behind most “unknown education” entries in early years). For example, whereas in 2010 all states reported education in death certificates, and entries with missing data were less than 5% of the total, in 1989 (the first year in which educational attainment entered death certificates) only 21 states reported education with less than 10% missing data. We can nonetheless try to impute an educational attainment level to each entry without one using simply the death counts with education recorded and the exposures. We generalise the method explained in Sasson (2016), based on Bayes theorem:

$$p(e|C, X) = \frac{p(C|e, X)p(e|X)}{p(C|X)} \quad (1)$$

where e represents education levels, C is CoD and $X=\{\text{Year, age, gender}\}$ captures the characteristics of the underlying live population as well as the matching characteristics at the time of death for those who die.

As long as we can provide an estimate for all three quantities on the right-hand side, we will be able to calculate the probability that an entry corresponding to a person who died from cause C , with characteristics X , had attained education level e .

The denominator on the right-hand side is simple to compute. It is the aggregate mortality rate by CoD across all education groups and can be calculated using the death counts and total population data.

As the second term in the numerator, we have the probability that a person with characteristics X has a certain level of education. This would be calculated as the exposures for education level e divided by total exposures for each X . As a first approximation, and because there is no “unknown education” category in the CPS, we can use the (unsmoothed) survey data to give an

initial estimate of this quantity. Once we have death counts with complete education, we will proceed to the smoothing of the exposures described in the next subsection and obtain a better estimate for $p(e|X)$, which we will use for a second round of imputation. After this, we will recalculate our exposures and repeat the process until we reach convergence, and the difference in the estimated death rates between two consecutive rounds of imputation-smoothing becomes negligible.

The first term in the numerator can be problematic, since we need the mortality rates by education level. In later years, and because the number of entries with missing information is relatively small, we can use the whole dataset to give an initial estimate for this quantity and proceed to the imputation from there. But for the earlier years, which concentrate most of the entries with unknown education, the large proportion of missing data means that an estimate obtained using all of the death counts could be biased. In this case, we have to use a subset of all states that have always reported educational attainment in at least 90% of death certificates to get a first estimate of these death rates, restricting both our death counts and exposures to them. This criterion has been used by the NCHS themselves when reporting educational-related data, see NCHS (1993). From year 1997 on, the missing education information is below the 10% threshold for the whole of the US, and therefore the first estimate for $p(C|e, X)$ can be calculated using all the available mortality data. From 1989 until 1996, we use data from the following states: Arizona, California, Colorado, Delaware, Florida, Hawaii, Idaho, Illinois, Iowa, Kansas, Michigan, Minnesota, Missouri, Montana, New Hampshire, Oregon, South Carolina, Utah, Vermont, Wisconsin and Wyoming.

The entries without education are split evenly between all three education groups in the first iteration to initialise $p(C|e, X)$ for each X . This way we ensure that there will be an initial seed for all X and e for all CoDs. Otherwise, we could end up with deaths due to rare causes not being assigned to any group because there were no occurrences for which education was known. Comparing the value to which the death rates converge using different criteria for the initial assignment of these deaths shows that the effect of the specific choice is negligible. This is understandable since only in extremely rare cases the number of deaths in the three education groups will not be much bigger than in the unknown education group thanks to the 90% completeness criterion.

Note that, with this method, we are assuming that the deaths with missing education are equally distributed in e given certain X and C as the deaths for which education was recorded.

3.2 Education in the exposures

The raw exposures obtained from the CPS are rather noisy, and even total population estimates for individual ages vary quite widely for consecutive calendar years (as does the estimate of the number of people of different ages in a single calendar year). Figure 2 clearly illustrates this: the HMD gives a smooth evolution of the total population, whereas the CPS (and even the more accurate ACS) estimates are noisier. As a consequence, we decided to take an approach in which we use the more reliable estimates for the total exposures from the HMD and then separate them by education level using the ratios of people in each education group as extracted from the CPS.

Even when HMD total population estimates are used, the CPS noise is present in the ratios of people belonging to each education group. If death rates by single years of age and calendar time are to be computed, this noise needs to be smoothed out to obtain the underlying population ratios. Furthermore, the smoothing procedure we propose provides us with a natural way of extrapolating the education ratios beyond the age of 79 years, the highest age for which the CPS has consistently reported educational attainment data by single year of age in the period we analyse (this extrapolation procedure will be described in detail in section 5). The smoothing of the noisy CPS ratios of educated people is one of the core elements of this paper and will be explained in careful detail in the following.

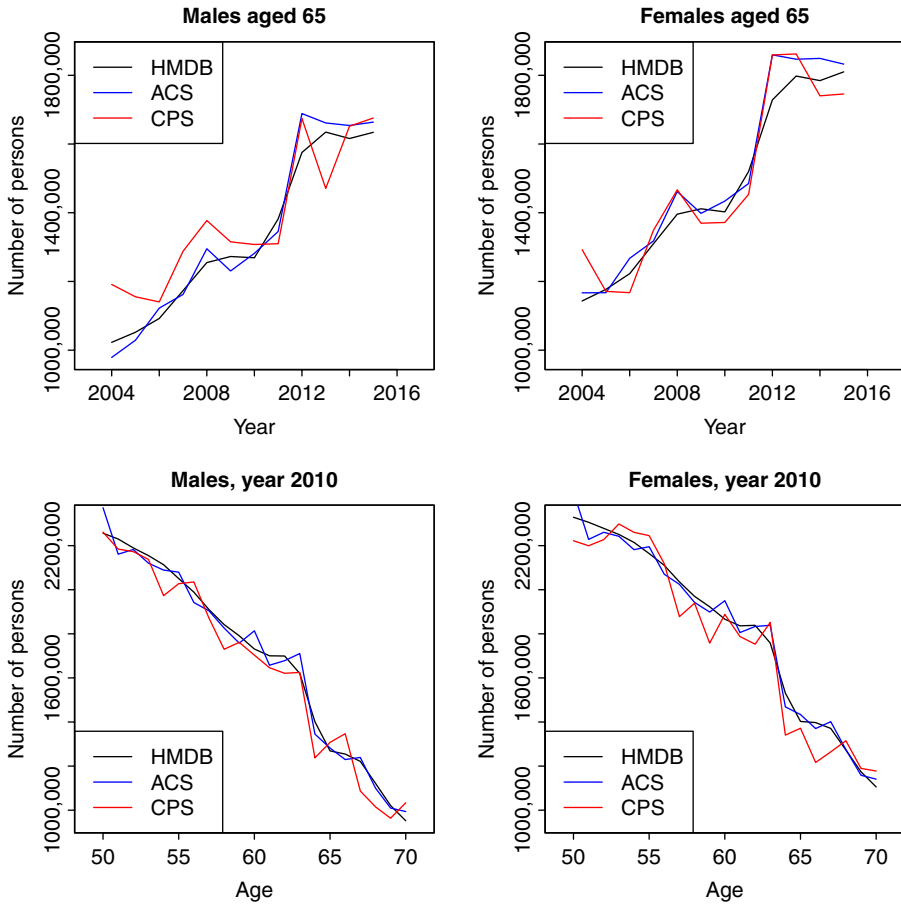


Figure 2. Comparison between the exposures coming from the CPS (red lines) and the ACS (blue lines) with the raw HMD exposures. The trends explained in the text are clearly visible in these examples, but they exist with more or less strength for most calendar years and ages.

The starting point will be calculating the evolution of the ratios of people in each education group for each of the cohorts that participate in the analysis. Crude estimates of these ratios based on raw CPS data will be quite erratic due to the sampling error. However, it is reasonable to expect that the true proportion of people with a certain education level *within a cohort* should evolve smoothly from year to year as the cohort ages.

We define the following quantities:

- Let c represent a cohort with a particular year of birth and gender.
- Let e represents the level of educational attainment.
- $E_C(i, e, c)$ = exposures for the cohort c with education level e , in year i of observation. $i = 0$ represents the first year that we observe this cohort. i then has a dual role as a proxy for both the calendar year and the age of the cohort.
- $E_C(i, c) = \sum_e E_C(i, e, c)$ = exposures for the cohort c across all levels of educational attainment, in year i of observation.

We now define

$$R_C(i, e, c) = \frac{E_C(i, e, c)}{E_C(i, c)}. \tag{2}$$

$R_C(i, e, c)$ is the true ratio of people in cohort c with education level e in year of observation i .

It is easy to derive the recurrence relation

$$\begin{aligned}
 R_C(i + 1, e, c) &= \frac{E_C(i + 1, e, c)}{E_C(i + 1, c)} \\
 &= \frac{E_C(i, e, c) - \Delta_C(i, e, c)}{E_C(i, c) - \Delta_C(i, c)} \\
 &= \frac{R_C(i, e, c)E_C(i, c) - \Delta_C(i, e, c)}{E_C(i, c) - \Delta_C(i, c)} \tag{3}
 \end{aligned}$$

where we have defined Δ_C as the decrease in the number of members of a certain cohort between years i and $i + 1$. There are three possible contributions to this Δ_C : the members of the group that died between years i and $i + 1$, net migratory flux, and people moving from a lower to a higher level of educational attainment (with the latter not affecting $\Delta_C(i, c)$, the decrease in the total cohort exposures). If we assume that at the ages analysed migration is small compared to the number of deaths, and the amount of people who choose to pursue further education is negligible, we can approximate Δ_C by the number of deaths:

$$\tilde{\Delta}_C(i, e, c) = D(t_0 + x_0 + i, x_0 + i, e, g) \tag{4}$$

where $D(t, x, e)$ is the number of deaths in calendar year t , age x and education level e , t_0 is the year of birth of the cohort, x_0 the age at which the cohort enters our analysis, and g is the gender (note that all these three parameters are collectively represented by c on the left-hand side). Additionally, we have

$$\tilde{\Delta}_C(i, c) = \sum_e D(t_0 + x_0 + i, x_0 + i, e, g). \tag{5}$$

This approximation will be accurate as long as the number of deaths is bigger than that of migrants or the education profile of these migrants is similar to that of the existing population. Although the foreign born population is found to be less educated on average than native citizens, see Ryan & Bauman (2016), the difference is relatively small. The result of compounding those two effects will increase the strength of our approximation. As we will see in the next section, statistical tests of the results do not seem to indicate that we have introduced any bias in our results by not including net migration in the analysis.

Once we have specified an initial estimate for $\hat{R}_C(0, e, c)$, we can calculate estimates of the remaining ratios, $R_C(i, e, c)$ for $i = 1, 2, \dots$ using the $\tilde{\Delta}_i$ (equations (4) and (5)) and the recurrence relation (3). This produces a full series $\hat{R}_C(i, e, c)$, for that particular cohort and education level. Alongside this, we also have the crude CPS ratios, $R_C^{CPS}(i, e, c)$.

To estimate this initial ratio, it seems reasonable to minimise the quadratic deviation between our smoothed ratios using recurrence relation (3) and the crude CPS ratios:

$$\mathcal{O}(e, c) = \sum_i \omega(i, e, c) \left(\hat{R}_C(i, e, c) - R_C^{CPS}(i, e, c) \right)^2 \tag{6}$$

where the $\omega(i, e, c)$ s are weights that we will discuss later on. Once we have chosen to use this model, there is one single parameter to be estimated for each cohort and education level: $\hat{R}_C(0, e, c)$.

The advantage of this method is that we make full use of all our data to estimate ratios that are compatible with the mortality observed. However, because we need the exact number of deaths for all education levels to estimate the ratios, and the ratios themselves are necessary for the education imputation discussed in the previous subsection, we need to proceed iteratively and repeat

the imputation-smoothing procedure several times until convergence of the final death rates is obtained.

In the next section, when the detailed results are shown, we will see that this cohort-by-cohort-based smoothing presents some issues. In particular, because we are smoothing the ratios for each cohort independently from its neighbours, we will see “ripples” in the estimated death rates (see Figure 6), as for some cohorts there will be a systematically over- or underestimation of the ratio of people within a certain education group compared to its immediate neighbours. These type of oscillating changes in the ratios of educated people between consecutive neighbours are highly undesirable and can potentially make genuine cohort effects more difficult to detect. Due to this, we will slightly modify the smoothing procedure so that we get rid of this problem making small adjustments to the ratios of educated people. Thanks to this we will obtain death rates that vary smoothly between consecutive cohorts.

In Cairns *et al.* (2016), one of the diagnostics of the quality of the exposures was concavity of the generated death rates. We will see that this is the most clear indicator that the ratios obtained by minimising function $\mathcal{O}(e, c)$ need to be corrected. The concavity function (suppressing reference to (e, g)) is defined as

$$C(t, x) = \log(m(t, x)) - \frac{1}{2} (\log(m(t, x + 1)) + \log(m(t, x - 1))) \tag{7}$$

where the $m(t, x)$ are empirical death rates. Assuming that $m(t, x)$ is smooth from age to age within year t , this should be close to zero. However, simply minimising the $\mathcal{O}(e, c)$ for each cohort in isolation can lead to oscillations in the $m(t, x)$ and larger values for the concavity than are reasonable. We should, therefore, modify our target function to include a term that penalises ratios that would generate a large concavity. This new target function is

$$\mathcal{O}_N = \sum_{i,e,c} \omega(i, e, c) \left[\left(\hat{R}_C(i, e, c) - R_C^{CPS}(i, e, c) \right)^2 + C(i, e, c)^2 \right]. \tag{8}$$

The new term, which gives a measure of the concavity, is

$$C(i, e, c) = \log(m(t, x, e, g)) - \frac{1}{2} \left[\log(m(t, x + 1, e, g)) + \log(m(t, x - 1, e, g)) \right] \tag{9}$$

where we define $t = t_0 + x_0 + i$ and $x = x_0 + i$, and the $m(t, x, e, g) = D(i, e, c) / \hat{R}_C(i, e, c) E_C(i, c)$ are the empirical death rates for cohort c (c encapsulating gender g and year of birth $(t - x)$) with education level e .

Given that we know the number of deaths D for a cohort and the total exposures E at every age, and that we know how to estimate the ratios $\hat{R}_C(i)$ given the initial ratio $\hat{R}_C(0)$, the only thing we need to give this estimate of the concavity is the log of the death rates of the neighbouring cohorts (those that at year t , when the cohort born at year t_0 is age x , are ages $x + 1$ and $x - 1$). Because we are already using an iterative method to estimate our death rates (remember we did this so that we could repeat the imputation-smoothing several times until convergence was achieved), we can use for each iteration the mortality rates calculated in the previous one. As long as the method converges (the change in the death rates between consecutive iterations becomes close to zero), the death rates of the previous iteration will help us compute the concavity that we would observe at the end of the current one. Thanks to this, optimisation of equation (8) would be useful to calculate ratios that generate exposures from which we derive death rates with a small concavity.

Robustness checks of the results obtained by minimising the target function \mathcal{O}_N were carried out to check if our specific choice of the penalisation affected the final results. In particular, changing the relative weights of the least square and the penalisation terms in equation (8), or using a different (even) power of the concavity for the penalisation, were found to have only a very minor effect on the resulting smoothed ratios.

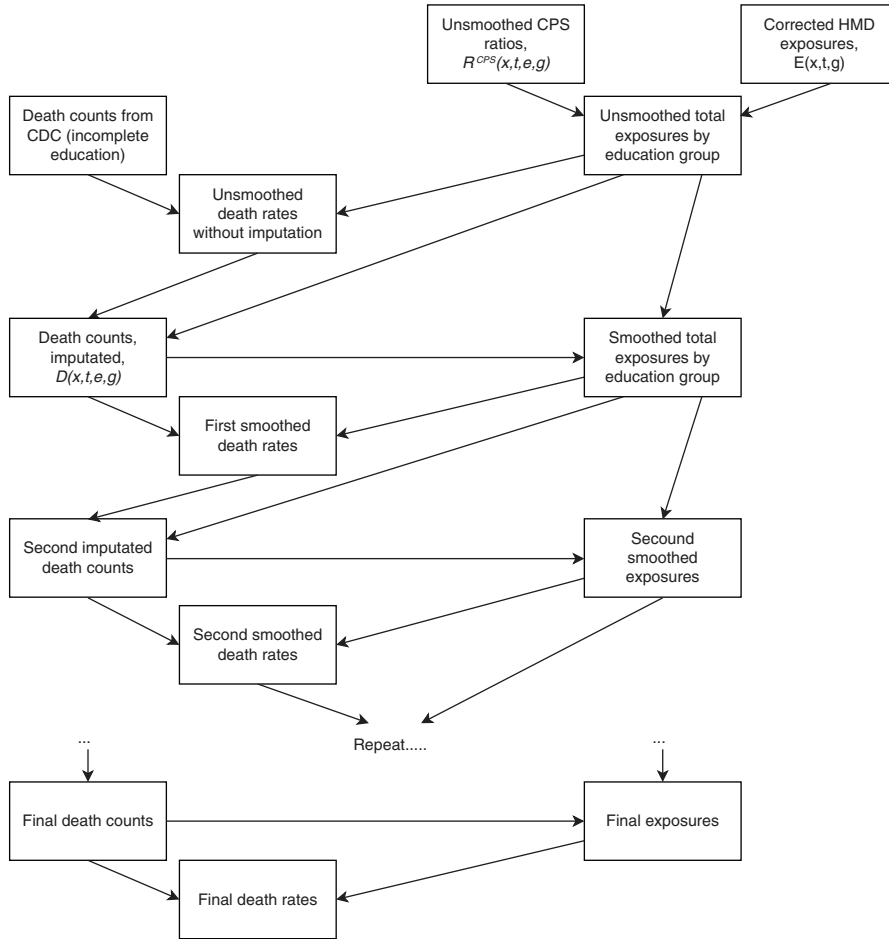


Figure 3. Flux diagram showing how the algorithm handles the initial data and obtains final death rates, death counts and ratios of educated people.

For reference, a diagram that graphically explains how our iterative algorithm works, transforming the input (HMD total exposures, noisy CPS ratios and death counts with missing information from CDC) into our results (smoothed exposures, death counts without missing education and death rates by CoD) is presented in Figure 3.

Lastly, it is worth mentioning that for both the unsmoothed ratios and the observed concavity, a standard error can be estimated. For the ratios we use, the number of counts in each education group follows a conditional binomial distribution with $c(t, x, g)$ trials (i.e. the number of persons in the sample with characteristics (t, x, g)) and probability $R(t, x, e, g)$ (the “true” ratio or proportion of the group with education level e) of success. The \hat{R} ratios we calculated before are an estimate of the true ratios R , and therefore their standard error, $\sigma(\hat{R}(t, x, e, g))$, is

$$\sigma(\hat{R}(t, x, e, g)) = \frac{\sqrt{c(t, x, g)\hat{R}(t, x, e, g)(1 - \hat{R}(t, x, e, g))}}{c(t, x, g)} = \sqrt{\frac{\hat{R}(t, x, e, g)(1 - \hat{R}(t, x, e, g))}{c(t, x, g)}}. \quad (10)$$

For the concavity, the fact that the total number of deaths observed is a random variable creates a certain variability of the death rates that will generate an uncertainty in the value of the concavity. If we assume that the number of deaths follows a log-normal distribution, then the logarithm of

the mortality would be normal. In particular, if we use the observed number of deaths D as the estimate of the mean of the distribution, and \sqrt{D} as its standard deviation, then the mortality $m = D/E$ would be log-normally distributed with mean $n = D/E$ and standard deviation $d = \sqrt{D}/E$. Through the relation between the parameters of the log-normal distribution and the distribution of its natural logarithm, we know that $\log(m)$ would be normally distributed with mean μ and standard deviation σ , where

$$\mu = \log \left(\frac{n}{\sqrt{1 + \frac{d^2}{n^2}}} \right) \tag{11}$$

$$\sigma = \sqrt{\log \left(1 + \frac{d^2}{n^2} \right)}. \tag{12}$$

Since at the end of the iterative process that imputes education and estimates ratios of educated people we get a total number of deaths and exposures for each group, we can estimate the standard deviation of $\log(m)$ for each calendar year and age. Because three log mortalities (that are, additionally, conditionally independent of each other) enter the calculation of the concavity, the conditional standard deviation of the concavity is

$$\sigma(C) = \sqrt{\text{Var}(\log(m(t, x))) + \frac{1}{4}\text{Var}(\log(m(t, x + 1))) + \frac{1}{4}\text{Var}(\log(m(t, x - 1)))}. \tag{13}$$

This allows us to add confidence intervals to the plots of the concavity so that we can see if the dispersion of the points for the different education groups can be attributed to random fluctuations in the number of deaths or to issues with the data. This should be interpreted as a lower bound of the error, since here we are assuming that the only variable generating the error is the number of deaths, that is, we assume that we know the exact exposures without any uncertainty. If we had added a certain error to E , the value of $\sigma(C)$ would be bigger than the one obtained here.

4. Smoothed Education Ratios and All-Cause Death Rates

This section will present the results obtained. Because the computation of the death rates and the ratios of people in each education group are intertwined in our model and neither can be calculated without the other, we will show the results for both the death rates and the ratios of educated people alongside each other.

Figure 4 shows, as an example, how the unsmoothed ratios look for certain cohorts. We can see the noisy nature of the data, with oscillations that are compatible with the sampling error, as shown by the dashed lines. The death rates extracted from this data will therefore be noisy as well, and thus the need to introduce the smoothing explained in the previous section.

In Figure 5, we show the smoothed version of these same ratios when the \mathcal{O} -smoothing and \mathcal{O}_N -smoothing procedures are used (with all weights ω_i being equal). We see that we obtain a smooth curve that falls really close to the centre of the 2σ intervals, and that the difference in the ratios due to the concavity penalisation is rather small. However, this small correction makes the smoothed death rates behave much better.

By way of example, resulting all-cause death rates for medium-educated are shown in Figure 6. We can clearly see the effect of each step of the smoothing procedure: when we go from the raw to the \mathcal{O} -smoothed exposures, we see that the surface becomes smoother, but with some ripples (or ridges) due to the independent adjustment of neighbouring cohorts; when going from \mathcal{O} - to \mathcal{O}_N -smoothed ratios anomalous cohorts that give rise to these ripples are brought closer to their neighbours, resulting in a much smoother and more plausible mortality surface.

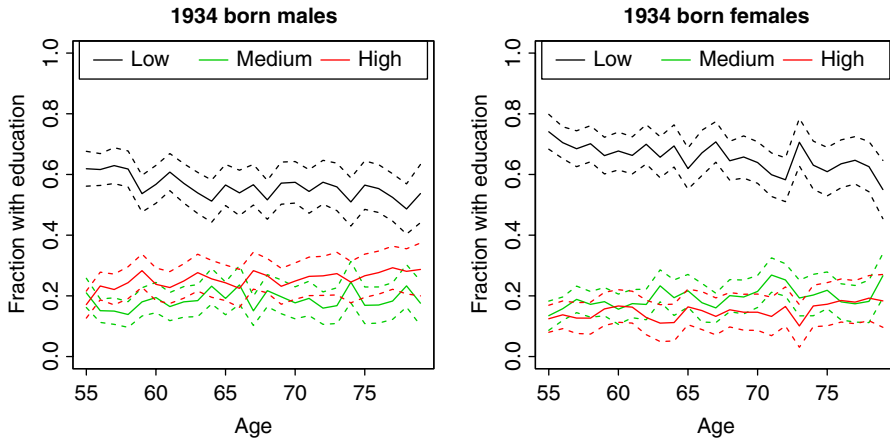


Figure 4. Unsmoothed ratios using the raw CPS data for male and female cohorts born in 1934, for the three groups low, medium and high education. The ratios (solid lines) and the $\pm 2\sigma$ confidence intervals (dashed lines) are shown.




Figure 5. \mathcal{O} -smoothed (solid lines) and \mathcal{O}_N -smoothed (dashed lines) ratios for the same cohorts as in Figure 4. The $\pm 2\sigma$ confidence intervals around the unsmoothed ratios from Figure 4 are shown as dot-dashed lines.

We now proceed to analyse the results by applying several diagnostic tools to our results. We will start with the concavity function. Figure 7 shows heat maps for the concavity function of medium educated males. The result is exactly as we would expect: for the unsmoothed ratios (Figure 7(a)), the concavity is fairly random and large (top left, note the different scale), for the \mathcal{O} -smoothed ratios (Figure 7(b)) we have a smaller concavity but patterns along some cohort diagonals, while for the \mathcal{O}_N -smoothed ratios (Figure 7(c)) the concavity is random and fairly small everywhere.

Figure 8 shows the detailed process for a specific cohort, the medium educated males born in 1925. In the top left panel we can see that, when unsmoothed ratios are used, the concavity is large and positive, and that there is quite a lot of noise. Not a single point, or the average concavity, falls within the $0 \pm \sigma$ confidence interval marked by the blue lines. In the top right panel, where \mathcal{O} -smoothed ratios have been used, we see a decrease in the total concavity (note the change of scale), and all the dots come closer to the average marked by the red line. But even though we managed to reduce the noise, the average concavity is still higher than zero and outside the $\pm \sigma$ interval. Finally, in the bottom panel, we see what happens with the \mathcal{O}_N -smoothed ratios. The

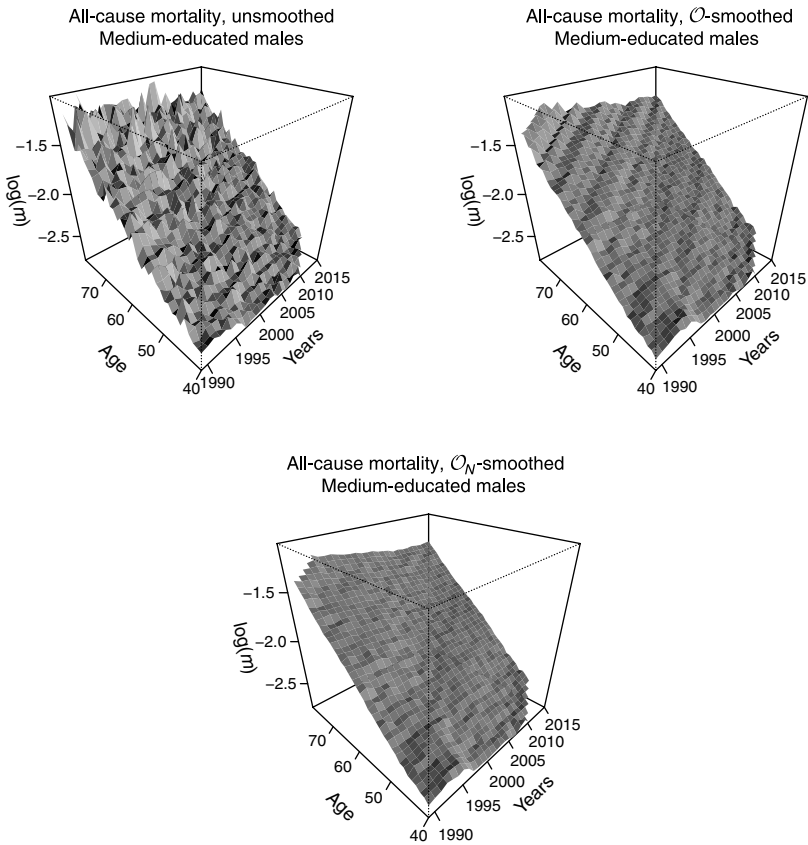


Figure 6. Death rates for medium educated males using the unsmoothed ratios (top left), smoothed ratios obtained minimising \mathcal{O} (top right) and smoothed ratios obtained minimising \mathcal{O}_N (bottom).

pattern of the dots is extremely similar to the \mathcal{O} -smoothed case, but they have been brought down so that the average concavity nearly vanishes. The dispersion of the points is compatible with our estimation of the standard error coming from the expected random variations in the number of deaths.

The reason for this can be seen in Figure 9. In this figure, we see the death rate for the 1925 birth cohort of medium educated males as a solid black line, along with the death rates for the two previous (dashed black) and the two following (dashed red) cohorts. A systematically positive concavity would appear if the death rate for the central cohort is higher than for its neighbours, which would be due to an underestimation of the exposures (underestimation of the initial ratio, $\hat{R}_0^{(e,c)}$). Indeed, when unsmoothed ratios are used the death rate for the central cohort is not only higher than for its neighbours, but has also some big jumps at higher ages, as expected from such noisy death rates as the ones seen in Figure 9. When the \mathcal{O} -smoothing is applied to the ratios, death rates for each cohort become smooth independently that eliminates all the big oscillations in both the death rates and the concavity, but the rate for some cohorts might be pushed “out of place.” This could easily happen if just a few points in a ratio time series are systematically lower than their “true” value, either because of some systematic bias of the sample used in the CPS or due to pure random variation, with the latter probably playing a big role for cohorts with few observations. In particular, for the cohort we are studying, we see that for ages 73 and 74 years the unsmoothed death rate is much higher than at any other age for any of the cohorts shown. It is probably due to this that \mathcal{O} -smoothing estimates a death rate that is higher at all ages for this

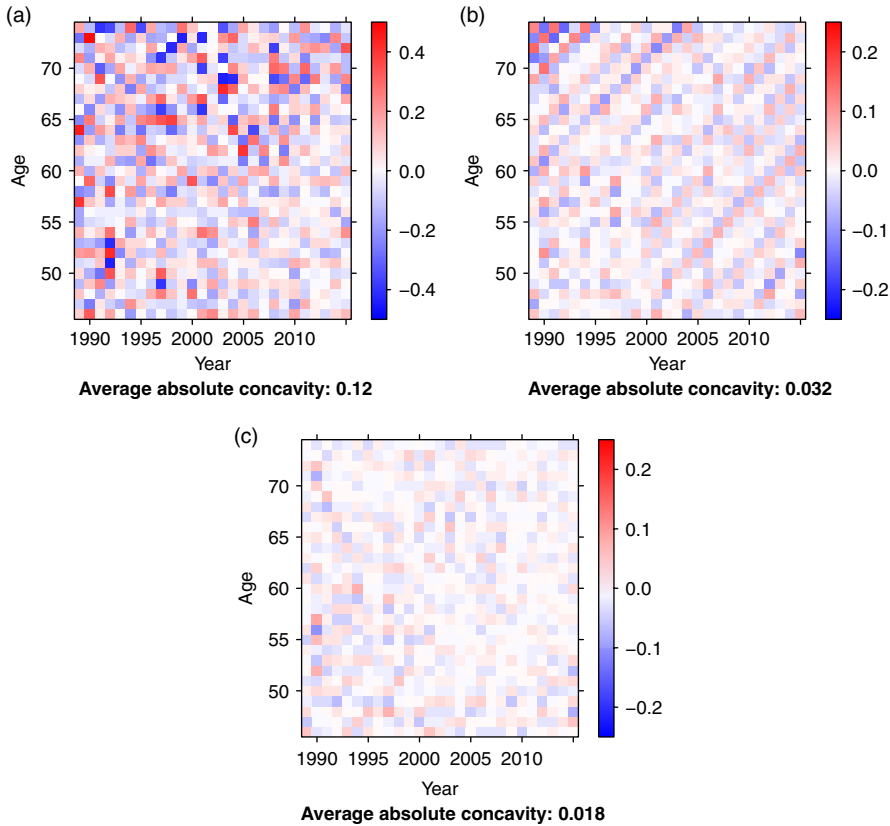


Figure 7. Heat maps of the concavity of the log death rates for medium educated males. (a) Concavity, unsmoothed ratios. (b) Concavity, \mathcal{O} -smoothed ratios. (c) Concavity, \mathcal{O}_N -smoothed ratios.

specific cohort than those of its neighbours, since these “rogue” points are driving the whole curve up. However, and as expected, the older 1924 and 1923 cohorts have death rates higher than the 1926 and 1927 cohorts.

And here is when the “interaction” term in \mathcal{O}_N becomes justifiable, we cannot let a few unusual data points affect our final results so heavily. By adjusting our ratios not only on a cohort by cohort basis, but comparing with their neighbours, strange oscillations like the ones seen at high ages for the medium educated 1925 males will be smoothed out to behave in a way closer to the one observed for similar cohorts. This is a reasonable expectation as long as we don’t have any reasons to believe that the big jumps in the death rates are real and not a problem with the exposures (e.g. a period effect that is also seen at the same calendar year in neighbouring cohorts; however, from Figure 9, this does not seem to be the case here). The bottom panel of Figure 9 shows how, when \mathcal{O}_N -smoothed ratios are used, the death rate for the 1925 cohort has been “brought down” to be in a more reasonable position, quite close to the red and black dashed lines.

Now that we have produced smoothed education ratios we can test whether or not the observed, unsmoothed ratios from the CPS are compatible with them by means of statistical tests. We will briefly explain three different tests that can give us information about how good the smoothed ratios are and then discuss the results obtained from their application to our problem.

The goal of the tests is to check how the CPS results compare to a binomial distribution with number of trials equal to the number of people interviewed and probability of success the smoothed ratios we calculated. Tests will be implemented following the cohort dimension, that is, for each cohort (people with the same gender, education level, and year of birth), we have an

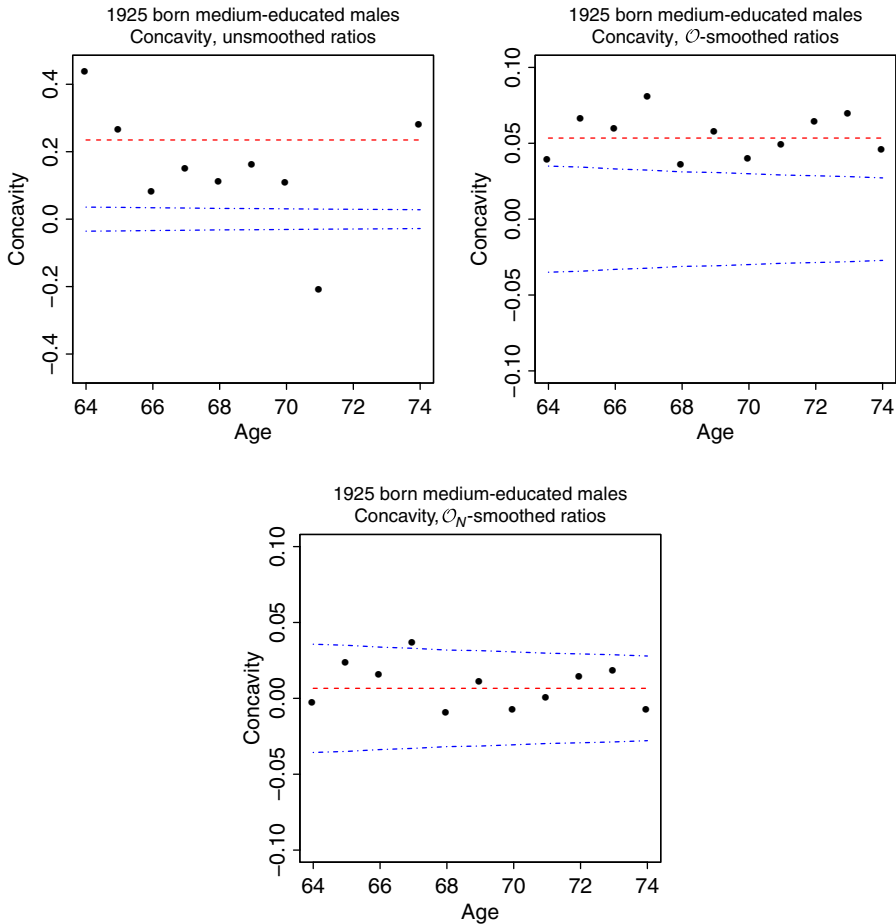


Figure 8. Concavity of the log death rates for the medium educated males born in 1925 (dots), along with its average (red line) and the $\pm\sigma$ interval around zero (blue lines), assuming underlying linearity.

unsmoothed trajectory and a smoothed one (the CPS and \mathcal{O}_N -smoothed ratios of people of that gender born in that year that achieved a certain level of education, respectively) that we will try to analyse.

The first test we will use is the likelihood ratio test. We implement this test by computing the deviance for each cohort, which is $D_C = 2(l_1 - l_2)$, with l_1 being the log-likelihood of a saturated model and l_2 the log-likelihood of the model we want to test (see, e.g. Macdonald *et al.*, 2018). For a binomial distribution, we have

$$\begin{aligned}
 D_C = 2 \sum_i & \left[c_c(i, c; e) \log \left(\frac{c_c(i, c; e)}{c_c(i, c) \hat{R}_C(i, e, c)} \right) \right. \\
 & \left. + (c_c(i, c) - c_c(i, c; e)) \log \left(\frac{c_c(i, c) - c_c(i, c; e)}{c_c(i, c) - c_c(i, c) \hat{R}_C(i, e, c)} \right) \right] \quad (14)
 \end{aligned}$$

where $c_c(i, c; e)$ is the number of successes (number of people that reported the education level e in the interview), $c_c(i, c)$ is the total number of people interviewed and $\hat{R}_C(i, e, c)$ is the smoothed ratio. The C subscript means that all these quantities are defined for cohorts, and the sum over i that we are aggregating over the cohort dimension, similar to the sums in equations (6) and (8). Now we assume this test statistic D_C follows a χ^2 distribution with as many degrees of freedom

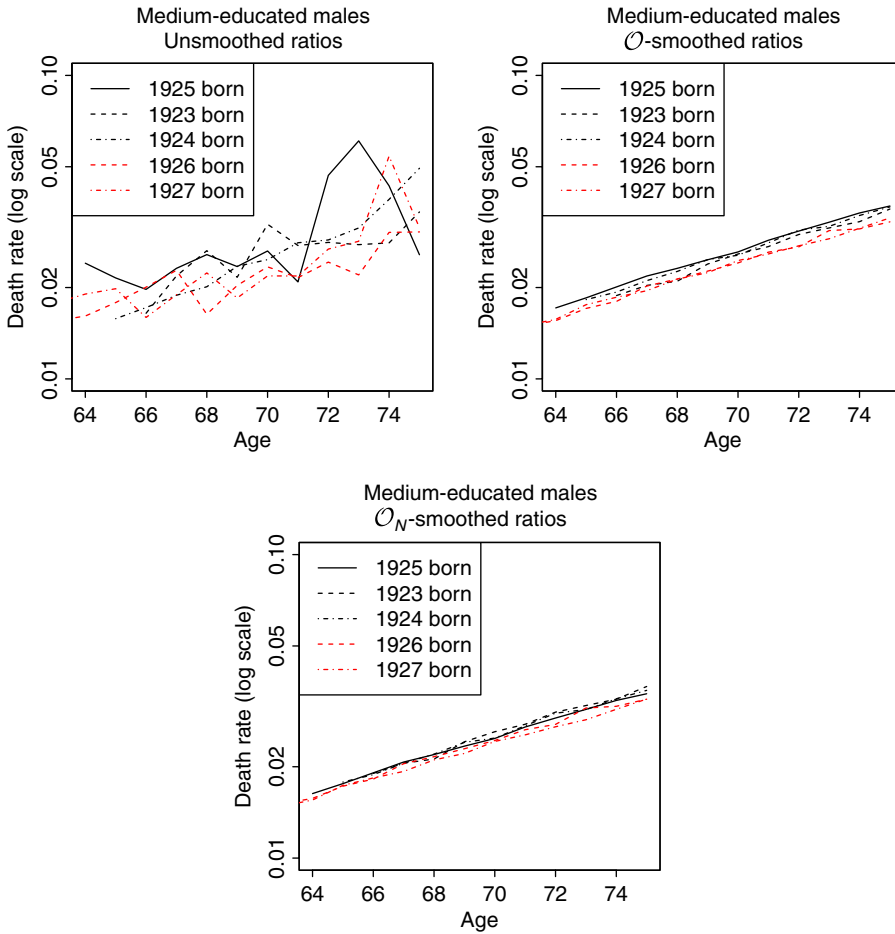


Figure 9. Death rates for the 1925 medium educated males cohort, along with the two previous (dashed black) and two next (dashed red) cohorts using unsmoothed, \mathcal{O} -smoothed and \mathcal{O}_N -smoothed ratios.

as data points we have for the cohort minus one (since the initial ratio $\hat{R}_C(0)$ was fitted using the unsmoothed ratios given by $R_C^{CPS} = c_c(i, c; e)/c_c(i, c)$). This will give us a p -value for the probability that, if the smoothed ratios \hat{R}_C were representative of the whole population, unsmoothed trajectories like the ones in the CPS, R_C^{CPS} , have a deviance like the one observed or greater. If this p -value is small, then it is not likely that our smoothed ratios are close to the true ratios for the whole population.

The second test we will use is the signs test. We compute the residuals $r_c(i) = \hat{R}_C(i) - R_C^{CPS}(i)$, the difference between the CPS ratios and the smoothed ratios for each cohort c at all times observed, i . We then check how many of these residuals are non-negative, $r_c(i) \geq 0$. If the model fits the data well, we expect the probability of finding non-negative residuals to be equal in practice to the one of finding negative ones (the probability of a residual being exactly zero being negligible). Therefore, we can calculate the probability that, in a binomial distribution with success probability 0.5 and number of trials equal to the data points for the cohort, we have as many successes (non-negative residuals) as we observe or more and check if we have too many or too few non-negative residuals compared with the expectations. In this case, very small or very large p -values (representing too many and very few non-negative residuals, respectively) would mean

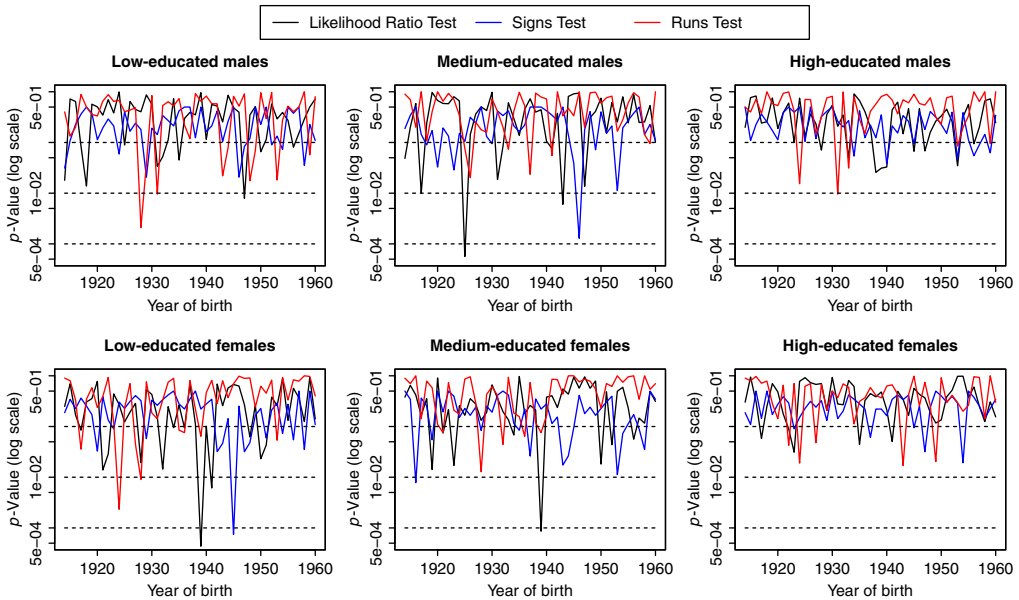


Figure 10. *p*-Values obtained from applying the three tests to the ratios of educated people on a cohort by cohort basis, in logarithmic scale: likelihood ratio test (black lines); signs test (blue); runs test (red). The 90%, 99% and 99.9% significance levels are marked by the dashed lines.

that there is a bias in the smoothed ratios, which would be systematically above or below the unsmoothed curve.

Finally, we have the runs test. The runs test starting point is the sign of the residuals calculated for the signs test. We want to check how many times the sequence of $r_c(i)$ for a cohort changes signs along the trajectory (because we are dividing the residuals in negative and non-negative, any residuals equal to zero are considered to have positive sign; however, as mentioned earlier, the probability of this happening is negligible). This number will be one less than the number of runs of residuals with the same sign, u . If the residuals are truly random, we don't expect them to show any structure. This structure would be represented by clusters of contiguous data being systematically above or below the smoothed curve. Therefore, the test should check whether the number of runs observed is too small and not compatible with random ordering. The equations for the probabilities of observing a specific number of runs, $P[U = u]$, which can be consulted in Macdonald *et al.* (2018), can be used to compute $P[U \leq u]$, the probability that the number of runs in a random sequence is less than or equal to our observation. This will give us a *p*-value representing how likely we were to find such ordered residuals given the number of positive and negative signs observed.

The *p*-values for all tests can be seen in Figure 10. For each gender education group, we plot the *p*-value obtained using the three tests against calendar year, in logarithmic scale. The horizontal dashed lines mark the 90%, 99% and 99.9% significance levels. Note that, due to the symmetric nature of the signs test (we would reject the null hypothesis if the *p*-value is either too big or too small), we do not plot *p* for this test, but $\min(p, 1 - p)$.

From the figure, it is obvious that in most cases the tests cannot reject the null hypothesis that our unsmoothed ratios were produced by true population ratios given by our smoothed results with 99% confidence. We are analysing a total of 47 years of birth, two genders and three education levels, which makes up a total of 282 cohorts. On these we are conducting three tests, so the number of *p*-values plotted in Figure 10 is 846. It is perfectly understandable then that some of these *p*-values are smaller than 0.01 simply due to high number of tests conducted.

However, it is encouraging that for not even a single cohort more than one p -value is smaller than 0.01 at a time, that is, for cohorts with a relatively high difference between the unsmoothed and smoothed ratios (which therefore fail the log-likelihood ratio test), there is no systematic bias in the smoothed ratios (otherwise the p -value for the signs ratio would be small too); and in a few cohorts for which the p -value given by the signs or runs test is small (which could mean our smoothed ratios are biased), the p -value of the log-likelihood ratio test is relatively high (therefore, if this bias truly exists it must be really small or the deviances would be high). From this picture, it looks reasonable to assume that our smoothed ratios are a good representation of the whole population.

5. Extrapolation to Higher Ages

Up to now, we have restricted the maximum age in our analysis to 79 years. In fact, due to the smoothing procedure, and because the CPS only provided estimates of the population by single year of age up to age 79 years in the 1989–2015 period, the oldest cohort included was the 1914 birth cohort. This choice ensured that for the oldest cohort analysed (whose members would be aged 75 years in 1989) we had at least five CPS data points to which we could fit our smoothed ratios. The choice of the starting age was arbitrary, but age 40 was chosen so that we can expect the death rates to increase exponentially for the whole range of ages considered (otherwise the vanishing concavity condition would not be applicable). However, our recurrence relation for the ratios, equation (3), means we have a systematic way to extrapolate the ratios of educated people in a cohort up to any age for which we have death counts (separated by education level) and total exposures by single year of age. The CDC data and HMD respectively give us that information for ages beyond 100 years.

Extrapolating is obviously useful on itself as a way to extract information about mortality at high ages, but it will also prove to be important as a test of the quality of the results. A small over- or under-estimation of the initial ratios of educated people, or any issues with the quality of the data, could potentially be hidden at relatively young ages when exposures are big and mortality low. However, any such issue will have a very significant effect at very high ages, when the cohorts are almost extinct. As we will see, the results of our extrapolation will force us to further modify our analysis due to concerns about data quality.

The extrapolation process is almost trivial to apply: once we have calculated our smoothed ratios of educated people for any range of ages, we use equation (3) to estimate those ratios beyond the highest age for which we have CPS data. We can do this because the HMD gives us the $E_i(c)$ and the CDC data the $\Delta_i^{(e,c)}$ we need up to ages much higher than 79 years. In fact, and since we use a cohort-based approach to extrapolation, we will obtain death rates for the cohorts born between 1914 and 1970 at any attained age between 1989 and 2015.

Note that, unlike other extrapolation methods, such as the widely used Gompertz model, we are not making any assumptions on the functional form of the death rate at higher ages. We are letting the data “speak for themselves,” and the resulting death rates for each education group will only depend on the raw data supplied to the algorithm. The single assumption we need to make is the previously explained negligible effect of migration on the education profile of the population, which in fact becomes more accurate for higher ages.

This non-parametric approach has the advantage of not forcing death rates at high ages to behave in a predetermined way, but as a consequence it can be quite sensitive to problems in the data. For example, if death counts for one of our groups were overestimated at higher ages, we would “run out of exposures” for that group and the death rate would quickly increase for a certain subpopulation as a cohort ages. This means that after applying this extrapolation technique some kind of test of the results needs to be made, for example a robustness analysis, to make sure that the results obtained are sensible and not driven by accidents in the data.

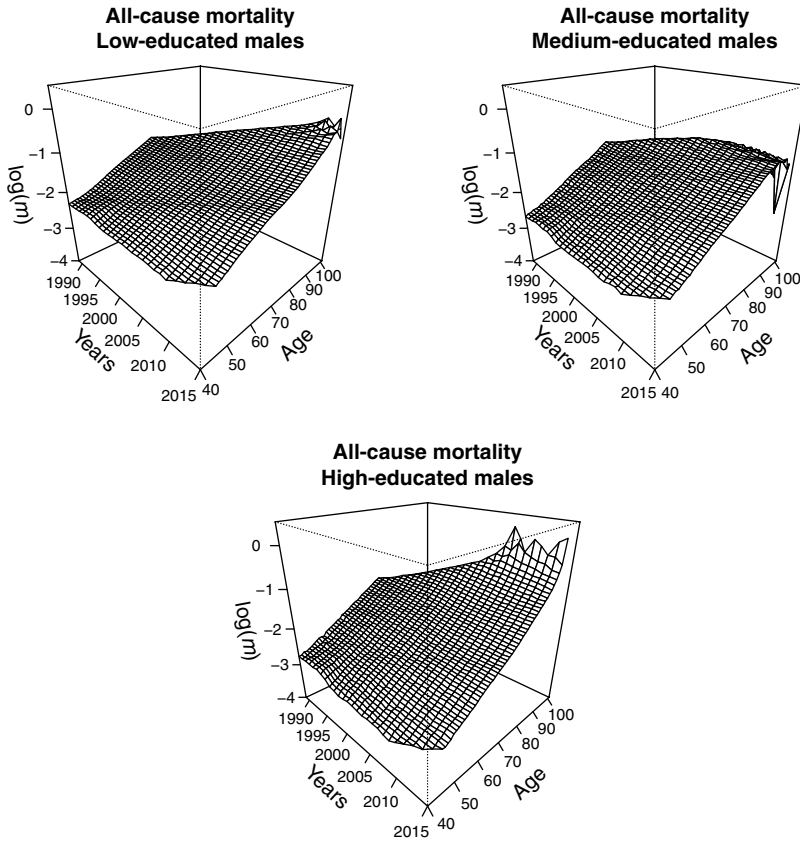


Figure 11. Death rates for males of all three education groups extrapolated to the highest ages possible.

Figure 11 shows the resulting extrapolated death rates for males. For each cohort, we have extrapolated the death rates up to the highest age possible (i.e. the age of that cohort in the calendar year 2015). Simple visual inspection shows that there is a problem in the extrapolated area. The mortality for group 2 reaches a *plateau* at very high ages and stagnates at a relatively low value, whereas the mortality for group 3 explodes at those same ages and even surpasses that of group 1. This is not a feasible behaviour, and we can only conclude that the extrapolation is failing for the medium- and high-educated groups.

This strange behaviour is extremely obvious if we plot the mortality of a cohort as it ages. A few examples are shown in Figure 12. We see that for some cohorts the mortality of group 2 stops increasing, or even decreases, at very high ages; conversely, the mortality of group 3 grows very quickly as the cohorts age. For the first three cases, the mortality of group 3 overtakes the mortality of group 2 around or before the 80 years of age, and in the last one both mortalities are quickly converging by that age (note the different scale in the x-axis for each cohort). This can be caused by either an underestimation of the exposures or an overestimation of the number of deaths.

Something that can help us understand where the problem lies is comparing our smoothed ratios with the ACS data. If the problem is a systematic underestimation of the ratio of people in group 3 (with the corresponding overestimation of group 2), then this should show up clearly as a discrepancy between our smoothed ratios and the relatively noise-free ratios of the ACS. This comparison is shown in Figure 13. Because of the restricted availability of ACS data, we can only do the comparison for two of the cohorts shown in Figure 12. Even with that the results seem to be clear, we are not underestimating the number of people in group 3. If there is no systematic

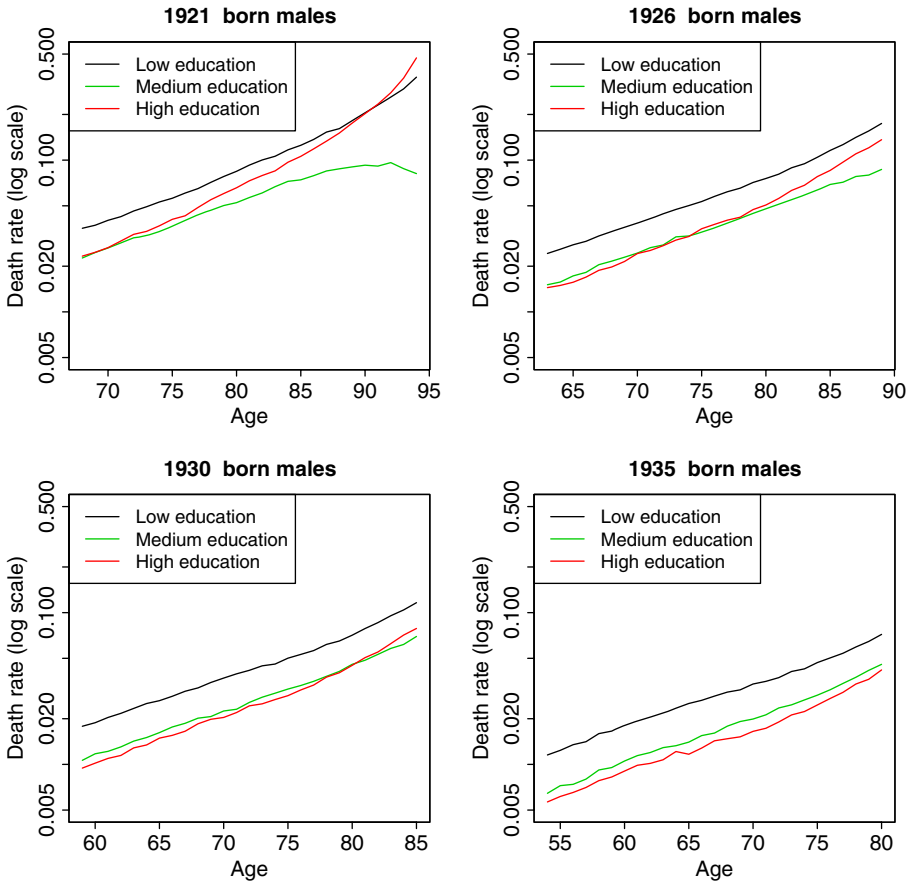


Figure 12. Death rates for males of different cohorts for all three education levels.

bias in our smoothed ratios, then the only reason why mortality would grow so quickly at high ages would be the misreporting of education in death certificates.

Misreporting of education in death certificates has been well studied, since comparability with other data sources was a key element to take into account when this new item was introduced. Rostron *et al.* (2010), building on the work of Sorlie & Johnson (1996), contains very valuable information regarding comparability between education reported in the CPS and in the death certificates. Thanks to the National Longitudinal Mortality Study, death certificate information for CPS respondents from 1992 to 1998 who died during that period could be compared with their interview responses. With that data in hand, the authors compare both sources of information on educational attainment of the individuals.

Of course, the first thing they note is that “[...] decedents without death certificate education information generally had lower educational attainment in the CPS than decedents with this information. For example, 49% of decedents without death certificate education information did not graduate from high school according to the CPS compared with 35% of decedents with death certificate education information.” This means that our simple imputation, in which we assume CDC certificates are randomly missing education, is not completely accurate. People with lower education are more likely to have an “unknown” education in their death certificate than higher educated people. However, we have modified our imputation method to analyse two extreme cases (namely, deaths with unknown education are all ignored or all assigned to group 1) and neither of them changes the qualitative behaviour of the death rates for groups 2 and 3 seen in Figure 12. This




Figure 13. Comparison between raw ACS and smoothed CPS ratios.

means that the problem with the excess of deaths in group 3 does not come from the imputation procedure.

The authors also explore the problems of correspondence between the degree-based and years-of-schooling-based systems. Thanks to this they produce results that are very relevant for us: Figures 1 and 2 of their paper show how their adjustment for misreporting corrects death rates as obtained from deaths reported using the 1989 and 2003 death certificate standards. For both males and females, regardless of how education was recorded, the corrected life expectancy at the age of 25 years for high school graduates is bigger than before correction, whereas in turn the life expectancy for people without high school completed is smaller. This is caused by funeral directors writing a decedent’s education down as “completed high school” when the person had never achieved their degree, as it is a very well-known bias. However, we also see that the life expectancy for people with bachelor’s degrees or graduate degrees (which together form our education group 3) as calculated using death certificates in the 1989 standard is increased when the correction is applied. However, life expectancy for people with some college is decreased by the correction. This is compatible with our observations: raw death counts of people with bachelor’s degrees or more are overestimated, whereas deaths of people with some college but no degree seem to be underestimated. For death certificates with the 2003 standard, the authors chose to use three different education categories. In this case, their “more than high school graduate” category, which corresponds to our groups 2 and 3 merged together, seems to be quite consistent.

There is another easy way to further diagnose our data. When every member of a cohort we are observing has died, the number of people who died in each education group divided by the total number of deaths should be equal to the initial number of people in each group divided by the initial number of people. This is, the quantity

$$R_C^D(i, e, c) = \frac{\sum_{j=1}^i D(t_0 + x_0 + j, x_0 + j, e, g)}{\sum_e \sum_{j=1}^i D(t_0 + x_0 + j, x_0 + j, e, g)} \tag{15}$$

which measures the proportion that each education group has contributed to the total number of deaths in a cohort from the point they enter our analysis up to a time i , should converge to the initial value of the smoothed ratios $\hat{R}_C(0)$ as i increases ($D(t, x, e, g)$ is the number of deaths, as defined in section 3). Figure 14 contains plots of this quantity for certain cohorts. Solid lines represent how much each education group has contributed to the total number of deaths up to that age, dashed lines are the smoothed ratios of educated people and the horizontal dot-dashed line marks the value of the initial ratio for each education group. Solid lines should approach the

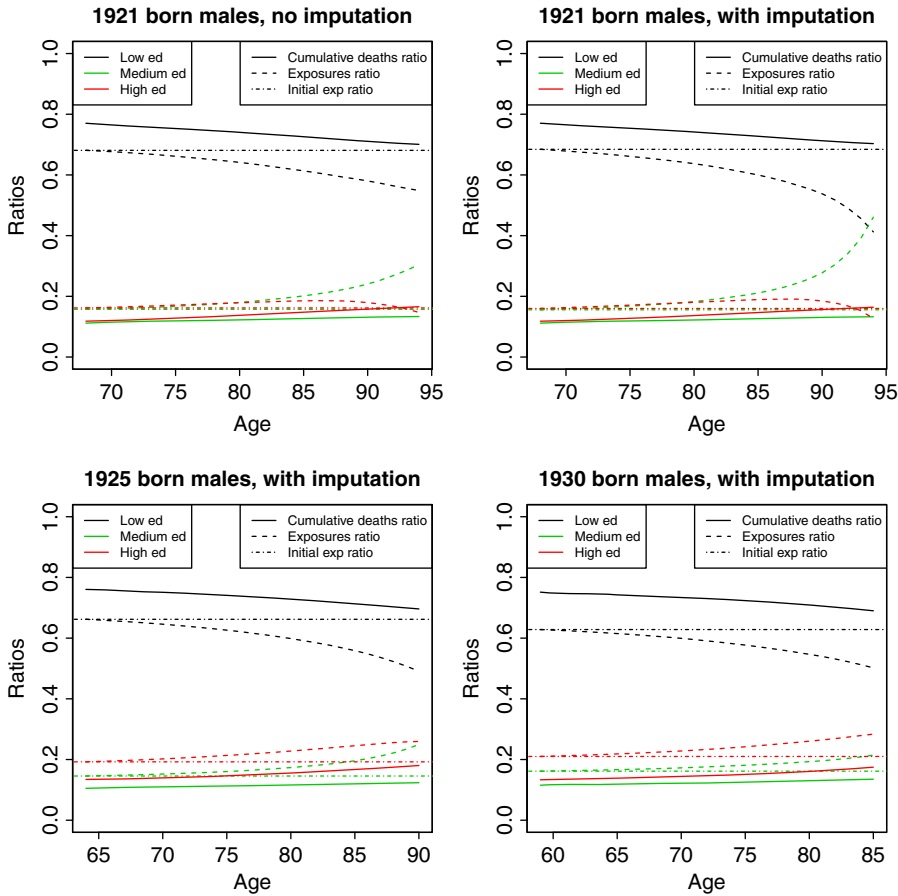


Figure 14. Figures comparing the ratios of cumulative deaths with the initial ratios of educated people.

horizontal line as the cohort ages, but crossings are not expected and mean problems with the data.

The top two panels show what happens with the 1921 cohort. As expected, because lower educated people die younger, they make the bulk of the deaths at younger ages (the black solid line starts much higher than the dot-dashed line, whereas for the two other groups the opposite is true). As the cohort ages all three solid lines start approaching their theoretical limiting values, but around 90 years of age the solid line for group 3 has crossed over the dot-dashed line, whereas the solid line for group 2 seems to be getting stagnant still far from its supposed value at very high ages. This happens in both panels, which means the problem is not being introduced by the education imputation procedure. As we saw, our smoothed ratios are comparable to values extracted from other sources, which means the problem is likely to be people of group 2 being misreported as having a bachelor’s degree or more in their death certificates.

The two lower panels of Figure 14 show what happens with two younger cohorts. Although the crossover has not happened yet for these, the tendency of the red line to approach its supposed asymptotic value much faster than the two other solid lines is clearly visible, and especially in the 1925 cohort (bottom left) it is clear that the crossing will eventually happen. This shows that the issue always appears at very high ages, not only for the older cohorts.

For all of this, we should conclude that it is the quality of the death counts data that is limiting the validity of our results. One possible solution, seeing how the results for the lower educated

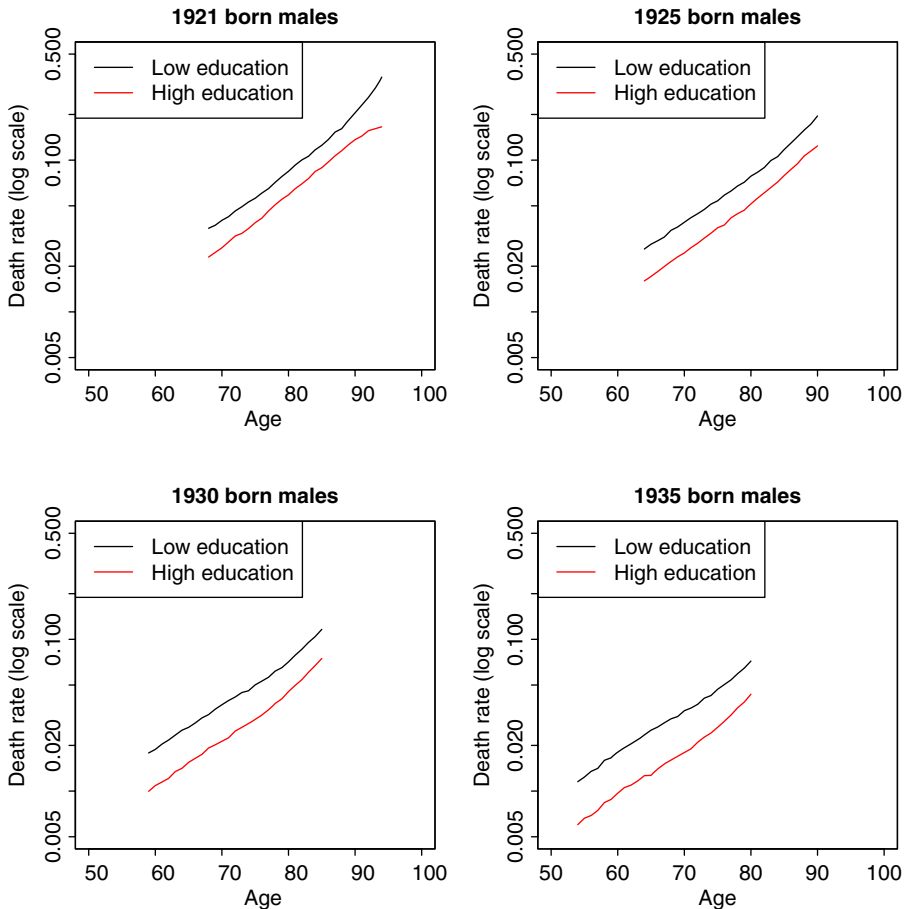


Figure 15. Death rates for males of different cohorts for the new two education levels.

group are consistent in our analysis (the black solid lines do indeed approach its theoretical asymptote), and the conclusions of Rostron *et al.* (2010), is to combine the two higher educated groups and divide the whole population in just two categories, “never went to college” and “at least some college completed.” For the cohorts shown in Figure 12, we obtain the new death rates plotted in Figure 15. In these, there is still a clear hierarchy between the education groups which is maintained at all ages. Therefore mixing groups 2 and 3 in a single, broad higher educated group lets us obtain reliable results up to very high ages.

One last concern is how robust these results are. If our procedure is robust, then small changes to the raw data or the input parameters would not affect the results. However, if they produce a significant change in the final results, then they should not be trusted. If we have any robustness problems we would expect to find them specially in the oldest cohorts, the ones for which very few data points are available in the CPS, since the lack of data for long periods of time would make them more vulnerable to small changes in the data available and its treatment.

We introduce two changes in our analysis. Firstly, we move the starting year from 1989 to 1995. This will change the number of data points available for some of the oldest cohorts. We will also change the weights in our smoothing procedure that we had (perhaps naively) overlooked so far. Instead of setting all of the ω_i in equation (9) to a constant, we will use the estimated error in the unsmoothed ratios (given by equation (10)) to give a weight to each point. We will make the obvious choice $\omega_i = \sigma(\hat{R})^{-2}$ and see if this weighting affects the results obtained or not.

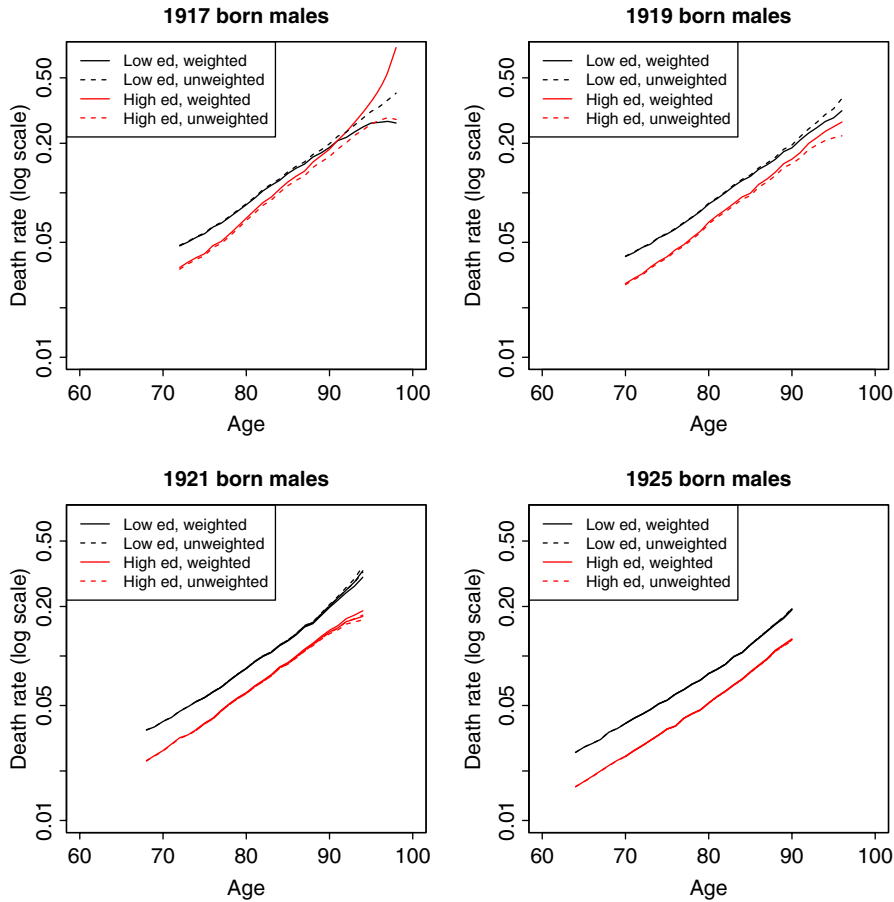


Figure 16. Death rates for males of different cohorts calculated in four different ways (1989–1995 as starting year, and weighted-unweighted fit of the smoothed ratios).

In Figure 16, we see how the death rates for certain cohorts change as we change the method used to calculate them. The solid lines are the death rates when $\omega_i = \sigma(\hat{R})^{-2}$, and the dashed lines when $\omega_i = \text{constant}$. For the cohorts for which it is possible (1921 and 1925) there are two sets of lines, one of which starts at a later age. These second death rates are obtained starting the analysis in 1995.

It is obvious that we have a problem with robustness in the 1917 and 1919 cohorts. It is also clear that the problem is only visible at high ages. As expected, small changes in the initial ratios of educated people create “trolls” or “phantoms” (as defined in Cairns *et al.*, 2016) that do not manifest themselves until the total cohort membership is very small. We also see how for younger cohorts the results seem to be robust for a wider range of ages: for the 1917 cohort, we already have a noticeable discrepancy between the different death rates at the age of 90 years, whereas for the 1925 cohort there is barely any difference between them. There are two possible causes: one is that younger cohorts maintain higher memberships for longer (i.e. not as many people live for 90+ years if born in 1917, compared with people born in 1925); another difference is that, because the 1917 cohort enters our analysis at age $1989 - 1917 = 72$, we only have eight points to which we can fit our smoothed ratios, whereas for the 1925 cohort (which enters the analysis at age $1989 - 1925 = 64$) we can use 16 unsmoothed ratios in the fitting procedure. More data points mean of course a better, more reliable fit.

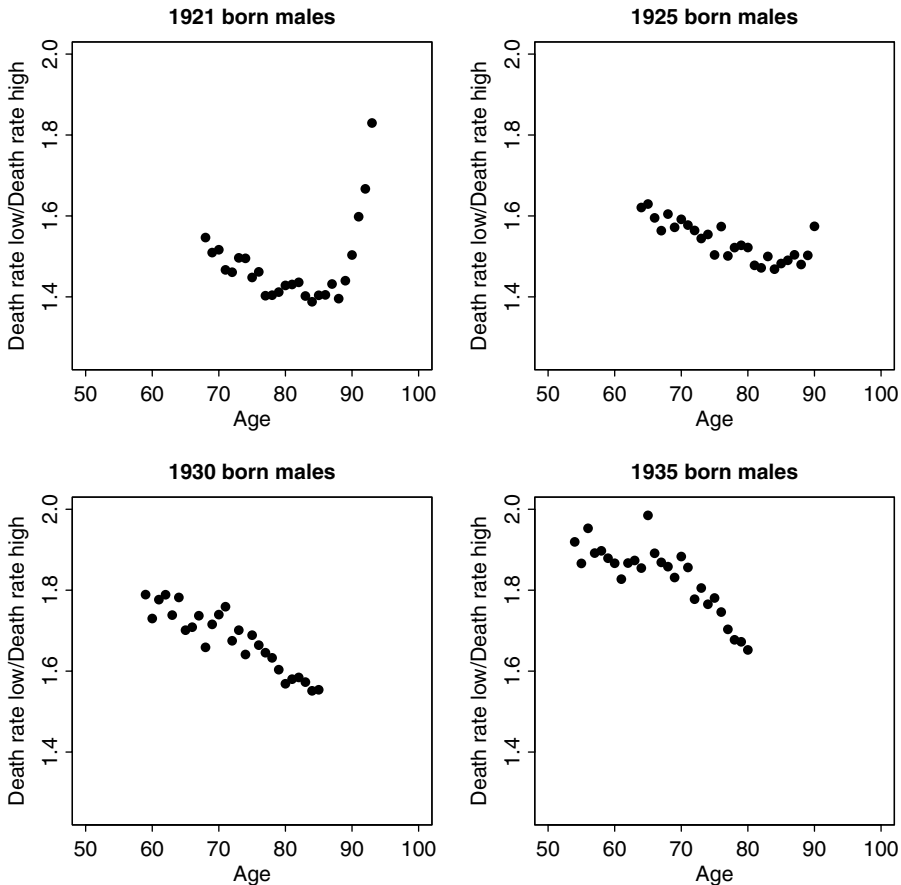


Figure 17. Ratio between the low and high-educated group death rates for selected male cohorts. The convergence between the mortality of the two groups as we move towards higher ages is clearly seen up to roughly 90 years of age. The rapid increase in the gap at very high ages for the oldest cohorts shows that the extrapolation fails to produce reasonable results in that region.

We can also use the evolution of the mortality gap as a diagnosis of the quality of the results. We expect the gap between the two groups to close as cohorts age, since socio-economic status mainly affects mortality at younger ages. Visual inspection of Figures 15 and 16 show that for some cohorts there is a period of convergence up to roughly 90 years of age, followed by a divergence afterwards. Looking at the ratio of the mortality between the two groups would help us see this behaviour more clearly and find the point at which the trend changes. The reversal of the convergence at a very high age would tell us at which point our extrapolated results become unreliable.

These ratios are plotted in Figure 17 for selected cohorts. In these, it is clear that the death rates initially converge as a cohort ages, as expected. However, for the cohorts for which we have data at very high ages (1921 and 1925), this converging trend changes as we approach 90 years of age, signalling the point at which the calculated death rates cannot be trusted any more.

The above analysis shows that for the US data our extrapolation procedure only produces reliable results for ages up to 90 years approximately (as clearly demonstrated by Figures 16 and 17). This is still roughly a 10-year improvement over what we could have achieved without extrapolating the CPS data (which only reaches the age of 79 years). Of course, applying this method to different populations could lead to different results, producing good results for a longer or shorter

range of ages. An analysis on the lines of what we have shown here should be always performed after using the proposed method to diagnose the results obtained.

6. Final Results: All-Cause and Cause-Specific Death Rates

In the previous sections, we have looked at the construction of the death counts and exposures datasets in detail. Now we have everything we need to produce death rates by CoD and for all-cause mortality, for two different education groups, and for males and females separately.

Table 2 shows the list of CoDs analysed, as well as the ICD-9 and ICD-10 codes used to identify them. As explained earlier in this paper, the correspondence between CoDs in the different ICDs is not perfect, not only because of the change in the coding of certain illnesses but also because of the modification of the process by which the underlying CoD is determined when there are several contributing factors. We believe, however, that it is much more desirable to produce a database with a high level of granularity, as this allows us to differentiate between CoDs with specific behavioural/lifestyle risk factors. Additionally, the categories here can be easily aggregated, if so desired.

The CoDs chosen cover all deaths in the US in the period analysed, and approximately 98% of them if we exclude the “Other causes of death” category. Some of them are heavily connected to lifestyle factors, and deaths due to those are therefore highly avoidable (e.g. it is estimated tobacco smoking is the direct cause of between 80% and 90% of deaths due to lung cancer, see US Department of Health and Human Services, 2004). Others, however, do not have very well-known causes (e.g. Alzheimer’s disease or prostate cancer). In the rest of this section, we will present plots of the evolution of the death rates for all-cause and some selected CoDs, with a mix of CoDs from the “highly preventable” and “not preventable” groups.

Figure 18 shows the all-cause death rate for males and females in years 2000 and 2010. The education gap is very clearly seen in all four panels, with the difference being much bigger for younger ages, and a quick convergence of the death rates for older people. The time evolution of this gap for ages 60 and 75 years is shown in Figure 19. The top panels show the death rates for males and females at the age of 60 years, and the bottom panels at the age of 75 years. It is very interesting to see that mortality improvements at the age of 60 years have occurred almost exclusively in the high-educated group, whereas death rates for the low educated have remained relatively constant throughout the period analysed. The stagnation of the improvements for the low-educated group, together with the decreasing rates for the high educated, have resulted in an increasing mortality gap between our two groups. The situation at the age of 75 years is not as dramatic for people in the low-educated category, as death rates have kept decreasing slowly; however, faster mortality improvements for the high-educated group have resulted in an increasing mortality gap as well for this age.

To understand these trends in all-cause mortality, we will now look at CoD-specific death rates. In Figure 20, we plot death rates by age for selected CoDs in males, year 2010. We have chosen a wide selection of different CoDs with different underlying causes and risk factors, which explains the different behaviours observed.

The top three panels show CoDs that are very heavily linked to one risk factor: lung and bronchus cancer (smoking), accidental poisonings (substance abuse), and alcohol liver disease and cirrhosis (alcohol abuse). We can see that all three CoDs disproportionately affect the lower educated group. It is also important to notice that these CoDs make up a higher *percentage* of all deaths at relatively young ages, since their rates stagnate (lung cancer and cirrhosis) or even decrease (poisonings) for older people. This behaviour of the lifestyle-linked CoDs, so different from most other CoDs (for which death rates grow roughly exponentially with age), partly explains why the mortality gap is particularly big in the middle ages.

The middle panels show three CoDs for which risk factors are either not very well known, or such that an individual cannot easily modify them through simple lifestyle changes. These are

Table 2. List of causes of death used with the corresponding ICD-9 and ICD-10 codes.

	ICD-9 codes	ICD-10 codes
Infectious diseases	001-139	A00-B99
Cancer in mouth and gullet	140-150	C00-C15
Stomach cancer	151	C16
Cancer in gut or rectum	152-154	C17-C21
Cancer in larynx	161	C32
Cancer in trachea	162.0	C33
Cancer in bronchus or lungs	162.2-162.9	C34
Breast cancer	174-175	C50
Cancer in uterus, cervix and neighbouring	179-180, 182	C53-C55
Cancer in ovary	183	C56
Cancer in other female genital organs	181, 184	C51-C52, C57-C58
Prostate cancer	185	C61
Cancer in other male genital organs	186-187	C60, C62-63
Cancer in liver	155	C22
Cancer in pancreas	157	C25
Skin cancer	172-173	C43, C44
Cancer in urinary organs	189	C64-C66, C68
Bladder cancer	188	C67
Cancer in lymphatic or blood-forming tissues	200-208	C81-C96
Benign tumours or tumours without specification	210-229, 235-239	D10-D49
Cancer in other locations	Rest of 140-239	C23-C24, C26-C31, C37-C39, C40-C42, C45-C49, C69-C80, C97, D00-D09
Diseases in blood and blood-forming organs, endocrine sufferings and deficiency diseases	240-289	D50-99, E00-E07, E15-E90
Diabetes	250	E08-E14
Vascular and other dementias	290, 294	F01-F03
Other mental illnesses	Rest of 290-319	F04-F99
Diseases in the nervous system and sense organs	Rest of 320-389	G00-G29, G32-G99, H00-H99
Alzheimer's and related degenerative diseases	331	G30-G31
Increased blood pressure or rheumatic fever	390-405	I00-I16
Ischaemic heart diseases	410-414	I20-I25
Other heart diseases	415-429	I26-I28, I30-I52
Cerebrovascular diseases	430-438	I60-I69
Circulatory diseases	440-459	I70-I99
Influenza, pneumonia and other acute lower respiratory tract infections	466, 480-488	J09-J22
Chronic lower respiratory diseases	490-496 (except 495)	J40-J47
Other lungs and breathing diseases	Rest of 460-519	J00-J06, J30-J39, J60-J99
Alcohol liver disease and cirrhosis	571	K70, K73-K74
Digestive diseases	Rest of 520-579	K00-K69, K71-K72, K75-K99

Continued

Table 2. Continued.

	ICD-9 codes	ICD-10 codes
Diseases in skin, bones and connective tissue	680-739	L00-L99, M00-M99
Urine, kidney, genital organs and breast gland diseases	580-629	N00-N99
Road and other accidents	Rest of 800-949	V00-V99, W00-W99, X00-X39, X50-X59, Y40-Y86, Y87.1-Y89, Y90-Y99
Suicide	950-959	X60-X84, Y10-Y34, Y87.0
Accidental poisonings	850-869	X40-X49
Other causes of death	Rest of codes	Rest of codes

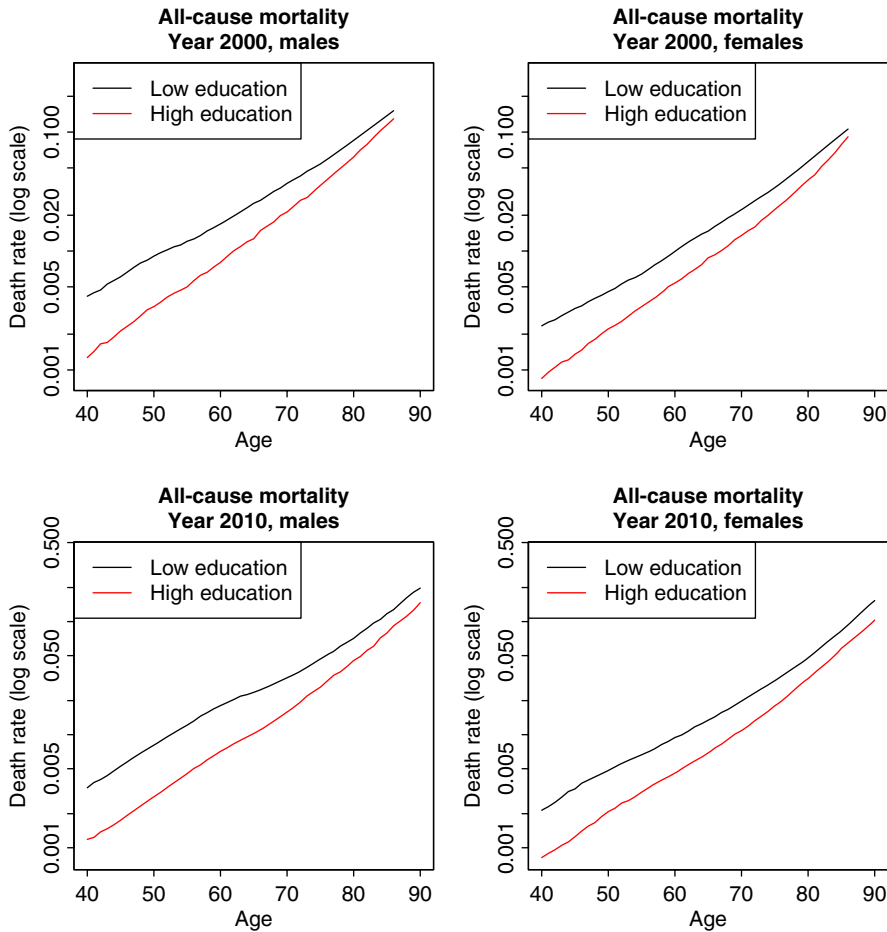


Figure 18. Plots of the all-cause death rates for males (left) and females (right), in years 2000 and 2010. The education gap and its narrowing with age are very clearly visible in all four panels.

prostate cancer, Alzheimer’s disease and other similar dementias, and skin cancer. These have the behaviour expected (roughly exponential growth), and the rates are particularly small at young ages, but overtake those of the lifestyle-related CoDs at very high ages, except for skin cancer not reaching the level of lung cancer. We can see that some small gap still remains, potentially due to uncontrollable risk factors (like genetics, different ethnic make up of each group, etc.) or

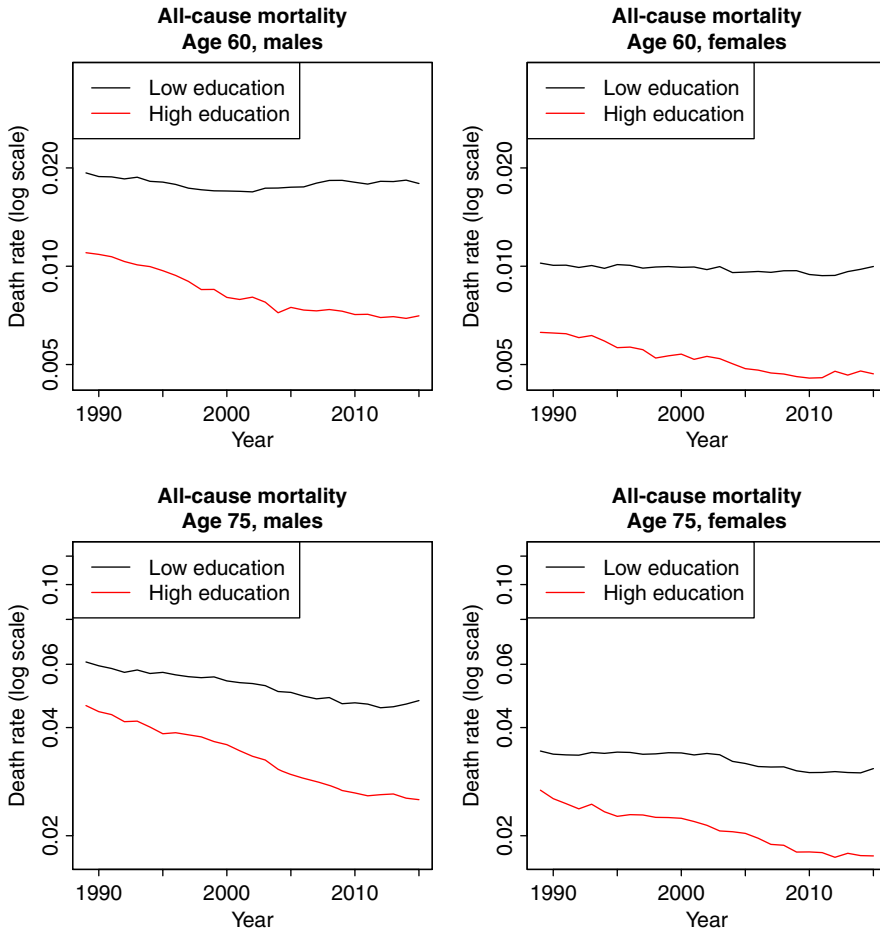


Figure 19. Plots showing the evolution of the death rates with calendar year for ages 60 and 75 years in males (left) and females (right). We can clearly see the stagnating death rate for the low-educated group at the age of 60 years, and the growing gap with their high-educated counterparts in all four cases.

unknown underlying causes of the disease (e.g. a link between certain dietary factors and prostate cancer might exist but has not been proven, see Grönberg 2003). Different levels of engagement with treatment and access to healthcare can also be behind the observed gap, see for example, Freeman (2004). In any case, the narrower gap for non-lifestyle-related CoDs, and their larger share of the total deaths in the older population, are at least in part behind the narrowing of the mortality gap at higher ages.

Lastly, in the bottom three figures we show the death rates for three big killers for which death rates at high ages are comparable to those of the previous CoDs. These would be in a category between the previous two, in which lifestyle does have a significant impact on the risk of developing the disease, but there is not one single risk factor that accounts for the majority of the deaths. For these diseases, the general behaviour is not unlike the all-cause death rates, increasing rapidly with age and with a narrowing gap as we move towards older ages. The main difference is the appearance of small “quirks” particular to each CoD, for example, the increased mortality in middle ages for infectious diseases, probably due to the burden of hepatitis C in the US middle-aged population, as reported in Zibbell *et al.* (2018).

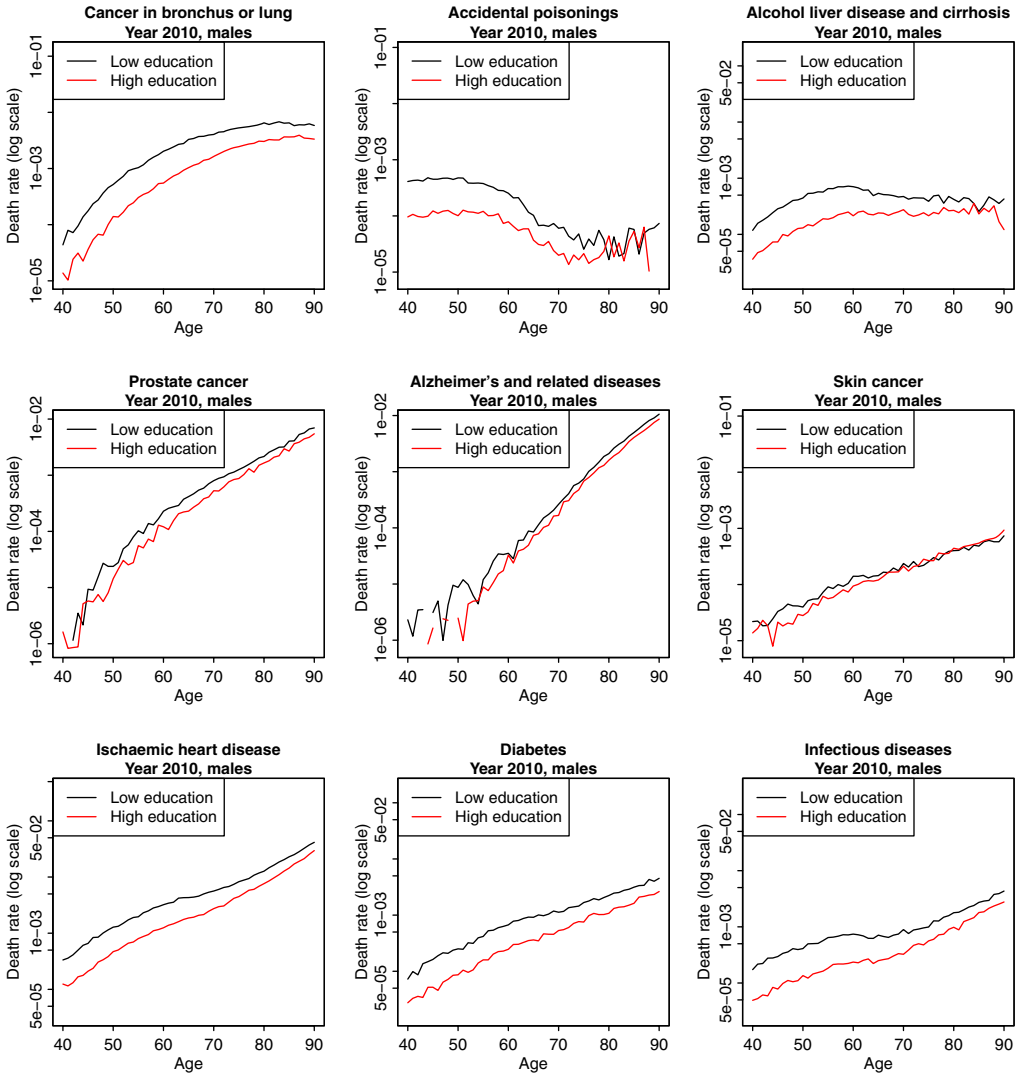


Figure 20. Cause-specific death rates with age for males in year 2010, selected causes.

Figure 21 shows the temporal evolution of the death rates in males aged 70 years for the same CoDs as the previous figure. It is interesting to note that the trends differ greatly between different CoDs, even within each of the three categories we described before (i.e. heavily linked to one risk factor, not easily preventable, and linked to complex lifestyle factors). For some of the biggest killers, such as lung cancer or ischaemic heart disease, there is a clear decrease in death rates, although improvements stagnate for heart disease in the late 2000s for the low-educated group and after 2010 for the higher educated. The faster decrease for the high-educated group in these CoDs, which are some of the main contributors to all-cause death rates, explain much of the increasing gap we observed.

An important conclusion of the results presented is that there is not a single cause that makes up most of the contribution all-cause mortality trends presented earlier. That is, it would not be correct to deduce that the stagnation in mortality for the low-educated group is due to an increase in accidental poisonings, or the slowdown of improvements in ischaemic heart disease mortality

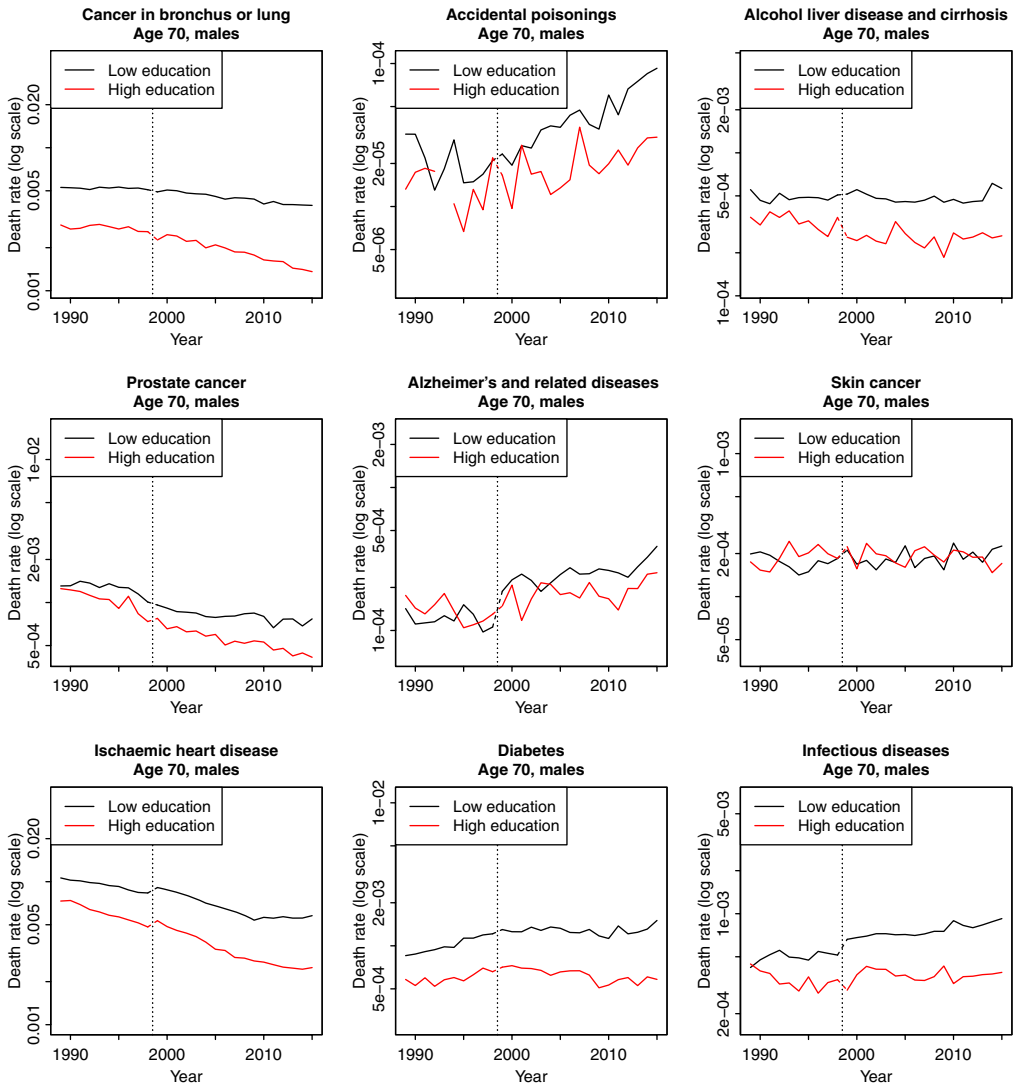


Figure 21. Cause-specific death rates with calendar year for males aged 70 years, selected causes. The vertical dotted line marks the point where the ICD standard changed.

only. It is in fact a combination of very different trends for different causes that is behind the all-cause trend presented before. For example, without the increase in mortality due to diabetes and infectious diseases, it is very likely that the decrease in prostate and lung cancer death rates would have masked the effects of the trends of ischaemic heart disease and poisonings, and an overall decrease in all-cause death rate would have been observed. Care must therefore be exercised when trying to answer the question of what is behind the trends in all-cause death rates, and we must avoid looking for oversimplified explanations of a very complex issue.

It is also noteworthy that even in a non-preventable CoD, like prostate cancer, the gap seems to be increasing. In this case, there is a mortality improvement for both groups, and the divergence is caused by a faster decrease in the death rate for the higher educated population. This could be explained by medical advances that benefit the general population (hence, the overall decrease), but that have a bigger impact in higher socio-economic strata, due to better engagement with




Figure 22. Plot showing the evolution of the proportion of people in the high-educated group (black, left-hand scale) and the ratio between the low- and high-educated groups death rates (red, right-hand scale).

treatments, more awareness of the latest health developments, treatment cost and insurance considerations, etc. (see Freeman 2004 for a discussion on the factors that affect cancer mortality in different social groups). However, there could also be a link to changes in the underlying but unknown risk factors (e.g. poor diet or a sedentary life could potentially cause prostate cancer, as discussed earlier, although the effect is small).

Another possible reasons behind the increasing gap are the changes in the relative compositions of both groups, this is, the lower educated people might be becoming an increasingly negatively selected group in which poor health and predisposition to early death is becoming more prevalent simply because through time more healthy individuals have access to higher education. The problem of broad categories like “low educated” and “high educated” being inconsistent over long periods of time has been addressed in the literature before, see Hendi (2015) or Brønnum-Hansen & Baadsgaard (2012). The main issue can be summarised in a simple question: how can we be sure that the growing gap in the death rates is due to changes in the death rates of people within a group, and not because through time more healthy individuals “move” into the higher group, making the lower group more and more negatively selected?

Our analysis includes cohorts of people born between 1914 and 1975. Over this period, the percentage of people going into higher education has increased significantly, and a higher percentage of people will be highly educated in the younger cohorts, see for example, Ryan & Bauman (2016). In Figure 22, we try to assess the impact of this increase in the ratio of high-educated people through time, or grade inflation, in our results. In black, and the left scale, we have the ratio of people in the high-educated group as a function of birth cohort, for males aged 50 and 70 years. The red line and right axis show the ratio between the death rates of the low- and high-educated groups as a function of birth cohort, again for males aged 50 and 70 years. We see that the increase in the mortality gap is highly correlated with a higher proportion of males in the high-educated group up to the 1949 cohort, for which education peaks. From that point on the ratio of high-educated males decreases and then stagnates as we move into the 1960s, whereas the gap keeps increasing until well into the 1950s born cohorts. This figure shows that grade inflation probably plays a part in the results observed, but the continuing increase in the gap even after the ratio of higher educated males levels off makes us confident that not only the changing demographics of each group, but behavioural lifestyle changes within each group, are contributing to the growing mortality inequality in the US. This is consistent with the findings reported in Hendi (2015). Quantifying the exact contribution of each factor to our observations would require more detailed data than the one we have available and is therefore out of the scope of this work.

7. Summary and Future Research

With the procedure described in this paper, we have managed to obtain death counts, exposures, and, combining the two, death rates for the US population born between 1914 and 1970 for years 1989–2015, separately for two groups of educational attainment and for 32 different groups of CoDs. Even though some concerns remain about the reliability of our results at very high ages, particularly for the oldest cohorts, some important insights have been gained thanks to the final data obtained.

As possibilities for further research, a systematic analysis of the data using statistical modelling would provide us with further insight on the observed trends, helping us understand which are the drivers of the diverging mortality scenario we describe here. These data will also be useful in other applications, for example, in comparisons of total and CoD mortality between different countries, since its granularity (single years of age and calendar years) offers greater flexibility than the age-standardised rates available elsewhere. We are confident that this dataset will be useful in many future applications and will be made available by the authors for any interested researcher to use.

Acknowledgements. This work forms part of the research programme “Modelling, Measurement and Management of Longevity and Morbidity Risk” funded by the Actuarial Research Centre of the Institute and Faculty of Actuaries, the Society of Actuaries and the Canadian Institute of Actuaries. The authors gratefully acknowledge feedback from the respective steering and oversight groups.

References

- Anderson, R.N., Minino, A.M., Hoyert, D.L. & Rosenberg, H.M. (2001). Comparability of cause of death between ICD-9 and ICD-10: preliminary estimates. *National Vital Statistics Reports*, **49**(2), 1–32.
- Bronnum-Hansen, H. & Baadsgaard, M. (2012). Widening social inequality in life expectancy in Denmark. A register-based study on social composition and mortality trends for the Danish population. *BMC Public Health*, **12**, 994–1000.
- Cairns, A.J.G., Blake, D., Dowd, K. & Kessler, A. (2016) Phantoms never die: living with unreliable population data. *Journal of the Royal Statistical Society, Series A*, **179**, 975–1005.
- Case, A. & Deaton, A. (2015) Rising morbidity and mortality in midlife among white non-Hispanic Americans in the 21st century. *Proceedings of the National Academy of Sciences*, **112**, 15078–15083.
- Freeman, H.P. (2004). Poverty, culture, and social injustice: determinants of cancer disparities. *CA: A Cancer Journal for Clinicians*, **54**(2), 72–77.
- Grönberg, H. (2003). Prostate cancer epidemiology. *The Lancet*, **361**(9360), 859–864.
- Hendi, A.S. (2015). Trends in U.S. life expectancy gradients: the role of changing educational composition. *International Journal of Epidemiology*, **44**(3), 946–955.
- Hoyert, D.L., Heron, M.P., Murphy, S.L. & Kung, H.-C. (2006). Deaths: final data for 2003. *National Vital Statistics Reports*, **54**(13), 109.
- Jemal, A., Ward, E., Anderson, R.N., Murray, T. & Thun, M.J. (2008). Widening of socioeconomic inequalities in U.S. death rates, 1993–2001. *PLOS ONE*, **3**(5), e2181.
- Macdonald, A.S., Richards, S.J. & Currie, I.D. (2018). *Modelling Mortality with Actuarial Applications*, 1 edition. Cambridge University Press, Cambridge.
- NCHS (1993). Public use data tape documentation: multiple cause of death for ICD-10 1990 data.
- NCHS (2016). National Center for Health Statistics, Vital Statistics Data. https://www.cdc.gov/nchs/data_access/vitalstatsonline.htm [accessed May 2017].
- Olshansky, S.J., Antonucci, T., Berkman, L., Binstock, R.H., Boersch-Supan, A., Cacioppo, J.T., Carnes, B.A., Carstensen, L.L., Fried, L.P., Goldman, D.P., Jackson, J., Kohli, M., Rother, J., Zheng, Y. & Rowe, J. (2012). Differences in life expectancy due to race and educational differences are widening, and many may not catch up. *Health Affairs*, **31**(8), 1803–1813.
- Park, J.H. (1999). Estimation of sheepskin effects using the old and the new measures of educational attainment in the current population survey. *Economics Letters*, **62**(2), 237–240.
- Rostron, B.L., Boies, J.L. & Arias, E. (2010). Education reporting and classification on death certificates in the United States. *Vital and Health Statistics*, **2**(151), 1–21.
- Ryan, C.L. & Bauman, K. (2016). Educational attainment in the United States: 2015 (current population reports). *Current Population Reports, US Census Bureau*.
- Sasson, I. (2016). Trends in life expectancy and lifespan variation by educational attainment: United States, 1990–2010. *Demography*, **53**, 269–293.

- Sorlie, P.D. & Johnson, N.J.** (1996). Validity of education information on the death certificate. *Epidemiology*, 7(4), 437–439.
- US Census Bureau** (2000). Current population survey design and methodology, Technical Paper 63.
- US Census Bureau** (2006), 'Current Population Survey Design and Methodology, Technical Paper 66'.
- US Department of Health and Human Services** (1995). Technical appendix from vital statistics of United States. *National Vital Statistics Reports*, p. 15.
- US Department of Health and Human Services** (2004), The health consequences of smoking: a report of the surgeon general.
- Zibbell, J.E., Asher, A.K., Patel, R.C., Kupronis, B., Iqbal, K., Ward, J.W. & Holtzman, D.** (2018). Increases in acute hepatitis C virus infection related to a growing opioid epidemic and associated injection drug use, United States, 2004 to 2014. *American Journal of Public Health*, 108(2), 175–181.