


ALPHA SOURCES

JULY, 2021



MERE MORTALS NO LONGER?

The evolution of mortality through the demographic transition is as close as we come to a deterministic process in the analysis of population dynamics. Science and technology have become increasingly better at keeping people alive, a benefit that still seems to drive the human experience to this day. It's possible to identify milestones through history such as the development of modern sanitation to defeat contagious air- and waterborne illnesses, the development of vaccines for specific illnesses, as well as overall technological development in the field of healthcare. It is a story

about pinning down the causality between rising national income and technological development and the improvement in the human living condition in the past 250 years. Researchers still debate the relative importance and merits of specific drivers, but it's possible to generalize, all the same. The story about of human mortality is contained in a few relationships, for the individual, between, and possibly within countries. It is a story about Nike swoosh-shaped, logarithmic and asymptotic curves, and the extent to which we observe deviations from such stylized relationships over time, and why.



THE BIG PICTURE

Angus Deaton's *The Great Escape* from 2015 is as good a polemic on the shifts in human living conditions through time as you'll find. The book makes with two contradicting points. It's never been a better time to be alive, *provided* you live in the developed world. This is an intuitive, and even trivial point, but important to stress at the outset. The divergence between high life expectancy in richest countries, and low life expectancy in poor countries remains a key feature of the overall human living condition. The fact that many countries—mainly in sub-Saharan Africa—are still stuck in something resembling a poverty trap is, in itself, an astonishing reality, given how far the richest nations have come. Deaton (2015) optimistically talks about catch-up via learning by doing, though evidence at the start of the 2020s suggests that significant and sustained divergences persist.

In a grander perspective, Deaton (2015) begins with the point that hunter-gatherers actually didn't have it so so rough, despite enjoying relatively short lifespans. In fact, the Neolithic revolution around 12000 BC, characterized by the transition from nomadic hunter-gathering to agricultural settlements, was initially associated with a *deterioration* in the human living condition. The book describes a significant reduction in life expectancy and life quality driven by famine, drought and diseases blighting humans' early attempts to settle down. In the now immortal words of Nassim

Taleb, it would seem that the hunter gatherers initially were anti-fragile, to an extent, while the early iterations of agricultural settlers were not.

As far as more recent history is concerned, Deaton (2015) relies on the standard demographic transition model to argue that a rapid decline in child mortality—driven by improvements in healthcare, nutrition, and disease prevention—and a general betterment in our ability to expand life spans later in life were key drivers of improving living conditions. Cutler, Deaton et al. (2009) dives deeper into the relative importance of these drivers, identifying three phases of the mortality decline through the demographic transition.

The first, from the middle of the 18th to the middle of the 19th century, in which improvements in mortality was driven by better nutrition, economic growth, and advances in public health initiatives. The second, in the final part of the 19th century, and into the 20th, is dominated by public policy on health. This effect was initially negative, due to high mortality in large cities, and the inability of public health measures to accommodate that. In the latter part of this period, however, public health made great strides in the areas of sanitation, nutrition and general bacterial health. Third, and finally, the period from the 1930s onwards, is described as the era of "big medicine" in which the development of vaccines and antibiotics are heralded as the two most significant innovations.



It is important to emphasize that this sequence is one of many currently battling for supremacy in the literature. It is impossible for me to pass judgement on the relative merit of each of the main factors put forward to explain how and why life expectancy has increased. It is fair to assume, however, that it was a combination of factors, which is exactly what Deaton et al (2009) suggest.

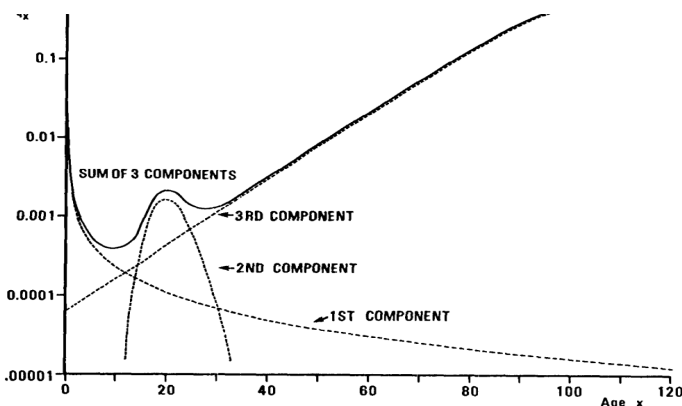
A (BRIEF) SURVEY OF THE THEORY

Our search for models of mortality that stand the test of time starts in the dusty world of actuarial science. It doesn't have to start there mind, but it is as good as place as any. Actuaries are in the business of putting numbers on risk, and one of the biggest risks, at any point in time, is the risk of death. Using post-war mortality data from Australia, Heligman and Pollard (1980) sets out to develop an "age pattern of mortality", or more specifically a continuous variable, that accurately reflects "the underlying mortality pat-

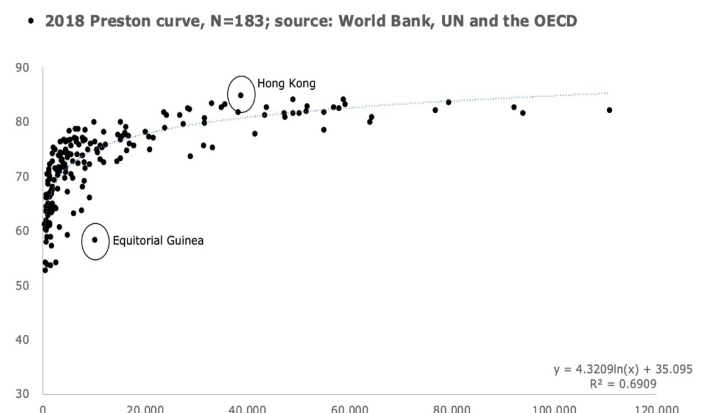
tern." In doing so, Heligman and Pollard (1980) extend an inquiry that started in the 17th century with the first life tables collected by John Graunt, in 1662, and Edmund Halley in 1693. The first formal law of mortality was proposed by Abraham de Moivre in 1725, though Heligman and Pollard (1980) credits Benjamin Gompertz' work in 1825 as the "the best known early contribution". As far as more recent, 20th century, contributions, Elston (1923) and Benjamin and Haycocks (1970) stand out.

The gist of the model is captured in the first chart below, reproducing the key figure from Heligman and Pollard (1980). It presents an intuitive model of mortality for the average individual over time, which resembles Nike's famous swoosh logo, or, for the economists, a variant of a J-curve. The beginning of life is precarious. Humans are at their most fragile when born, and it doesn't take much to extinguish life. Once this initial hurdle is conquered, however,

fig. 01 / The swoosh of human mortality - fig. 02 / The Preston curve



Source: Heligman and Pollard 1980, figure 1, p. 51
Y: mortality risk, X: age. PDF can be found [here](#).



Source: Author's own calculation. The original Preston curve comes from Preston (1975), p. 235. PDF can be found [here](#).



the risk of death declines steadily, hitting a trough at the age of 15, before an accident hump briefly raises mortality risk beyond what can be explained by aging. Finally, the probability of dying increases steadily towards 1 as the individual's age advance towards 100 and above. The overall curve is then a product of three distinct processes—each depicted separately — briefly sketched below.

The first, an asymptotic declining function, reflects the decline in mortality during infancy, which is a function of rising adaptability over time. Humans adapt quickly to their environment, and over time the infant-fragility mortality risk, and the risk of an existing environmental factor causing death, independent of senescence fragility, tends towards zero.

The second is the adult-accident hump, which temporarily raises mortality in later adolescence and into the 20s. For men, this is best thought of as risk factors involving violence and excessive risk taking, while for women, a temporary rise in mortality risk occur during child birth. Third and finally, senescence drives a linear increase in mortality as a function of age with a concave or logarithmic form—as the probability of death converges on 1—over time.

We always need to treat the idea of unbreakable laws with skepticism in the cross-section between evolutionary theory, biology and social sciences, but the idealized curve presented in Heligman and Pollard (1980) comes

close. Specifically, to the extent that we do *not* observe this relationship—in a given environment and population—it seems important to ask why. More generally, it's plausible that the shape of the curve changes over time, within and between countries and communities, and figuring out why is important.

The so-called Preston curve is another near-universal relationship in the analysis of mortality. It first appeared in Preston (1975) and compares life expectancy at birth with income. The chart in Preston (1975) plots this relationship for a sample of countries in 1900, 1930 and 1960, and is reproduced above. It proposes two key relationships, between countries and over time. First, it suggests that the positive relationship between income per capita and life expectancy is concave over time. The initial increases in income per capita convey a significant lift in life expectancy across countries, but then flatlines as income rises above a certain threshold. Secondly, the curve seems to be shifting up over time, reflecting a gradual increase in life expectancy across all countries. Indeed, a key test for this model is the extent to which the relationship can be replicated *over time*.

The original work by Preston (1975) suggests that it is. In the appendix I update the Preston curve over time with data from the World Bank, confirming the original model and intuition behind it. The cross-sectional relationship between income per capita—in constant prices—and



life expectancy is indeed concave, and stable over time. In addition, the curve also shifts higher in my sample, which changes over time, reflecting data availability, in the same way as the original chart in Preston (1975).

The appendix could easily form the basis of a separate paper, two observations stand out. The biggest outliers are countries with relatively low life expectancy compared to their income. Resource-rich economies stand out here. Additionally, even a casual look at the data hints at a structural break in the data across time. Beyond a certain level of income the relationship between life expectancy and rising income per capita is linear and strong, and after that it deteriorates, significantly. It is obviously the *combination* of these two statistical relationships that produce a concave, or logarithmic, function, but it's possible that the analysis should be separated across countries with different levels of income. Finally, it is *possible* that this relationship holds within countries too. Evidence presented in Muney and Moreau (2021), suggests that it does.

Infant mortality is an independent area of research that also deserves attention. It is effectively the study of what determines the slope on the first part of the mortality curve in Heligman and Pollard (1980). The study of child mortality is devoted a lot of attention in its own right, mainly because infant mortality is one of the key empirical lines of demarcation be-

tween non-developed and developed countries, or more specifically, between countries that are yet to start their demographic transition and those who have.

The theoretical and empirical literature on the determinants of child mortality is vast, but fortunately, Mosley and Chen (1984) has stood the test of time. The paper suggests two main innovations in the work on child mortality. First, it points out that social sciences and medical sciences tend to approach the question of child mortality in different ways. The former often draws a straight line from socioeconomic variables to mortality, omitting the specific health outcomes through which such variables operate to affect mortality. The latter, by contrast, focus solely on these proximate variables. The two, according to Mosley and Chen (1984) must be reconciled.

The paper begins its analysis with the assumption that 97% of children survive to their fifth year in an "optimal setting", implying a natural rate of child mortality—defined here as a fatality between year zero and five—of 3%. I have no a priori objection to this number, though it seems evident that it is subject to divergence across environments and social settings.

Mosley and Chen (1984) argues that the probability of child survival is a function of four broad independent variables; social, economic, biological and environmental. Socioeconomic factors, in particular, must be operationalized via their link with



proximate determinants. In turn, specific disease and other health deficiencies can be viewed as biological outcomes following from the proximate determinants. As far as the dependent variable is concerned, Mosley and Chen (1984) suggests the definition of a continuous variable of child survival and health, "growth faltering," is superior to a binary mortality indicator. This framework sounds reasonable to me though two broad qualifiers are needed. Firstly, the explanatory variables might not in fact be independent, indicating that researches must take care to set up the right-hand side of the model in a statistically appropriate fashion. Secondly, the choice of a continuous dependent variable is likely to be just as much a question of data variability as anything else.

Mosley and Chen (1984) identifies four broad major categories of proximate determinants.

- 1) **Maternal factors** – Age, parity—the number of times a mother has given birth—and birth interval.
- 2) **Environmental containment** – air, food, water, soil, insects etc.
- 3) **Nutrient deficiency** – calories intake and diet composition.
- 4) **Injury** – Accidental or intentional, the latter presumably related to violence and self-harm.

The paper also defines a fifth variable, personal illness control, capturing individual education, or more specifically, how socioeconomic factors impact individu-

als' ability to prevent, or treat, their own and children's illnesses. The final piece of the puzzle is to define the socioeconomic variables impacting the proximate determinants

Even if you are not a statistician, or economist, a simple empirical framework should now be taking form in your head.

$$Y = F(x, y, z \dots)$$

Where Y is a measure of mortality or health and x, y, z are independent variables, capturing the impact of the proximate determinants on Y. As a dependent variable, Mosley and Chen (1984) proposes to combine the perspective from a binary mortality indicator and a continuous variable, measuring children's health. Based on contemporary research, they suggest using weight-for-age, with an upper grade signifying death, claiming that such a variable offers the researcher ample opportunity to construct a rich dataset, in most circumstances. I have not verified this empirically, but my intuition is that any comprehensive analysis would seek to study both binary variables, via a probit or logit model, and a continuous cross-sectional, time-series or panel data study.

As far as the proximate variables, data availability and research creativity set the limit.

There is one more model we need to deal with before concluding on the stylized facts of human mortality, mainly because it shows up in most economic studies of mortality. I am



talking about the so-called Grossman model, based on the widely cited Grossman (1972). The model is a set in a standard neoclassical world. Health is treated as a durable good, following a simple discrete law of motion for net capital accumulation, which consumers optimized for under a budget constraint. In Grossman's model, the representative agent draws utility from investment in health for two reasons; it reduces the disutility of poor health, and because it increases the time available for market and non-market activities, both of which lead to higher utility.

The formula for the evolution of health according to Grossman (1972) is:

$$H_{(t+1)} - H_{(t)} = I_{(t)} - dH_{(t)}$$

↔

$$H_{(t+1)} = I_{(t)} - (1-d)H_{(t)}$$

↔

$$H = A(1-d)^t + I_{(t)}/d$$

Where H is the stock of health, I is investment, and " d " is depreciation. The price of health is positively related to the price of medical care, age, and the rate of depreciation—itsself assumed to be a positive function of age—and negatively with education. Mortality occurs when the stock of health falls below a certain level. Counterintuitively, in this model households choose how long they live, and can live forever, depending on the parameters and their available resources. This doesn't make any

sense, and it is fair to say that the Grossman model isn't uniformly accepted in the world of health economics.

Zweifel (2012), for example, isn't impressed with the model, despite its widespread use, describing it as "elegant, very inspiring, but of limited relevance to the real world." The main criticism is that the Grossman model is ill equipped to deal with the stochastic element of health shocks, and even when authors have tried to correct for this, the results have been poor.

Zweifel (2012) also makes the point that health in period t is negatively correlated with investment in health in $t-1$, though this is partly due to the narrow definition of investment. If investment is limited to spending on outpatient visits or outlays for medical care, it is pretty clear that the higher the investment the more ill you're likely to be. You don't spend a lot of time at the doctor to prevent future illness, or at least, most people don't. It seems to me, though, that a broader more holistic definition of capital expenditures could restore the positive relationship between health and investment.

It is not my intention to get bogged down in a long discussion about the merits of the Grossman model. It is clear, however, that anyone studying the topic of mortality, and what drives it, will have to contend with a big volume of literature that uses it. I tend to take the charitable view that the correct prism through which to view the



Grossman model is whether some version of it can match the empirical facts, effectively generating the curve postulated by Heligman and Pollard (1980).

Muney and Moreau (2021) shows that it can, and is probably as good an example as any of an all bells-and-whistles modern model of human mortality. Using a sample of French birth cohorts from 1816 to 2015, drawn from the Human Mortality Database, Muney and Moreau (2021) develops a version of the Grossman model to generate mortality patterns closely matching empirical reality. Specifically, the model produces a mortality pattern closely resembling the model presented by Heligman and Pollard (1980)—a result that is replicated in a sample of chimpanzees—and also manages to capture the so-called “rectangularization” of mortality across cohorts, in effect describing the fact that survival rates have increased across age groups over time, except for in very old age.

TOWARDS AN IDEAL MODEL

Heligman and Pollard (1980) and Preston (1975) provide foundation for the first chart below, which plots my version of two theoretically optimal mortality curves through the demographic transition.

The idea that life expectancy is a concave function of time follows directly from the Preston curve. It is based on the simple idea that improving advances in living conditions, healthcare and technology to keep death at bay face diminishing returns over

time. This perspective is reflected in the Preston curve via the fact that the positive relationship between higher life expectancy and income flattens above a certain level of income.

This perspective assumes that there is, in fact, a biological limit to the human life span. This, in turn, raises the question about how far along the path towards this hypothetical endpoint humans have travelled. More specifically, the question is whether we are still on the linear part of the curve, or whether are we now at the point where diminishing returns to improved healthcare are setting in, as a collective.

That’s a difficult question to answer. A 2016 article in *Nature*, claims that 120 years represents the maximum of the human life span, concluding that maximum life expectancy has already plateaued. The study finds that life expectancy for super-centenarians peaked in 1997, with the death of death of 122-year-old French woman Jeanne Calment. At a first glance, this analysis then seems to suggest that diminishing returns set in more than two decades ago, though it’s important to be clear about the distinction between two mathematical perspectives.

The research published in *Nature* makes its argument at the margin of the age distribution, trying to determine the probability of finding humans living beyond 120 years. Even if this probability is vanishingly small, it doesn’t mean that average life



expectancy still can't increase substantially from this point on. It is one thing saying that the probability of finding humans living beyond 120 years is very low, it is quite another trying to move the mode and median of the distribution towards that theoretical limit.

But why 120 years? I don't have an answer to that question, though I think that I can explain why a limit somewhere around that level plausibly exists. I'll lean on the work by evolutionary biologist Bret Weinstein whose explanation of the issues at play—featured in his PhD thesis—is the most elegant account that I know of. What follows is my attempt to re-create his explanation.

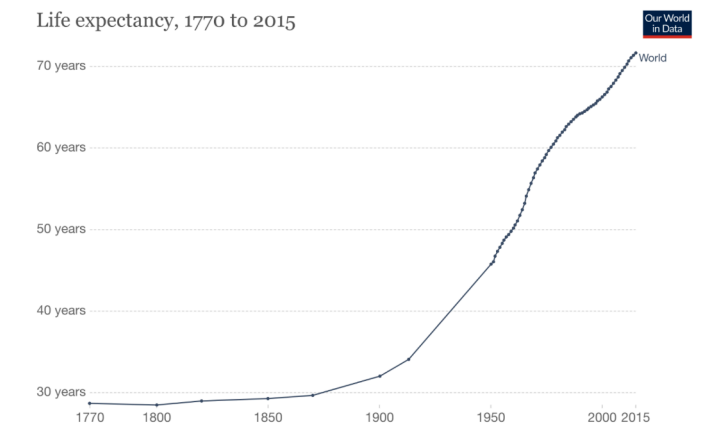

One way to explain the process of aging is to relate it to our cells' declining ability to reproduce themselves. This ability, in turn, is a function of the length of the so-called telomeres, which, according to Wikipedia represent: a *region of repetitive*

nucleotide sequences associated with specialized proteins at the ends of linear chromosomes.

All we need to know about telomeres for the purpose of this discussion is that they shorten every time a cell reproduces itself, eventually reaching the hayflick limit, beyond which no further reproduction is possible. It stands to reason that as more and more of your cells reach this point, your body will eventually stop functioning, and you will die, hopefully, old.

The solution to eternal life is now clear, even for the layperson. We just need to find a way to create infinitely long telomeres allowing our cells to reproduce with no limit. A brief Google search will reveal that research into extending telomere length, or abolishing the hayflick limit, for the purpose of extending human life is hot stuff. I will not pass judgement on this line of inquiry here. As Bret Weinstein neatly explains, however, there is a

fig. 03 / Mortality, "in theory" - **fig. 04 / Up and away**






catch; cancer. Put differently, a cell with the ability to reproduce itself infinitely, and unpredictably, is exactly what a cancerous cell is, among other things. Suddenly, hayflick limited telomeres isn't such a bad thing after all. From this perspective, it is exactly the hayflick limit in our cells, which prevent them from reproducing uncontrollably, in most cases that is. Sometimes, nature, or an adverse external environment, throw a spanner in the work and cancer is the result. In short, you can stay young forever, by coming up with a formula for extending telomere length, but in order to enjoy life, you would need to cure cancer too. As far as I know, this challenge still eludes us.

There are two alternatives to the hypothesis of a positive, but concave, relationship between improving healthcare, technology, or income and life expectancy. One is that the relationship is linear, though that's not very likely, and also to some

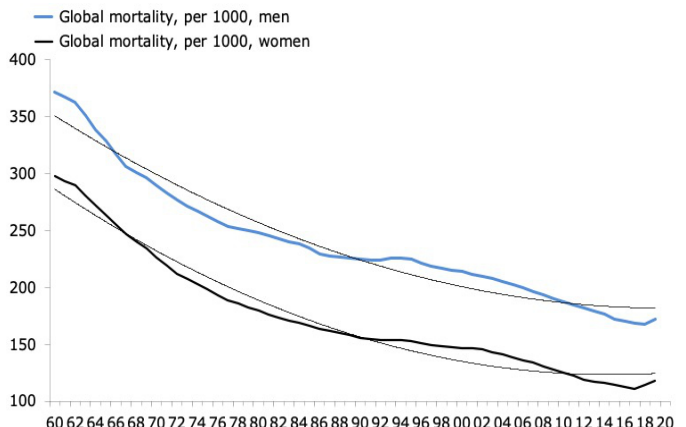
extent captured by concave function. Even if the true relationship is concave over time, we might still be on the near-linear part of the a curve. The second alternative is an exponentially rising function, which would occur in the context of a technological jump that allowed humans—probably a select few with a lot of resources to begin with—to live for a very long time, potentially even indefinitely. The annals of science fiction are full with suggestions for how such a technology could look or work. Allowing for a redefinition of what "life" is—if you upload your conscience to a robot are you still alive in a normal sense?—it should at least be considered that humanity at some point might achieve a technology that allows for long-term existence of the entity we today call conscience, probably in some kind of embodied form.

The second line, drawn as a function asymptotically declining towards zero over time is

fig. 05 / Still rising... - fig. 06 / ...Stabilizing?



Source: World Bank



Source: World Bank



the best-case scenario, though it's probably a reasonable assumption for the evolution of mortality in a modern capitalist economy in an environment with no war or conflict. One way to think about it is that this line is the sum of successive shifts *lower* in the swoosh-shaped line plotted by Heligman and Pollard (1980) and *upward* shifts in the Preston curve.

STILL IMPROVING

Theory is one thing, but what does the empirical evidence match the hypothesized stylized evolution in mortality. The answer after a quick glance at the numbers is, kind of.

Chart 04 above is of one the best attempts that I have seen charting the cause of life expectancy across the demographic transition. It draws on a number of sources, and clearly shows the point at which humanity broke the Malthusian chains, and it doesn't appear as if diminishing returns are setting in, yet, at least not in aggregate. The first chart below confirms this conclusion, indicating that life expectancy in aggregate still exhibits a positive linear correlation with time, here based on numbers from 1960 to the present. The latest full sample from from the World Bank shows that global life expectancy at birth was 72.7 years in 2019, split between extremes such as some 60 years years in Sub-saharan Africa and just over 82 years in the OECD. Across genders, mortality rates are higher for men than women, driven by the fact that men tend, on average, to be more

engaged in dangerous activities and occupation, and that they tend to be more "successful" in the context of suicide.

The level of life expectancy in 2019 compares to 67.5 years in 2000, 65.4 in 1990, and 62.8 years in 1980. In that 39 year period, global life expectancy at birth has increased by an average of 0.25y per year. Extrapolating this trend suggests that global life expectancy will have increased to 80 by the end of the 2040s, and 100 by the year 2128, for those looking further ahead. Such linear extrapolation probably isn't worth much, though it goes to show that even in a world where we assume a biologically binding age limit at 120 years, there is still of room for improvement.

At present, the most benign conditions for life expectancy and mortality are found in Western Europe, North America, Australia and many parts of South East Asia.

So far so reassured, it is worthwhile sketching the four horsemen of the apocalypse that could alter that picture.

Natural disaster - Imagination is the limit in this category. Asteroids, earthquakes, volcanos, storms, floods all have the ability to drive a significant increase in human mortality, if not extinguish human life altogether. As these are discrete events, they have to be understood in their correct context. A natural disaster leads to a temporary increase in period-mortality, but not necessarily a lasting shift in



the trend in mortality. The larger the disaster is, however, the higher is the risk that the underlying drivers of the trend changes too, especially in the context of permanent changes to the ecological environment.

Climate change and pandemics

These two could have been included in natural disaster bracket, but I think they deserve their own category. Unlike the true exogenous nature of the natural disasters mentioned above, climate change and pandemics are, at least in part, endogenous to human existence. In other words, we are to blame for both. The discourse on climate change is particularly controversial in this regard. Man-made climate and environmental change is a fact, and it certainly has the potential to alter human life for the worse in areas that are least able to adapt. Whether it is a threat worthy of the alterations in human life that are being proposed to halt it, however, is more debatable.

In the extreme, the threat of rising mortality from climate change—either directly or indirectly via conflict or war—is best solved by an increase in mortality. This is a provocative interpretation of the contemporary discourse, but it is also one that is difficult to escape in the ultimate analysis. It is, put simply, the idea that from the point of view of the climate, maybe the issue is that there are too many people on earth. If that turns out to be the case, the solution, while cruel, is straightforward, and we have the tools to implement it. For the record, I be-

lieve that such fatalism in the context of climate change analysis is morally wrong, not to mention rather unambitious, but it is a perspective of the discourse that is increasingly difficult to ignore.

Meanwhile in the context of a the life-eradicating superbug, it is a threat that comes with varying levels of risk. I would argue, for example, that the idea of such a bug killing the entire human race is remote, though it doesn't have to, in order to make a big dent in mortality. Indeed, the Covid-19 epidemic provides an interesting realtime experiment of how much aggregate mortality reacts to such an event—after all, the CFR is estimated to be around 1%, at worst—and whether it leads to a sustained shift—for example on a cohort basis—or just a blip.

Following Muney and Moreau (2021), the literature talks about “harvesting effects”, which cover environmental effects that displace deaths. The 2003 heatwave in France, for example, drove up mortality sharply for the elderly population in the reference year, which was then followed by a sharp decline in the subsequent year. This shift is driven by the fact that the environmental shock increases mortality in the weakest part of the population, leaving a more resilient and stronger population in the subsequent periods, with a lower period mortality, at least for a time. This logic can be applied to pandemics, and almost surely, Covid-19 too.



War and conflict - Again, imagination sets the limits, though in most cases, the real-life horror and cruelty are far worse than anything we can imagine. It's worthwhile distinguishing between two overall categories of war from the perspective of human mortality; nuclear war and everything else. The former is a potentially human life-extinguishing event, and can, in the extreme, be analyzed along the same lines as an earth-destroying asteroid strike or similar. However narrow in scope, a nuclear war or exchange of any kind would likely drive a significant shift in human living conditions, and therefore mortality, especially in territories where such weapons were deployed.

The latter—non-nuclear war and conflict—comes in so many forms that it is difficult to categorize, though they have significant effects on mortality too. The First and Second World Wars, for example, had significant impact on mortality, and birthrates, ostensibly by wiping out a significant portion of young men in Europe. More localized wars and conflict don't show up in global macro data, but they have the ability to radically change regional and country-specific mortality. Indeed, in the context of ethnic conflicts, some population groups face annihilation altogether. Finally, even excluding nuclear exchanges, the potency of modern weaponry, and the force with which advanced nations can utilize such tools, show that the civilian population in a territory playing host to a modern armed conflict face two choices; flee,

and if that's not possible, near-certain death.

Muney and Moreau (2021) draw on literature discussing so-called "scarring effects" of war to estimate that WWI lowered life expectancy of the 1896 male cohort in their sample by a whopping 16 years, and that WWII lowered it by a further 2 years. These numbers follows Wilson (2014), which shows similar effects in a study of morality in among men in New Zealand born in 1896 compared to those born in 1900.

Endogenous shifts in health -

The analysis thus far assumes that mortality is, and always will be, a declining function of modernity, absent war or natural disaster. We need to consider that it isn't. Obesity, cancer and other lifestyle ailments, as well as drug addiction—often linked to depression and suicide—are all threats. To the extent that such ailments are a direct consequence of modernity, we are getting better at treating them—ostensibly keeping people alive despite their prevalence—but that balance isn't necessarily always going to be in our favour. The clearest example at time of writing is the fact that life expectancy in the U.S.—the world's most prosperous nation—has been falling since 2014, a trend linked to drug overdoses, alcohol abuse, suicides, and other ailments. The literature speaks of diseases and deaths of despair, covering drug overdoses, suicides and alcoholic conditions. The 2020 book, *Deaths of Despair*, by Case and Deaton offers



a timely overview of the state of play in the U.S., and the picture isn't pretty. The study builds on the couple's 2015 paper, *Rising morbidity and mortality in midlife among white non-Hispanic Americans in the 21st century*, which details an increase in mortality linked to an increasing number of deaths among low-income and poorly educated, mainly white, men and women due to "drug and alcohol poisonings, suicide, and chronic liver diseases and cirrhosis." In the middle of the 1990s, these so-called deaths of despair counted around 70,000 per year, a number that has since increased to 150,000-to-200,000, a significant increase even factoring in the 30% increase in the size of the U.S. population over the same period. So, why is this happening?

Deaton and Case (2020) identifies a "long-term drip of losing opportunities and losing meaning and structure in life", linked to changes in family structure, education and labour force participation rates, especially among white men.

Speaking to Vox.com research Roge Karma, Angus Deaton hits the nail on the head;

"when you look at these graphs of labor force participation and wages both trending down together, it's very hard to conclude anything except that it's the supply of jobs that has gone wrong, and there's simply less and less work for less skilled people."

The analysis and data presented in Deaton and Case (2020) pro-

vide a glimpse of an important, and uncomfortable, reality. The data in the U.S. seem to suggest that it is very possible for a combination of adverse factors to reduce life expectancy in rich societies, even in the context where society as a whole is getting richer and more technologically advanced.

A more recent analysis by Max Roser, founder of *Ourworldindata.com*, adds colour by listing the number of reasons why life expectancy in the US is lower than in the rest of the developed economies. A higher rate of obesity and opioid overdoses stand out as the most obvious driver. More generally, Mr Roser's analysis implies that the US is simply a stand-alone example of a bad mix between high healthcare costs and a poor outcome in aggregate.

Muney and Moreau (2021) attempts to generalize this perspective, drawing on the comprehensive Chetty et al (2016) detailing the drivers of different mortality outcomes in the US between 2001 and 2014. Muney and Moreau (2021) hypothesis that investment in health—negatively correlated with mortality risk in the Grossman model—is positively correlated with education and income. They further speculate that a higher socioeconomic status is linked to slower aging, thanks to lower exposure to pollution, stress and a more healthy lifestyle.

These findings sound imminently reasonable to me, and give rise to the idea that the Preston curve does indeed apply



within countries as well. That said, evidence from one country is not enough to draw that conclusion. In addition, and assuming that differences in socioeconomic status can be used to make conclusions about within-population dynamics across countries, the US data give rise to another interpretation.

This is the possibility that *within-country* dynamics, can spill over into a population level phenomena if the deterioration in one group's life expectancy overwhelms the improvement of longevity for other groups. Such contours of modern-day Dickensianism is closely linked to the broader, and increasingly publicized, story of income and wealth inequality. This is already a topic studied by polemicists on both sides of the political spectrum. As per usual, to the extent that everyone agrees on the diagnosis, commentators vehemently disagree on the underlying roots of the issue, and, as a result, what to do about it.

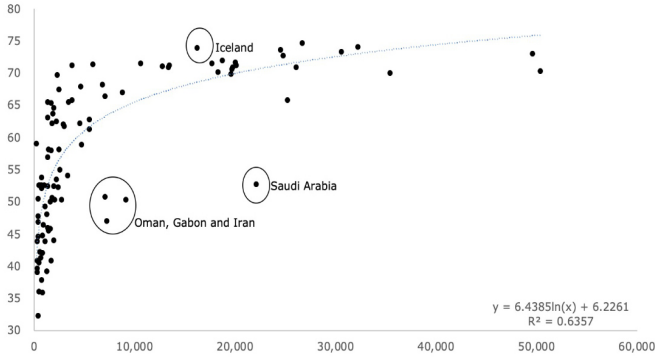
It is not this essay's intention to get bogged down in a detailed analysis about the extent to which either of the risks described above will drive a sharp shift in global mortality and life expectancy anytime soon. At this point, the trends in life expectancy and mortality are up and down, respectively. I am not willing to bet on a change of this, though it seems fair to note that the focus on existential risk is on the rise, across many disciplines and many different levels of societal discourse. This doesn't necessarily mean such risks are rising, though the notion of "no smoke without fire" come to mind. Existential risks are in part a question of simple probability—which rises towards one as time goes by—such as in the case a life-destroying asteroid hit, while in other cases, it is endogenous to human behavior. Time will tell which part of the story wins out. As far as this project goes, it proceeds assuming that the status quo, falling mortality and rising life expectancy, persists for a little longer.

This space has intentionally been left blank


Find the landing page for this piece, and an overview of the project as a whole, [here](#).

Appendix 1 – Preston curves, 1970 to 2018


• 1970 Preston curve, N=95; source: World Bank, UN and the OECD



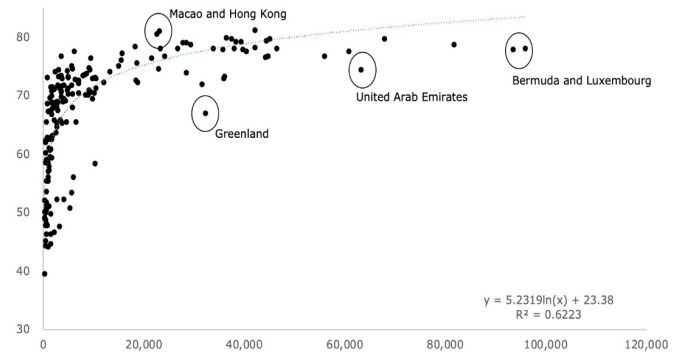
• 1980 Preston curve, N=120; source: World Bank, UN and the OECD




• 1990 Preston curve, N=161; source: World Bank, UN and the OECD



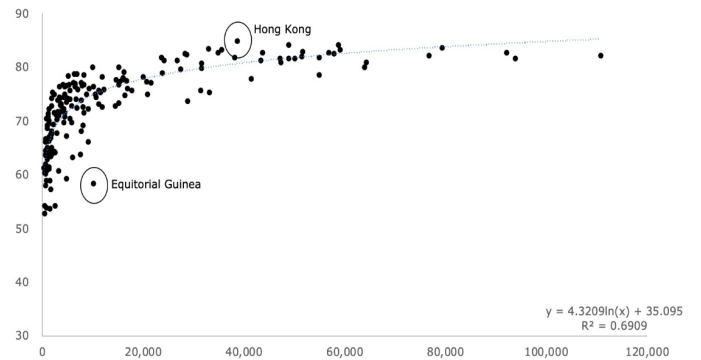
• 2000 Preston curve, N=186; source: World Bank, UN and the OECD



• 2010 Preston curve, N=193; source: World Bank, UN and the OECD



• 2018 Preston curve, N=183; source: World Bank, UN and the OECD



Source: World Bank database. Original source are World Bank, UN population stats, and OECD national accounts. All curves and models created in Excel.