

The value of longevity

Politics, Philosophy & Economics

1–19

© The Author(s) 2019

Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/1470594X19880279

journals.sagepub.com/home/ppe**Greg Bognar***Stockholm University, Sweden and La Trobe University, Australia*

Abstract

Longevity is valuable. Most of us would agree that it's bad to die when you could go on living, and death's badness has to do with the value your life would have if it continued. Most of us would also agree that it's bad if life expectancy in a country is low, it's bad if there is high infant mortality and it's bad if there is a wide mortality gap between different groups in a population. But how can we make such judgments more precise? How should we evaluate the harm of mortality in a population? Although philosophers have written a lot about the harm of death for individuals, very little work has been done on the harm of mortality for populations. In this article, I take the first steps towards developing a theory of the harm of population mortality. Even these first steps, I argue, lead to surprising results.

Keywords

age discrimination, comparativism, harm of death, life expectancy, population mortality, value of life

Introduction

Longevity is valuable. As long as life is not an unbearable burden, we believe it is good to go on living. We hold that it is a great misfortune to die young. We think it is better to pass away after a long and fulfilling life, rather than in your prime. And we believe what is true of individuals is true of populations. We consider it a bad thing that life expectancy in war-torn Afghanistan is almost 25 years less than in affluent Japan. We agree that it is bad that a newborn African American boy in the US can expect to live over 16 years less than a newborn Asian American girl. And we believe it is tragic that in most of sub-Saharan Africa, almost one in 10 children die in the first year of their life. But

Corresponding author:

Greg Bognar, Department of Philosophy, Stockholm University, 10691 Stockholm, Sweden.

Email: greg.bognar@philosophy.su.se

we agree that it is good that infant mortality has halved in that region in the last 25 years.¹

These judgments are straightforward; they are almost like platitudes. But other judgments are more difficult to make. Most of us would be unable to say just *how bad* the low life expectancy of the Afghan population is. Should it be compared, say, to Japan's? If so, how much worse off the population of Afghanistan is in terms of longevity than the population of Japan? Or should the harm of mortality be compared to some other baseline? What should that baseline be? And how should we make the comparison?

Population mortality raises many puzzling questions. Some of these are practical. If you have to decide how to allocate scarce medical resources between, for instance, treatment of malaria and HIV/AIDS, one thing you need to know is the harm of premature mortality due to these conditions. When governments consider introducing road safety regulations, they need to consider how many years of life are lost due to fatal road traffic accidents. And when it comes to environmental policy, determining the harm of premature mortality is indispensable. For instance, one of the harms of climate change is that it kills people through its various effects. You can't respond to climate change effectively unless you can evaluate policies in terms of their impact on averting premature death.

But questions about the value of longevity are also philosophical. In many societies, the number of old people has been increasing as people live longer and have fewer children. Gains in longevity are likely to continue and people in future generations will live longer – and perhaps substantially longer due to biomedical breakthroughs. Is there a point at which increased longevity ceases to be a benefit? Does an extra year of life have the same value at every age, regardless of whether you are young or old?

It is astounding therefore that longevity and population mortality have been almost completely ignored by philosophers.² There has certainly been a fair amount of discussion on the evil of death, but it has focused only on the case of individuals; the harm of premature mortality for populations has received little attention.³ Perhaps philosophers have thought that once you know how to evaluate the harm of death for individuals, all you have to do is to count up those individual harms and the sum will give you the harm of premature mortality for the population. Even if it's difficult to do this in practice, there is no theoretical problem.

One reason this 'scaling up' strategy doesn't work is that we don't have a good theory for measuring the harm of death for individuals. There are different, competing accounts, but they all have their own problems, and there seems to be no reason to choose one over the others. Consequently, the scaling up strategy has difficulties getting off the ground. It is therefore worth tackling the problem of the harm of population mortality directly. If a plausible account for the case of populations can be developed, it might be used to make progress for the case of individuals.

I begin by explaining the problem for the case of individual mortality in the next section. After that, I take the first few steps of developing a theory of the harm of population mortality. I show that even these first steps lead to surprising results.

The harm of mortality for individuals

According to the prevalent view of the evil of death, death is bad for the person who dies because it deprives her of further goods of life. On this view, you compare how good overall a person's life would be if it ended at some time with how good overall it would be if it continued. Thus, the harm of death for person p at time t is determined by the difference between the value of p 's life if she died at t and the value of her life were she to continue to live. This view is known as *comparativism*.

Many philosophers accept some variant of comparativism.⁴ Even some of its main rivals share its basic idea. For instance, some prefer *gradualism*, on which the harm of death for a person p at time t is determined by the difference between the value of p 's life if she dies at t and both the value of her life were she to continue to live and the degree of psychological unity relations between herself at t and her future selves. Gradualism extends comparativism to what matters with respect to personal identity.⁵ For my present purposes, I'm going to adopt comparativism. Gradualism would introduce too many complications to my argument. However, it would change the conclusions of this article only in minor ways. The account I develop can be reformulated in gradualist terms.

Although many philosophers agree on the general idea of comparativism, there is no agreement on the details. In particular, there is no agreement on how to make the comparison between the value of the person's actual life and the value of her counterfactual life. Which of the many ways that p 's life could have gone should be the basis of comparison for determining the magnitude of the harm of her death at t ? Since the comparison is between the value of a person's life in the actual world and the value of her life in some possible world, I call this the *contrast world problem*.

There are various proposals for specifying the contrast world. One is to hold that the person's life should be compared to her life in the *nearest possible world* in which the event that caused her death does not occur.⁶ However, there are several problems with this idea. For one thing, a death may be caused by a chain of events, such that if any of the events did not occur, the death would not have occurred. But it's often not clear how to decide which of these events you should imagine did not occur. And depending on your choice, different possible worlds will be specified as the contrast world, with no apparent reason to choose between them.⁷

Even if this problem is set aside, another remains. Sometimes the nearest possible world does not seem to be the appropriate contrast world. Jeff McMahan imagines a young pedestrian who absentmindedly steps off the curb and is killed by a bus. It is discovered during the autopsy that he had a cerebral aneurysm that would have certainly ruptured within a week. If the accident did not take place, the pedestrian would have died soon anyway; but the accident and the rupture are not part of the same causal chain of events. Now if the rupture would have happened in the nearest possible world, it seems this person's death is not particularly harmful for him. Yet his death was surely a great harm for him. Consequently, the nearest possible world may not always be the appropriate contrast world.⁸

It seems that the appropriate contrast world for the young pedestrian is one in which he survives to old age (even if that would not have happened in a nearby or probable possible world). Often, it seems that a person's death should be compared to what could

be *expected to happen* to the sort of person he is. But it's hard to make this idea more precise. As McMahan concludes, there is no rationally required solution to the problem of specifying the contrast world.⁹

Consequently, in the absence of a solution to the contrast world problem, the 'scaling up' strategy that starts from the harm of individual mortality and aggregates it for all individuals in the population cannot get off the ground. It cannot be used for determining the harm of population mortality. If there is no uniquely rational account of the contrast world in the case of individuals, the harm of individual mortality cannot be used to determine the harm of mortality for populations. You may as well begin from the population case. A theory for the harm of population mortality might, as a bonus, have some lessons for the case of individuals too.¹⁰

The harm of mortality for birth cohorts

For the sake of simplicity, I am going to first consider only *birth cohorts* – a group of people born at a particular place in a given time period. For example, the group of people born in Tokyo in 1919 constitute one birth cohort; the group of people born in Afghanistan in 1969 constitute another. Evidently, every population is made up of several birth cohorts.

As I will show, once you have an account of the harm of mortality for birth cohorts, it is straightforward to extend it to populations. It is just a matter of adding up the harm of mortality for all the birth cohorts in the population.

I begin by making some simplifying assumptions. First, I take a year as a unit of time. This is purely for convenience. Second, I assume that all the people in a population have a life worth living from birth to death. All years of life are worth living. Third, I assume that all people in all the years of their lives have the same level of well-being. Thus, each person's life is equally good at all times, and the only thing that makes a difference to the value of a life is its length. Overall lifetime well-being is a linear transformation of the number of years one lives. These assumptions can be relaxed. But for now they will serve as convenient starting points.¹¹

If you list every person and their age of death in a birth cohort, you get the *mortality profile* of that cohort. The mortality profile can be represented by a *survivorship function*. For an example, see Figure 1. The horizontal axis shows age; the vertical axis shows the percentage of the initial birth cohort that is still alive as time passes. Everyone in the cohort is alive at year zero, and no one is alive by the hundredth year. Life expectancy at birth equals the area under the curve.

The central idea of comparativism is that the badness of death is determined by comparing a person's actual life to the life that person would have in the contrast world. The harm of mortality is determined by the difference between the values of these lives. The account of the harm of population mortality that I propose extends this idea first to birth cohorts and then to populations. It compares the actual mortality profile of a birth cohort to the mortality profile it would have in a contrast world. So, the first question to address is this: what is the appropriate contrast world for a whole birth cohort?




Figure 1. A survivorship function.

To examine some possible answers to this question, it is useful to borrow some ideas from demography. One way to calculate the harm of premature mortality is on the basis of *potential years of life lost*. Formally:

$$\sum_{x=0}^L d_x(L - x)$$

where x is age, d_x is the number of people who die at age x and L is a specific number of years – an *ideal age* to which we compare people’s actual deaths. Thus, when a person dies, the harm of her death is the gap between L and the age at which she dies – the potential years of life she loses. To determine the harm of mortality for a birth cohort, you aggregate the harms of death for all individuals in that cohort.

There are different accounts for specifying L . One is to equate it with life expectancy at birth. After all, life expectancy at birth is often used in practice as one indicator of human development.¹² It can be interpreted as a measure of longevity.

For an example, consider Figure 2. In this birth cohort, life expectancy at birth is 55 years, as shown by the vertical line. A person dies prematurely when she dies before reaching this age, and the harm of her death is given by the difference between her age and life expectancy at birth. For instance, if you die at 25, the harm of your death is proportional to 30 years; if you die at 40, it is proportional to 15 years. The overall harm of mortality in this cohort can be represented by the shaded area over the survivorship function. On this view, the contrast world is one in which everyone reaches at least their life expectancy at birth.

The ideal age can be given other interpretations. For instance, L may be equated with the age that is necessary for a *complete life*. The idea of a complete life has been suggested in discussions of distributive justice: it is the lifespan that is sufficient for experiencing all the different goods of life, including those associated with growing up, embarking on a career, gaining recognition for one’s accomplishments, starting a family




Figure 2. Potential years of life lost in terms of life expectancy at birth.

and seeing one's children grow up and start families of their own. It has been suggested that 70 years might be sufficient for a complete life.¹³ On this view, the harm of mortality for a birth cohort would be given by the area over the survivorship function extending out to 70 years of age on the right side of the figure. Thus, the contrast world is one in which everyone is given a chance to live a complete life by surviving at least to this threshold. The harm of mortality will be greater than on the life expectancy at birth view as long as the complete life threshold is greater than life expectancy at birth.

On yet another view, the harm of mortality would be even greater. The ideal age might be equated with the *natural human lifespan*. Perhaps there is a biologically determined age beyond which very few people survive, but which most people could reach given the right social and environmental conditions. The contrast world is one in which everyone reaches at least the end of the natural human lifespan.¹⁴

These are some of the views that are based on calculating the harm of mortality on the basis of potential years of life lost. Extending them from birth cohorts to whole populations is straightforward: you calculate the harm of mortality for each birth cohort whose members are present in the population and aggregate the harms. You can take all the deaths occurring at a time to arrive at a 'snapshot' view of the harm of mortality at that time – in a given year, for example. It is then straightforward to extend the calculation over time: you aggregate the harm of mortality in each year for the extended period you are interested in.

These views specify a threshold below which death harms you, but above which it is no harm – more precisely, above which your death *does not contribute* to the harm of population mortality. Hence, these views suffer from the same problem: they ignore the harm for those who exceed the threshold. Once you have reached life expectancy at birth, or the age which is sufficient for a complete life, or the natural human lifespan, your death does not count. But this is implausible. It is reasonable to insist that your death




Figure 3. Life expectancy profile.

matters even if you are one of the lucky few who exceed the natural human lifespan, or the complete life threshold, or life expectancy at birth. Since death can still harm you over the threshold, it should contribute to the harm of population mortality. For this reason, we should reject these views.

There is another way of calculating the harm of premature mortality in demography. It is based on the idea of *expected years of life lost*. Formally,

$$\sum_{x=0}^k d_x e_x$$

where x is age, d_x is the number of people who die at age x , e_x is life expectancy at age x and k is the greatest age to which the last person in the birth cohort survives. Note that e_x varies for different values of x . This is because life expectancies change as people move through life due to the fact that some members of their birth cohort have already died. Figure 3 shows life expectancies at different ages for the example cohort that I have been using. They define the cohort's *life expectancy profile*.

The harm of birth cohort mortality on this kind of view is given by the area between the survivorship function and the plot of life expectancies at different ages. It is the gap between the cohort's mortality profile and its life expectancy profile. Thus, the contrast world is such that everyone survives to the age they would have expected to, had they not died at the time they did. Given the mortality and life expectancy profiles of the birth cohorts that make up the population, it is straightforward to calculate the harm of population mortality. And just as before, it is straightforward to aggregate it for longer time periods.

One advantage of this kind of view is that all deaths count. There is no threshold over which a death does not contribute to the harm of population mortality. Even the very last person in a birth cohort has some probability of surviving to an additional period, hence her death will contribute to the overall harm of mortality.

Age	Life expectancy
0	86.02
1	85.21
5	81.25
10	76.27
15	71.29
20	66.35
30	56.46
40	46.64
50	37.05
60	27.81
70	18.93
80	10.99
90	5.05
100	2.23

Figure 4. Life table from the Global Burden of Disease studies.

Another problem, however, remains. The life expectancy profiles of two cohorts are likely to be different. Normally, people born at a particular time and place have different life expectancies than people born earlier or later in time or at different places. This is going to be equally true of cohorts in the same population and cohorts in different populations. Measuring the harm of population mortality on the basis of different life expectancies introduces unfairness. Why should the harm of mortality in one population be smaller than in another just because its members have worse life prospects? Why should a particular person's death *matter less* just because she happened to be born into less fortunate circumstances?¹⁵

This is a problem that demographers have already faced. The United Nations Population Division publishes *life tables* for different regions, on the basis of which years of life lost can be calculated uniformly in those regions, even when the data are incomplete.¹⁶ (A life table describes a life expectancy profile. It is a list of ages and life expectancies at those ages. The latter are called *period life expectancies*.) And the Global Burden of Disease project attempts to quantify the harm of disease, disability and premature mortality across the globe. Its measure is called the DALY, or *disability-adjusted life year*. DALYs are made up of two components: *years lived with disability* and *years of life lost* due to premature mortality. The DALYs associated with a particular health outcome are the sum of the years of life lost and the years lived with the disability that is associated with that health outcome.¹⁷

To measure years of life lost, the Global Burden of Disease project uses a uniform life table. The harm of every death is calculated by subtracting the age at which it occurs from the life expectancy associated with that age. Figure 4 presents an abridged version of the life table from the most recent version of the Global Burden of Disease studies.¹⁸

Thus, if a person dies at 20, the years of life lost are 66.35 years. If she dies at 60, the years of life lost are 27.81 years. This is regardless of the birth cohort or population to which she belongs. All deaths at a given age, no matter where and when they occur, make the same contribution to the harm of premature mortality. They all matter equally.

This view is based on expected years of life lost, but it uses standardized, or, as they are also called, *ideal* life expectancies. Formally,

$$\sum_{x=0}^k d_x e_x^*$$

where e_x^* is ideal life expectancy at age x , and the other terms are as defined above.

Ideal life expectancies solve the problem of unfairness that occurs if different life expectancy profiles are used for different birth cohorts and populations. All deaths count towards the harm of population mortality, and all deaths at a given age count equally. An ideal life table, such as the one used by the Global Burden of Disease studies, describes in minimal but sufficient detail the contrast world to be used for calculating the harm of population mortality.¹⁹

Note that this approach has no difficulty incorporating some of the theories I discussed above. If you find attractive the idea of calculating the harm of mortality in terms of complete lives or the natural human lifespan, you can interpret them in terms of expected years of life lost. There is no reason to hold that a complete life or the natural human lifespan must be defined in terms of rigid thresholds. These formulations would avoid the problem that some deaths do not contribute to the harm of population mortality.

This leads to an obvious question. Where do the numbers in the ideal life table come from? If you take the natural human lifespan view, you might turn to biology or medicine to fill in the cells. If you take the complete life view, you can define them in terms of the time needed for a chance for a complete life, taking into account that life plans are open-ended: even those who have reached old age tend to have plans and projects for the future, even if they are modest. For instance, people at the end of life often prefer to survive just long enough to put their affairs in order, or to live long enough to see a grandchild graduate or get married. It is a mistake to think that a complete life has to have a rigid threshold. As Thomas Nagel put it:

A man's sense of his own experience . . . does not embody this idea of a natural limit. His existence defines for him an essentially open-ended possible future, containing the usual mixture of goods and evils that he has found so tolerable in the past.²⁰

The Global Burden of Disease project built its life table on the basis of actual life expectancy data. Life expectancies are taken from wherever they are the highest in the world. The argument for constructing the life table this way is that if some societies are able to reach these life expectancies, then ideally all societies should be able to reach them too. The harm of mortality should be calculated on the basis of whatever is possible in some part of the world. No death, whether it occurs in Western Europe or sub-Saharan Africa, should matter less just because local life expectancies differ.

x	n_x	d_x	e_x	n_x	d_x	e_x	n_x	d_x	e_x
0	100	20	50	100	30	40	100	10	60
20	80	20	40	70	25	32.8	90	15	45.5
40	60	20	30	45	20	25.5	75	20	32.6
60	40	20	20	25	15	18	55	25	20.9
80	20	20	10	10	10	10	30	30	10
100	0	0	—	0	0	—	0	0	—
	<i>Leibniz</i>			<i>Huygens</i>			<i>Halley</i>		

Figure 5. Three contrast worlds.

The harm of mortality for populations

I have argued that the harm of population mortality should be calculated on the basis of expected years of life lost – in particular, in terms of *ideal* expected years of life lost. The appropriate basis of comparison is provided by an ideal life table. The ideal life table defines a contrast world in minimal but sufficient detail. Its period life expectancy data can be filled in on the basis of the most suitable comparativist theory of the harm of death.

I will not argue for any particular theory. Nor will I attempt to provide specific figures for period life expectancies in the ideal life table. I do, however, need to address some general questions concerning the life expectancy profile in the ideal life table. What are some of the properties of the contrast world that it defines? As I will put it, what should the contrast world *look like*?

To address these questions, I introduce some simple examples. I imagine three possible contrast worlds. They are illustrated in Figure 5. To make the discussion more vivid, I will speak of them *as if* they represented the actual mortality profile of a birth cohort – rather than possible ideal life tables. This will help draw out my point. Ultimately, what matters for the calculation of the harm of population mortality for any population are only the period life expectancy data included in them.

Let us suppose that each of these three worlds (called ‘Leibniz’, ‘Huygens’ and ‘Halley’) has only one birth cohort with a hundred people. As before, x is age. (In these worlds, no person is alive by their hundredth birthday.) n_x is the number of people who are alive at the beginning of a time period (or age). d_x is the number of people who die during a time period. e_x is life expectancy at the beginning of a time period.

In Leibniz, 100 people are alive at the beginning of the first year (or age 0) and 20 people die between years 0 and 20. Only 80 people are alive by year 20. In the next period, another 20 people die, so that by year 40, only 60 people are alive; by 60, only 40 are. Life expectancy at birth is 50 years; at year 20, life expectancy is 40 years, and at 40, it has fallen to 30, and so on.

Huygens and Halley are different. In Huygens, 30 people die by year 20, and life expectancy at birth is only 40 years. In Halley, 10 people die by year 20, and life expectancy at birth is 60 years.

The crucial difference between these three worlds is the way the number of deaths changes as time passes. In Huygens, the number of deaths (or the crude death rate) is decreasing: many people die in the first period, and fewer and fewer die in later periods. In Halley, the death rate is increasing: few people die in the first period, and more and more people die in later periods. In Leibniz, in contrast, the death rate is constant. The same number of people die in each period.²¹

Because of these patterns, the worlds have different properties. These properties make a difference to the way the harm of mortality is calculated. Therefore, the worlds can be evaluated on their basis. The properties can serve as criteria for assessing their suitability as contrast worlds. However, what matters to these criteria are not the particular period life expectancy data, but their pattern; what matters is what the contrast world *looks like*. Merely on this basis, it is possible to derive some conclusions for the theory of the harm of population mortality. Let me explain.

Start with some examples. One way to compare the three worlds is by looking at *life expectancy* at each period. Plainly, Halley does best in this respect. At each period except for the very last, it has a greater life expectancy than Leibniz, which in turn has, with the same exception, greater life expectancies than Huygens. On this measure of longevity, Halley weakly dominates the others.

You can also look at *population size* at different times. Given the same initial population size, it's better if the surviving population is larger at later times. Again, Halley does best in this respect. It has more people surviving to each age before the very last than Leibniz, which has more people surviving to each age before the last than Huygens. In addition, save for the last period, the *proportional death rate* – the ratio of the number of deaths in a given period and the population size at the beginning of that period – is smaller in Halley than in Leibniz, which has a smaller rate than Huygens in the same way. Halley emerges weakly dominant on these criteria too.

There are two more measures of longevity to consider. You can look at the *median age of death* – the age at which half of the population is already dead. You can also look at the *modal age of death* – the age at which the greatest number of people die. In Huygens, the median of age death is sometime before year 40. In Leibniz, it is sometime between year 40 and 60. In Halley, it is sometime after year 60. Once again, Halley is ranked highest on this measure, followed by Leibniz and then Huygens. Furthermore, the greatest number of people die before year 20 in Huygens, whereas the greatest number of people in Halley die in the last period, after year 80. Arguably, Halley does best on this measure too.

Plainly, Halley is best from these three possible contrast worlds. It does better than any of the other two on all the measures of longevity that I have considered. It would be rational to choose Halley if you were a self-interested decision maker behind the veil of ignorance and you had to choose which of these three populations you belong to when the veil is lifted. You would choose Halley if you were a benevolent impartial spectator concerned with maximizing the well-being of a population. Thus, the ideal life table that is used to calculate the harm of mortality for any population should have these properties too. Since longevity is valuable, it should do well on all of its standard measures.²²

We can conclude the following. The ideal life table that is used to determine the harm of mortality for populations should have certain properties. These properties are a

function of the pattern of the period life expectancies. In particular, the *death rate* – the number of people who die in each period – should be increasing to ensure the best life expectancies, population size, proportional death rate, as well as median and modal age of death. There might be further criteria that the ideal life table should have. But, other things being equal, it should score as well as possible on these measures. The contrast world should look like Halley.

The value of longevity

Nevertheless, adopting a world like Halley as the contrast world comes at a cost. To see this, consider the following example.

Suppose you have to choose between saving the life of a 20-year old person who would survive for another 20 years, on the one hand, and saving the life of a 60-year old person who would survive for another 20 years, on the other. That is, the 20-year old would die when she is 40; the 60-year old would die when he is 80.

If you interpret the three simple worlds I introduced in the previous section as a description of the mortality profile of birth cohorts, then the choice might be between saving a person's life now compared to saving the life of another person in the future. Only one of these acts can be performed. I assume there is no discounting of the value of life years, hence it makes no difference that one of the acts is further in the future. Suppose also that your aim is to minimize the harm of mortality.

What would be the harm in Leibniz? You can answer this question by looking back at Figure 5. If you choose to save the 20-year old, then you would have one person who survives until she is 40, and another who dies when he is 60. The harm of mortality for each is determined by the difference between their death at these particular ages and their ideal life expectancies at those ages. The overall harm of mortality is the sum of the harms of their deaths. That is, the harm of mortality for the person who dies at 40 is 30 years (since that is her life expectancy at that age). The harm of mortality for the person who dies at 60 is 20 years (his life expectancy at that age). The overall loss is 50 years. This is the harm of population mortality if the younger person is chosen.

Things are similar if the 60-year old person is saved. In this case, there is one person who dies at 20, and the loss is 40 years; and there is another person who dies at 80, with the loss of 10 years. The overall loss is 50 years. The harm of mortality is the same in both cases.

Consider Halley now. If the 20-year old is saved, the overall harm of mortality is 53.5 years: one person dies at 40, losing 32.6 years, and another person dies at 60, losing 20.9 years. If, in contrast, the 60-year old is saved, you have one person who dies at 20, with a loss of 45.5 years, and another person who dies at 80, with a loss of 10 years. The overall loss is 55.5 years. It is greater than in the previous case. Consequently, saving the older person leads to a *worse* outcome – for the smaller the harm of mortality, the better.

Therefore, it is better to save the younger person, since the overall harm of mortality is smaller in this case. This is despite the fact that the same benefit would be provided to the two people – 20 years of additional life. But in Halley, this benefit has greater value if the beneficiary is younger. That is, additional years of life have greater value if they go to younger people. Longevity has decreasing marginal value. Halley is better

on all the measures of longevity that I have considered, but it discriminates by age, favouring the young.

Age discrimination might be considered another criterion for evaluating contrast worlds. You might think that it is a desirable property of an ideal life table that it does not imply age discrimination. For instance, similarly to Halley, Huygens also implies age discrimination – although it is in the reverse direction. On Huygens, additional years of life have greater value if they go to older people. There is discrimination against the young. I take it that few people would agree that this is a desirable property. Our account of the harm of population mortality should not lead to discrimination against younger people.

Alternatively, you may start out by being uncertain whether age discrimination is unfair. My argument can help settle the question. Huygens has no criteria to recommend it, hence it cannot be used to justify discrimination against the young. But Halley does best on all the criteria that I have considered. All things considered, it is the easiest to justify. In the light of this, you may come to believe that it provides an argument for age discrimination. It implies giving greater value to additional years of life for younger people. This form of age discrimination, therefore, is not unfair.

Furthermore, my argument may also provide a solution to the harm of death in the individual case. It is determined by the gap between the age of the person who dies and his ideal life expectancy at that age. Thus, the death of McMahan's young pedestrian was indeed harmful for him, even though he would have died a week later. The harm is determined by the gap between his age at the time he dies and his ideal life expectancy at that age. The theory of the harm of population mortality yields an answer to the question of the harm of individual mortality. It is able to capture the intuitive idea of what is tragic about the death of the young pedestrian: he died before he could have expected to. As I said earlier, intuitively it seems that the harm of a person's death should be determined by a comparison between the value of his life given his actual death, and the value of his life given what could be expected to happen to the sort of person of he is. My proposal makes this idea clearer by specifying a contrast world in which everyone would survive to the age they would have expected to, had they not died at the time they did.

Remember that the basic difficulty for theories of the harm of death in the case of individuals is that there seems to be no rationally required solution to the problem of specifying the contrast world.²³ However, *if* you start from the case of the harm of mortality for populations, there is arguably a rationally required solution: the contrast world should have a specific pattern – one that fares best on the measures of longevity that I discussed. These can serve as criteria of what the contrast world should look like. Consequently, my proposal can not only provide answers to questions about the harm of population mortality but it may also offer an account of the harm of death for individuals.²⁴

Implications

Before I conclude, it is worth briefly highlighting some of the further implications of the proposed theory and to put it into a broader context.

First, it should be pointed out that this is a normative theory of the harm of mortality – it operates with a moralized notion of harm in the sense that it specifies the contrast world on the basis of normative criteria. This is not a shortcoming, but a deliberate feature. As McMahan’s example of the young pedestrian shows, at least sometimes our intuitions about the harm of death are governed by considerations about what we take to be the *appropriate* basis of comparison, rather than considerations of what would have happened. The intuitions about appropriate contrast worlds are normative. The theory put forward in this article makes the normative nature of the comparisons explicit.

At the same time, the theory remains neutral among different substantive proposals about the contrast world. What I call substantive proposals provide the numbers to be used in the ideal life expectancy profile. As I have already mentioned, they may, for instance, rest on some view about a complete life, the natural human lifespan or some particular theory of well-being. The theory is agnostic among such proposals; you may plug your favourite into it. The only constraint is that the numbers that your favourite proposal yields must make up an ideal life expectancy profile that does as well as possible on the criteria I introduced.

Second, it should also be noted that my example in the previous section ignored important complications. I assumed that the only thing that mattered in the choice between saving the 20-year old and the 60-year old was the minimization of the harm of mortality. In real life, many further factors are relevant in the allocation of life-saving resources. Longevity is not the only value that matters. Thus, it is important to distinguish the narrower issue of the measurement of the value of longevity and the broader issue of allocating resources for extending lives. Resource allocation choices usually involve difficult trade-offs between different values or moral considerations, and I don’t mean to suggest that the value of longevity is the only one that matters. Perhaps, there are a few cases when only longevity is relevant. On the other hand, when it comes to allocating life-saving resources, there are only a few cases when longevity is irrelevant.²⁵

Be that as it may, the theory of the harm of population mortality presented in this article has several advantages over rival views of resource allocation. I have already mentioned some of them. It doesn’t operate with a crude, fixed threshold beyond which deaths do not contribute to the harm of mortality – or, alternatively, beyond which the benefits of life extension make no moral difference.²⁶ It also avoids the unfairness inherent in valuing life extensions in different populations due to the fact that the life expectancy profiles of different cohorts are different. While the first problem is avoided by ‘age-weighting’ views – theories of resource allocation that give decreasing weights to benefits as age increases – even these views face the second.²⁷ On an ‘age-weighting’ view – other things being equal – saving the life of a 60-year old in a cohort or population with greater life expectancy will be given priority over saving the life of a 60-year old in a cohort or population with smaller life expectancy. In contrast, the view presented here (again, other things being equal) gives the same weight to saving either of these 60-year olds.

Of course, other things are usually not equal, and the view can take account of this. In fact, it will *not* always give more weight to benefits to younger people. Suppose that

members of an earlier birth cohort have smaller life expectancies than members of a later cohort. Therefore, more people die at age 60 (say) in the earlier cohort than in the later cohort. Even though the death of any single individual from either of these cohorts contribute the same amount to the harm of population mortality, the overall harm is greater for the earlier cohort than for the later cohort; *as a group*, the earlier cohort contributes a greater amount to the harm of population mortality than the later cohort. This is easy to see if it is expressed numerically:

$$d_{60}^a e_{60}^* > d_{60}^b e_{60}^*$$

where d_{60} is the number of deaths at age 60, and the superscript a denotes the earlier cohort and the b denotes the later one. Since there are more deaths in cohort a at this age, the harm of mortality is greater in this group. This kind of inequity is readily apparent on this view. It is something that social policy may want to take into account. For instance, additional resources may be allocated to reduce premature mortality in cohort a . The theory of the harm of population mortality that I propose can take into account demographic change.

Conclusion

In this article, I set out to develop a theory for the harm of mortality for populations. I argued that we need a theory to answer both practical and philosophical questions. I accepted comparativism as a general theory for the harm of mortality, but I argued that you cannot simply ‘scale up’ any account of the harm of mortality for individuals. This is because there is no single account that adequately specifies the contrast world. I proposed a different approach to the problem, starting from an account of the harm of mortality for birth cohorts. I argued that any account must be formulated in terms of expected years of life lost, and I showed how it can be extended from birth cohorts to populations. I concluded that the harm of population mortality should be measured on the basis of an ideal life table.

I then set out to examine what the ideal life table should look like. Using simple examples, I constructed a number of criteria that can be used to evaluate different kinds of ideal life tables. But I discovered that the kinds of life table that do best on these criteria have a surprising implication: they entail that additional years of life have different values depending on the age of the beneficiary. In particular, the value of additional years of life diminishes with increasing age. I conclude by suggesting that at the end of the day, we should accept this implication and retain the theory, even though this would have a profound effect on our policies. Longevity, it turns out, has decreasing marginal value.

Acknowledgements

For comments on earlier versions of this article, the author would like to thank audiences at Bayreuth University, Linköping University, Oxford University, the Université Catholique de Louvain, the University of Helsinki, Uppsala University, as well as participants at the *Population and Ethics: Interdisciplinary Perspectives on Birth and Death* conference at Cumberland Lodge, Windsor, UK, at a *Workshop on the Ethics of Age Limits* at

the Institute for Futures Studies in Stockholm and at a *TADS Workshop* at the Institute of State and Law in Prague. The author would also like to thank Karl Pettersson, Shlomi Segall and Bruno Verbeek for detailed comments.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This research was supported by the Swedish Research Council (VR) (grant number 2018-01861) and the ‘Taking Age Discrimination Seriously’ (TADS) project, based at the Institute of State and Law (Prague) and funded by the Grantová Agentura České Republiky (grant number 17–266295).

Notes

1. See Murray et al. (2006) and Jamison et al. (2006).
2. For example, a recent 500-page handbook on the philosophy of death (Bradley et al., 2013) does not have a single chapter on population issues of longevity and mortality. The word ‘population’ does not even appear in the index. The same is true of another recent, 350-page collection on the topic (Luper, 2014). The situation isn’t much better if you look elsewhere. For instance, a recent handbook on well-being (Fletcher, 2016) has no chapter on longevity.
3. To be fair, there has been some discussion of related issues in population ethics, starting from Parfit (1984). But that discussion has mainly centred on comparing the value of populations of different sizes (for example, Blackorby et al., 2005; Broome, 2004; Ryberg and Tännsjö, 2004). My focus in this article is different.
4. See, for instance, Nagel (1979), Silverstein (1980), Feldman (1991), Kamm (1993), Broome (1999) and Bradley (2009).
5. Gradualism is defended by McMahan (2002). A recent defence is Millum (2015).
6. The proposal is based on the account of counterfactuals developed by Lewis (1973).
7. For discussion, see, for example, Feldman (1991) and McMahan (2002: Ch. 2).
8. See McMahan (2002: 117). The argument can also be formulated in terms of the *most probable* possible world, which may or may not be the nearest possible world.
9. McMahan (2002: 116). Ben Bradely puts the point in a different way: ‘when it comes to determinations of *overall values* of particular events in people’s lives, we must simply abandon the idea that there are always unequivocal answers’ (Bradley, 2007: 122, his emphasis; see also Bradley, 2009: Ch. 2).
10. In what follows, I intend no controversial metaphysical claim about groups. In particular, I do not mean that populations are welfare subjects that can be harmed apart from their members. Premature mortality is a harm for a population the same way an epidemic is: it harms people to different degrees, and there is no harm for the population over and above these individual harms. Yet it makes perfect sense to say that the epidemic harmed the population. I talk about the harm of population mortality exactly the same way.
11. It goes without saying that the third assumption is a very crude simplification. In fact, a life with constant well-being may be worse than a life that, for instance, becomes progressively

- better, even if the overall amount of well-being is the same. The shape of a life may matter. (On this, see Velleman, 1991.) Perhaps, there is a way to ‘disperse’ such additional value and assign it to particular times. For a detailed discussion of this strategy, see Broome (2004).
12. For instance, in the Human Development Index. See UNDP (2013).
 13. See, for example, Harris (1985). For discussion, see Bognar (2015a).
 14. See Callahan (1995).
 15. Interestingly, a similar sort of unfairness has been noted in the individual case. When comparing the death of a 12-year old progeria patient and the death of a geriatric patient at the end of the maximum human lifespan, Bradley remarks: ‘we might want to rule out certain comparisons as being in some sense “unfair,” by determining rules governing which comparisons count as fair comparisons. For example, when comparing the misfortunes of [these patients], there would be something odd about comparing [the progeria patient’s] actual life with a life where he dies of a stroke a moment later, then comparing [the geriatric patient’s] actual life with a life she lives in a world where she doesn’t suffer the stroke, the aging process is reversed, and she lives another fifty years, and concluding that [her] death is worse than [his] death’ (Bradley, 2007: 122).
 16. See, for instance, UN (1982). For a more formal introduction to life tables, see also Keyfitz and Caswell (2005).
 17. For comprehensive overviews, see Murray (1996) and Murray et al. (2012); for a short introduction, see Bognar (2015b).
 18. See Murray et al. (2012).
 19. This is not to say that the Global Burden of Disease studies do not face philosophical and methodological problems. (For details, see Murray et al., 2002.) However, I set these aside here. In a companion paper (Bognar, forthcoming), I apply the arguments presented here to the Global Burden of Disease study.
 20. Nagel (1979: 9–10).
 21. The names of these worlds reflect these patterns. Gottfried Wilhelm Leibniz (1646–1716) thought that the number of deaths at each year of age remained constant. Christiaan Huygens (1629–1695) seems to have believed that the number of deaths decrease. And in the life table published by Edmond Halley (1656–1742), the death rate was increasing. For details, see Robine (2011).
 22. This claim isn’t about the numbers. After all, you could make death a greater harm for a person by assuming that she would have lived for a thousand years. But this would be pointless. When we say that the death of McMahan’s young pedestrian is tragic, we make a comparison to the lifespan that is normal and desirable for people of that person’s age, even though there are people who die at 40 and he might have been one of them. Similarly, when we evaluate the harm of premature mortality in a population, we want to make a comparison to what is normal and desirable for populations. And surely, low period life expectancies, a rapidly decreasing population size or horrific death rates are not normal or desirable (even if they do occur).
 23. See footnote 9.
 24. At this point, however, it is worth recalling the crude simplifying assumptions that I made earlier: that people have a life worth living from birth to death, and that all years of life have the same level of well-being. Once these assumptions are relaxed, things get much more complicated. I cannot here address the questions that relaxing these assumptions would raise.

25. For an early discussion of further factors that are relevant for the allocation of life-saving resources, see Rescher (1969). I should perhaps add that I disagree that all of the factors Rescher lists are morally relevant.
26. For such views, see, for example, Harris (1985) and Callahan (1995).
27. Examples of such views include the egalitarian version of ‘fair innings’ proposed by Williams (1997) and the prioritarian version of fair innings that I developed in Bognar (2015a).

References

- Blackorby C, Bossert W and Donaldson DJ (2005) *Population Issues in Social Choice Theory, Welfare Economics, and Ethics*. Cambridge: Cambridge University Press.
- Bognar G (2015a) Fair innings. *Bioethics* 29: 251–261.
- Bognar G (2015b) QALYs, DALYs, and their critics. In: Arras JD, Fenton E and Kukla R (eds) *The Routledge Companion to Bioethics*. New York: Routledge, pp. 44–55.
- Bognar G (forthcoming) Age and time in the measurement of the burden of disease. In: Eyal N, Hurst S, Murray C, Schroeder SA and Wikler D (eds) *Measuring the Global Burden of Disease: Philosophical Dimensions*. Oxford: Oxford University Press.
- Bradley B (2007) How bad is death? *Canadian Journal of Philosophy* 37: 111–128.
- Bradley B (2009) *Well-Being and Death*. Oxford: Clarendon Press.
- Bradley B, Feldman F and Johansson J (eds) (2013) *The Oxford Handbook of Philosophy of Death*. New York: Oxford University Press.
- Broome J (1999) Goodness is reducible to betterness: the evil of death is the value of life. In: *Ethics Out of Economics*. Cambridge: Cambridge University Press, pp. 162–173.
- Broome J (2004) *Weighing Lives*. New York: Oxford University Press.
- Callahan D (1995) *Setting Limits: Medical Goals in an Aging Society*. Washington, DC: Georgetown University Press.
- Feldman F (1991) Some puzzles about the evil of death. *The Philosophical Review* 100: 205–227.
- Fletcher G (ed) (2016) *The Routledge Handbook of Philosophy of Well-Being*. New York: Routledge.
- Harris J (1985) *The Value of Life*. London: Routledge & Kegan Paul.
- Jamison DT, Feachem RG, Makgoba MW, et al. (eds) (2006) *Disease and Mortality in Sub-Saharan Africa*. Second edition. Washington, D.C.: The World Bank.
- Kamm FM (1993) *Mortality, Mortality. Volume I: Death and Whom to Save from It*. New York: Oxford University Press.
- Keyfitz N and Caswell H (2005) *Applied Mathematical Demography*, 3rd ed. New York: Springer.
- Lewis D (1973) *Counterfactuals*. Oxford: Blackwell.
- Luper S (ed) (2014) *The Cambridge Companion to Life and Death*. Cambridge: Cambridge University Press.
- McMahan J (2002) *The Ethics of Killing: Problems at the Margins of Life*. Oxford: Oxford University Press.
- Millum J (2015) Age and death: A defense of gradualism. *Utilitas* 27: 279–297.
- Murray CJL (1996) Rethinking DALYs. In: Murray CJL and Lopez AD (eds) *The Global Burden of Disease: A Comprehensive Assessment of Mortality and Disability from Diseases, Injuries, and Risk Factors in 1990 and Projected to 2020*. Cambridge, MA: Harvard School of Public Health on behalf of the WHO and the World Bank, pp. 1–98.

- Murray CJL, Ezzati M, Flaxman AD, et al. (2012) Comprehensive systematic analysis of global epidemiology: Definitions, methods, simplification of DALYs, and comparative results from the Global Burden of Disease 2010 study. *The Lancet* 380. Available at: [http://www.thelancet.com/journals/lancet/article/PIIS0140-6736\(12\)61899-6/supplemental](http://www.thelancet.com/journals/lancet/article/PIIS0140-6736(12)61899-6/supplemental) (accessed 10 October 2019).
- Murray CJL, Kulkarni SC, Michaud C, et al. (2006) Eight Americas: investigating mortality disparities across races, counties, and race-counties in the United States. *PLoS Medicine* 3: 1513–1524.
- Murray CJL, Salomon JA, Mathers CD, et al. (eds) (2002) *Summary Measures of Population Health: Concepts, Ethics, Measurement, and Applications*. Geneva: World Health Organization.
- Nagel T (1979) Death. In: *Mortal Questions*. Cambridge: Cambridge University Press, pp. 1–10.
- Parfit D (1984) *Reasons and Persons*. Oxford: Clarendon Press.
- Rescher N (1969) The allocation of exotic medical lifesaving therapy. *Ethics* 79: 173–186.
- Robine JM (2011) Age patterns in adult mortality. In: Rogers RG and Crimmins EM (eds) *International Handbook of Adult Mortality*. Dordrecht: Springer, pp. 207–226.
- Ryberg J and Tännsjö T (eds) (2004) *The Repugnant Conclusion: Essays on Population Ethics*. Dordrecht: Kluwer.
- Silverstein H (1980) The evil of death. *Journal of Philosophy* 77: 401–424.
- UN (1982) *Model Life Tables for Developing Countries*. New York: United Nations.
- UNDP (2013) *Human development report 2013 technical notes*. United Nations Development Programme, New York. Available at: <http://hdr.undp.org/en/2013-report> (accessed 10 October 2019).
- Velleman JD (1991) Well-being and time. *Pacific Philosophical Quarterly* 72: 48–77.
- Williams A (1997) Intergenerational equity: an exploration of the ‘fair innings’ argument. *Health Economics* 6: 117–132.

Author biography

Greg Bognar is a Senior Lecturer in Practical Philosophy at Stockholm University, a Senior Researcher at the Stockholm Centre for Healthcare Ethics (CHE) and an Honorary Associate at the Department of Politics, Media and Philosophy at La Trobe University in Melbourne.