NBER WORKING PAPER SERIES

DEMOGRAPHIC PERSPECTIVES ON MORTALITY OF COVID-19 AND OTHER EPIDEMICS

Joshua R. Goldstein
Ronald D. Lee

Working Paper 27043
http://www.nber.org/papers/w27043

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2020

We received helpful comments on this work from seminar participants in the Department of Demography. Magali Barbieri helped us find the relevant data for Europe. Yi Zhou helped with the Chinese data. Peter Choi sent us data for S. Korea. Will Dow made many helpful comments. We thank them all. Goldstein is supported by a grant from the National Institute of Aging # R01AG058940. Both authors are supported by the UC Berkeley Center on the Economics and Demography of Aging (National Institute of Aging 5P30AG012839), and the Berkeley Population Center (National Institute for Child Health and Human Development P2CHD073964). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Joshua R. Goldstein and Ronald D. Lee. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Demographic Perspectives on Mortality of Covid-19 and Other Epidemics
Joshua R. Goldstein and Ronald D. Lee
NBER Working Paper No. 27043
April 2020
JEL No. J1,J10,J11

ABSTRACT

What would a hypothetical one million US deaths in the Covid-19 epidemic mean for mortality of individuals at the population level? To put estimates of Covid-19 mortality into perspective, we estimate age-specific mortality for an epidemic claiming for illustrative purposes one million US lives, with results scalable over a broad range of deaths. We calculate the impact on period life expectancy (down 3 years) and remaining life-years (12.3 years per death), which for one million deaths can be valued at six to 10 trillion dollars. The age-patterns of Covid-19 mortality observed in other countries are remarkably similar and exhibit the typical rate of increase by age of normal mortality. The scenario of one million Covid-19 deaths is similar in scale to the decades-long HIV/AIDS and opioid-overdose epidemics but considerably smaller than the Spanish Flu of 1918. Unlike HIV/AIDS and opioid epidemics, the Covid-19 deaths will be concentrated in months rather than spread out over decades.

Joshua R. Goldstein
Department of Demography
2232 Piedmont Avenue
Berkeley, CA 94720
josh.goldstein@berkeley.edu

Ronald D. Lee
Departments of Demography and Economics
University of California, Berkeley
2232 Piedmont Avenue
Berkeley, CA 94720
and NBER
rlee@demog.berkeley.edu

A data appendix is available at http://www.nber.org/data-appendix/w27043

Introduction

As we write, societies around the world are struggling to protect their populations from the Covid-19 pandemic. Meanwhile, both citizens and policy makers are trying to make sense of the magnitude of the crisis and the wave of deaths that could be coming. In this paper, we present several different ways to think about the mortality of the epidemic. It is possible to portray the death toll in a way that feels overwhelmingly large, but it is also possible to describe it in a way that makes epidemic mortality seem almost negligible. Our view is that Covid-19 should be seen as an extremely large mortality threat, smaller in scale than the Spanish Flu, but equal in overall magnitude to the HIV-epidemic or the opioid crisis—but unlike these, with Covid-19 deaths concentrated into the span of a few months rather than spread out over decades.

The intended contribution of this paper is not to provide new forecasts of Covid-19 mortality. Instead, we combine interpretation of the age pattern of deaths observed to date with existing forecasts for the total impact of the epidemic to produce an estimated age-profile of Covid-19 mortality. This age-profile, which can be scaled up or down, enables estimation of the epidemic's impact on period life expectancy and loss of person-years of life at a population scale, as well as comparison with past epidemics.

A further contribution of the paper is to show that the age-pattern of deaths, when appropriately adjusted, is nearly identical across a wide-range of countries and stages of the epidemic. This age pattern is remarkably similar to the gradient of all-cause mortality for adults, for which the risk of death increases by about 10% per year of age. The age-profile of Covid-19 mortality may change over time, as treatment becomes more (or less) available. However, the age-gradient we see to date suggests that the risk factors for Covid-19 are similar to those for all causes of death.

Epidemiological models are predicting anywhere from 60,000 to more than 2 million additional Covid-19 deaths in the United States, depending on the success of mitigation, the severity of mortality, and the capacity to treat the sick (Murray et al, March 26, 2020; Ferguson et al, 2020, McAndrew 2020). For illustrative purposes we use an intermediate scenario of 1 million deaths in 2020 due directly to Covid-19 across all waves to make our discussion concrete, at times comparing it to a lower scenario of 250,000. The metrics we produce scale directly with the number of deaths, so readers can translate our results under different mortality scenarios. We do not include the indirect increase in deaths as healthcare systems are overwhelmed, as well as any potential lowering of mortality, for example, from decreased air pollution, traffic accidents and consumption of alcohol resulting from the economic slowdown. The age pattern of these changes might be quite different.

As we describe below, the most commonly used measure of mortality, life expectancy at birth, is not a good measure of transitory mortality shocks. We consider instead a combination of other measures, including the crude death rate, age-specific mortality, and the loss of remaining person years of life, which together give a better summary of the magnitude of epidemic deaths. These perspectives allow us to compare mortality impacts over time and across populations. They also allow policy makers to make more informed judgments about the value of saving lives.

The most direct indicator of mortality is the number of deaths. This count is often given relative to population size. Given the most recent U.S. Census Bureau projection of population in 2020 and applying the 2017 mortality, we find about 3 million deaths out of a total population of 330 million, a crude death rate of 9.1 per thousand. In the absence of the epidemic, it would have been less than 10 per thousand.

An additional one million deaths from Covid-19 would increase the total annual deaths to 4 million, raising the crude death rate to about 12.1 (or to .985 for 250,000 deaths). The increased risk to the average person, 3/1000, is small in absolute size but large in relative terms, with a proportional increase of a third. Epidemic mortality will likely be compressed into a small portion of the year: if most of the deaths occurred within a three month period, the daily risk of mortality would be more than double its normal level.

The age-pattern of Covid-19 mortality

The “case fatality rate” or CFR is the measure most commonly shown by age. It is measured in somewhat different ways in different places, but in principle it is the number of Covid-19 deaths at some age divided by the number of reported cases of Covid-19 at that age. In practice, Covid-19 deaths may be only those occurring in ICUs, omitting those occurring at home or in nursing homes, for example. In practice, the number of cases may be those reported by doctors, or confirmed by tests, or those admitted to hospitals. Evidently there are problems with both the numerators and denominators of the CFR measure.

An alternative is to measure the cause-specific mortality rate by age. We have compiled age specific death data from countries – China (Hubei), S. Korea, Italy, Spain, France, and the United States. Our approach does not require that the counts of deaths have the same level of completeness across countries, which vary both in the definitions they use and in the stage of the epidemic we observe them for. But we do assume that the age-distributions of reported Covid-19 deaths are accurate. For each country we divide deaths at each age by the population at each age, typically for the nation as a whole[1], to calculate (unnormalized) age specific death rates (ASDR) for Covid-19. (See figure 1 (a)) We then normalize each country's rates by dividing by the sum of the rates, so that the normalized rates sum to 1.0. Since the epidemic may be concentrated in one part of the country, and since some countries may be at early stages with few reported deaths, these ASDR may be extremely low in some countries and much higher in others, which does not necessarily signal a more or less severe epidemic and should not be so interpreted. Instead, we believe that the most reliable information is the shape or pattern by age of death rates, abstracting from the level. (See figure 1(b)). We can see that the age-patterns are remarkably similar across the 5 countries. The age grouping of deaths in the US data is different than for the other countries. However, when we adjust the age groupings and normalize the US data, the age pattern looks similar to the average of the other five countries.

[Figure 1 about here]

It is evident from Figure 1 that Covid-19 mortality risk is many times higher for the old than the young, and indeed the vast majority of Covid-19 deaths are of older people. But the same is true for normal mortality—the vast majority of deaths are of the elderly. We estimate 75% of all US Covid-19 deaths to be age 70 or above, somewhat above the 64% for normal mortality. In fact, the age-distribution of deaths attributed to Covid-19 so far is quite similar to that of all-cause mortality, which tends to increase by about ten percent every year of age after age 30. Figure 1(b) shows that in China (Hubei), South Korea, Italy, France, and Spain, virus-attributed mortality rates rise by about 11% per year (a bit slower in Hubei, where the rate is 9.5%), close

---

[1] The only exception is China where we used the population age distribution of Hubei because 840 of the 1,023 deaths in Chinese data were in Hubei. We thank Prof. Yi Zhou for these calculations.

to the 10 percent we would expect for all-cause mortality. However, the rate of increase in US Covid-19 death rates by age is somewhat lower, at 8.5%.

At ages under 40, Covid-19 mortality risk is low, but so is mortality from other causes. It appears in the lower panel of figure 1(b) that average mortality is less than would be predicted from the exponential pattern of mortality at older ages. However, the number of cases is so small that we are hesitant to draw a conclusion before more data become available.

The consistent age-profile of Covid-19 means that the number of deaths in a population will depend both on the severity of the epidemic and the distribution of the population by age (Dowd et al 2020). Population aging is largely driven by the history of population growth. Indeed, in stable populations (that is, populations with constant and equal growth rates at all ages), mathematical modeling tells us that a difference of 1% in the historical rate of growth produces about a 40-fold difference in the crude death rate (Lotka 1936). A country with a long-term growth rate 1% higher than another (for example, the United States compared to Italy, or India compared to the United States) would be expected to have 40% fewer deaths per capita, even if the age-specific mortality rates were identical. The enormous effect of age structure is a caution against attributing differences in aggregate death rates to success or failure in addressing the epidemic.

Epidemic mortality risk as temporary aging

The similarity of the age-pattern of Covid-19 mortality to all-cause death rates allows one to translate the short-term period of elevated risk into measures of “temporary aging,” units that are perhaps more meaningful for individuals than the usual relative risk measures. The “temporary aging” measure expresses the increased risk of an individual during the months of the epidemic into the age of a person with the same mortality during normal, non-epidemic times. Because normal mortality rises at a rate of about 10% per year of age, a temporary increase of 10% in risk translates into making a person effectively 1 year older for a limited amount of time. If epidemic mortality is concentrated in 3 months, then a death toll equal to 10% of the annual total would raise mortality during this window by 40%, so that each person is temporarily subject to the mortality of someone 3.5 years older.

Table 1 gives estimates for the impact of different epidemiological forecasts for the United States in terms of temporary aging, assuming a 3 month window of epidemic mortality. We provide a range from as low as 125,000 -- a bit more than the lowest estimate released by the White House -- to a high of 2 million, a bit less than the estimate of an uncontrolled epidemic by Fergusun et al (2020). If we consider the estimates in Table 1, we can understand why for an older person, even the most optimistic forecast implies substantial risk: a 78 year-old during the epidemic would be exposed to the risk faced by an 80 year-old in normal times. If the epidemic is larger, then even a younger person would face substantial risk. In the uncontrolled epidemic a 30 year old would temporarily experience the mortality rate of someone close to 60. Considering mortality risk in this manner allows, we believe, accurate communication of risk. However, the approach has its limits, applying neither to children nor, probably, to the oldest old.

Table 1: Hypothetical epidemiological impacts in terms of years of temporary aging. An impact of 1 million deaths in three months in the United States would have the effect of temporarily exposing a 60 year-old to the normal mortality of a 73 year old.

Impact in 3 months: deaths	Temporary aging
2 million	26.7 years
1 million	13.3 years
500 thousand	6.7 years
250 thousand	3.3 years
125 thousand	1.7 years

Notes: Estimates assume a 10% rate of mortality increase with age.

Life expectancy decline overstates the impact of temporary epidemic mortality

The “period” life expectancy at birth is a familiar way to summarize the mortality in a year. In 2018 -- the most recent year reported for the United States -- life expectancy at birth was 78.7

years, a statistic which assumes a person lives their entire life, from birth to death, under the mortality conditions of 2018. In the context of epidemic mortality, life expectancy at birth is a misleading indicator, because it implicitly assumes the epidemic is experienced each year over and over again as a person gets older.

With this important caveat in mind, it is possible to calculate the effect of our estimated age-profile of Covid-19 on period life expectancy by comparing it to the normal mortality in Figure 2 to the value for total mortality there. We find this difference to be 3.0 years, a reduction from 78.7 to 75.7.

[Figure 2 about here]

This decline is slightly larger than what mathematical models of the life table would predict (Keyfitz 1977). A new cause of death that increased mortality rates by the same factor Δ at all ages would reduce life expectancy by a factor of ΔH with H typically about 0.15 in low mortality countries. The scenario of 1 million deaths we are considering would increase mortality by about 1/3, suggesting a drop in period life expectancy in 2020 by about 3.9 years.

Loss of remaining life

Whereas period life expectancy in an epidemic year tells us how long a person would live if they were to experience an epidemic every year of their lives, what we would really like to know is how a one-time epidemic affects the remaining life expectancy of the actual population.[2]

Based on Social Security actuarial projections by Bell and Miller (2005), we calculate that the 2020 American population of 330 million people has on average about 44.4 years of remaining life expectancy, totaling more than 14.8 billion person-years. The Covid-19 epidemic, if it kills an additional million people, will result in a loss of 12.3 million years of remaining life expectancy. This represents a loss of less than 1/1000th of the population’s remaining years to live.[3] Older individuals age 70 to 89, taking those who die and those who survive together,

---

[2] We treat this as a one time event, although it is not known whether infection confers lasting immunity, nor whether this coronavirus will recur in future years.

[3] As small as it is, our calculation is likely to overstate the loss of remaining life in that it assigns the remaining life expectancy based only on age, without taking into account that Covid-19 deaths are disproportionately occurring among those with compromised health status.

would on average lose about .2 years of remaining life, and younger individuals would lose far less.

How could such an enormous loss of lives produce such a seemingly small loss of remaining life expectancy? Two factors play a role. First, even with additional Covid-19 mortality, death will still be a statistically rare event. Most will survive and they will, if mortality returns to normal, have many years of life ahead of them. Second, those who die of Covid-19 are older and have on average fewer years of remaining life expectancy than the average person (12.3 instead of 44.4).

Our calculations above used the observed U.S. population and our estimated pattern of Covid-19 mortality. However, a stylized mathematical model produces similar results for the proportional loss of life from an epidemic. In the case of a stationary population with epidemic mortality proportional to the baseline age-pattern, an epidemic that increases mortality by a factor Δ will result in a loss of remaining life expectancy equal to $\Delta H/A$, where $H$is Keyfitz's life table entropy and A is the mean remaining life expectancy of those alive. (See Appendix 2). For a stationary population with the U.S. 2017 life table, Δ = ⅓, H = 0.15, and A = 40, this model predicts the share of lost life would be 1/800, slightly more than the exact calculation.

Life years lost, relative to non-epidemic mortality

While life expectancy calculations overstate the loss of life from a one-time epidemic, calculating the loss of remaining life expectancy tends to minimize the impact. We propose that a more useful way to quantify life lost from an epidemic is to consider it relative to the person-years lost from mortality in a "normal" non-epidemic year. This calculation accounts for the number and age of deaths from the epidemic, and weights them by the loss of remaining person years, comparing the result to the person-years lost in a comparable non-epidemic year.

[Figure 3 about here]

Using this metric, we estimated the average age of death and computed the comparative loss of life from Covid-19 relative to the Spanish Flu, the HIV epidemic, and the recent opioid epidemic according to three different measurements. For Covid-19 we show scenarios for 1 million and 250 thousand deaths. In the top row of Figure 3, showing the counts of deaths, with one million deaths Covid-19 would be the largest threat we have faced. In the middle row of the figure, which takes population size into account, Spanish flu emerges as having produced the largest increase in crude mortality rates. Taking the age of those who die into account and their remaining life expectancy, the bottom panel shows that in terms of lost remaining life expectancy, our scenario for Covid-19 is much smaller than the Spanish flu. The dashed lines on the Covid-19 bars show the same calculations based on 250,000 deaths.

The HIV epidemic, which peaked in the 1990s, and opioid overdose deaths, which continue today, have each over the decades produced total mortality that is comparable (for one million deaths), depending on the metric, to the Covid-19 scenario we are considering and even to Spanish flu. However, on a deaths-per-year basis, they are an order of magnitude smaller. In per capita terms, 1 million deaths from Covid-19 will confront society in one year, or possibly in just three or four months, while the deaths we have experienced from HIV or opioid overdoses occur over the course of decades. This concentration of epidemic deaths in a short time creates a crisis which overwhelms the capacities of health care systems, morgues, and mortuaries leading to triage in hospitals and make-shift storage of bodies such as in an ice rink in Madrid. The most appropriate measure of such crises is probably epidemic deaths per year per population or the proportionate increase in normal deaths. The slower-moving epidemics bring other stresses and anguish, since afflicted individuals suffer for many years, and the prevalence of those afflicted at any moment is consequently much higher than for epidemics that kill quickly.

The value of life saved

In one of the most quoted lines of the Talmud, it is said that whoever destroys one life, destroys an entire world; and whoever saves one life is considered to have saved an entire world. Still, policy makers face the inescapable choice of how many lives to save at what cost. Federal policy decisions are guided by a substantial literature in this area. Current estimates for the United States by Viscusi (2018, 2020) give a single value of $10 million to each life regardless of age or alternatively $500 thousand per year of life. Applying these numbers to the scenario we are considering, we see that a million coronavirus deaths would be valued at $10 trillion, while 12.3 million person years of life would be valued at about $6 trillion.[4] Health economists use

---

[4] These numbers are in close agreement with $4.7 trillion for 1 million lives saves, estimated by Greenstone and Nigam (2020) in relation to Coronavirus deaths.

substantially lower values around $125,000 per year of life (Institute for Clinical and Economic Review 2020) which would imply $1.5 trillion.

It is very difficult to evaluate the cost of measures taken to mitigate the epidemic. This is an active area of work (Correia et al 2020; Eichenbaum et al 2020). Early downward revisions of GDP forecasts for 2020 predict about 5% less output than expected, or about $1 trillion (Goldman Sachs). Not all of this decline should be attributed to societal choices to slow the spread of the virus, because the economy would also suffer -- perhaps even more -- if the virus were uncontrolled. But even if we assign all of the drop in GDP to measures taken to save lives, the economic costs of the actions we have taken appear to be appropriate for the scale of the crisis.

Conclusion

It is possible to portray the covid-19 epidemic as unimaginably large -- the biggest killer in American history -- or unimaginably small, reducing our remaining life by only a bit more than 1 part in 1,000. However, when the loss of life is put into comparative perspective, we see that the scale of an epidemic with a million deaths would be as large as the recent opioid and HIV crises but much smaller than the Spanish flu. The 1918 epidemic killed more people relative to population size, and it also caused a much greater loss of remaining life expectancy because those who died were so young.

As a society, we are making major and costly efforts to reduce mortality. The anticipated economic costs appear appropriate, or perhaps low, when compared to the statistical value of lives that may be saved.

The death toll of Covid-19 is a terrible thing, both for those who lose their lives and for their family, friends, colleagues and all whom their lives touched. Those are real individuals, not the abstract statistics presented here. But the population perspective helps us to place this tragedy in a broader context. As we put our efforts into reducing the impact of the epidemic, it is important to know that we as a society have been through such mortality crises before.

References

Bell, Felicite C. and Michael L. Miller (2005) Life Tables for the United States Social Security Area 1900-2100, Actuarial Study no.120. Social Security Administration. Office of the Chief Actuary. August 2005 SSA Pub. No. 11-11536

Correia, Sergio and Luck, Stephan and Verner, Emil, Pandemics Depress the Economy, Public Health Interventions Do Not: Evidence from the 1918 Flu (March 26, 2020). Available at SSRN: https://ssrn.com/abstract=3561560 or http://dx.doi.org/10.2139/ssrn.3561560

Lotka, A. J.: Théorie analytique des associations biologiques, 2ème partie. Hermann, Paris (1939)

Keyfitz, N. (1977). What difference would it make if cancer were eradicated? An examination of the Taeuber paradox. Demography, 14(4), 411-418.

Dowd, J. B. et al (2020) “Demographic science aids in understanding the spread and fatality rates of COVID-19” Pre-print. doi: https://doi.org/10.1101/2020.03.15.20036293

Eichenbaum, Martin S., Sergio Rebelo, Mathias Trabandt (2020) “The Macroeconomics of Epidemics” NBER Working Paper No. 26882 (March)

Ferguson, N. M., et al. (2020). Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand. Imperial College, London. DOI: https://doi. org/10.25561/77482.

Greenstone, Michael and Vishan Nigam (2020) “Does Social Distancing Matter?” WORKING PAPER NO. 2020-26 of Becker Friedman Institute for Economics of the University of Chicago.

Institute for Clinical and Economic Review (2020) “2020-2023 Value Assessment Framework” https://icer-review.org/wp-content/uploads/2019/05/ICER_2020_2023_VAF_013120-4.pdf.

McAndrew, T. (2020) "COVID19-Expert Forecast-Survey6-20200325.pdf" http://works.bepress.com/mcandrew/3/

Murray, Christopher J.L. et al “Forecasting COVID-19 impact on hospital bed-days, ICU-days, ventilator-days and deaths by US state in the next 4 months” Preprint submitted to MedRxiv 03.25.2020- tracking ID MEDRXIV/2020/043752.

Jiaquan Xu, M.D., Sherry L. Murphy, B.S., Kenneth D. Kochanek, M.A., and Elizabeth Arias, Ph.D. (2020) NCHS Data Brief, No. 355, January 2020, U.S. DEPARTMENT OF HEALTH AND HUMAN SERVICES, Centers for Disease Control and Prevention, National Center for Health Statistics, Mortality in the United States, 2018

Viscusi, W. Kip (2018) Pricing Lives: Guideposts for a Safer Society Princeton University Press.

Viscusi, W. Kip (2020) "Pricing the lives saved by coronavirus policies" New IZA World of Labor commentaries on the coronavirus crisis, MARCH 27, 2020 https://wol.iza.org/opinions

Murphy, Kevin M. and Robert H. Topel. "The Value Of Health And Longevity," Journal of Political Economy, 2006, v114 (No. 5, Oct), 871-904.

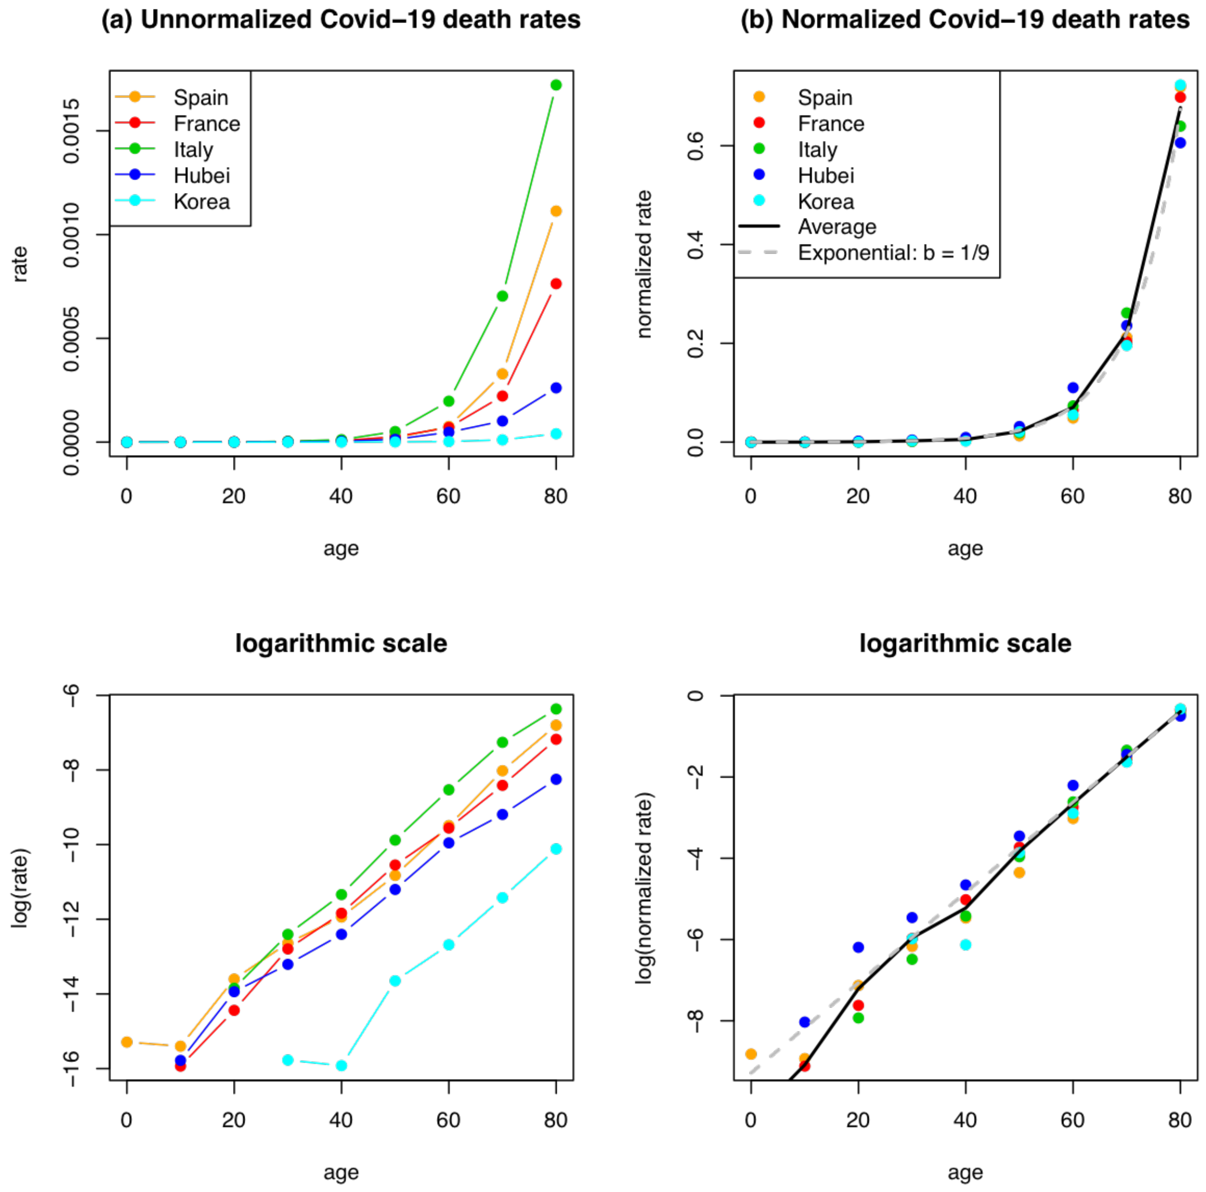

Figure 1: Similar Age Pattern of Covid-19 mortality by region. (a) Unnormalized age-specific mortality, calculated as the ratio of deaths by age to population by age. The levels should *not* be interpreted as reflecting real differences in mortality because of unaccounted variation in time scale, stage of the epidemic, and the extent of spread within the region (b) Normalized rates, giving all regions the same level, showing the common age pattern. The average is calculated across the five regions: it tracks the exponential increase of all cause mortality (b = 1/9 , leveled to intersect average at age 70). Lower panels show the same data as upper panels but in logarithmic scale. (Sources and methods in appendix 1.)

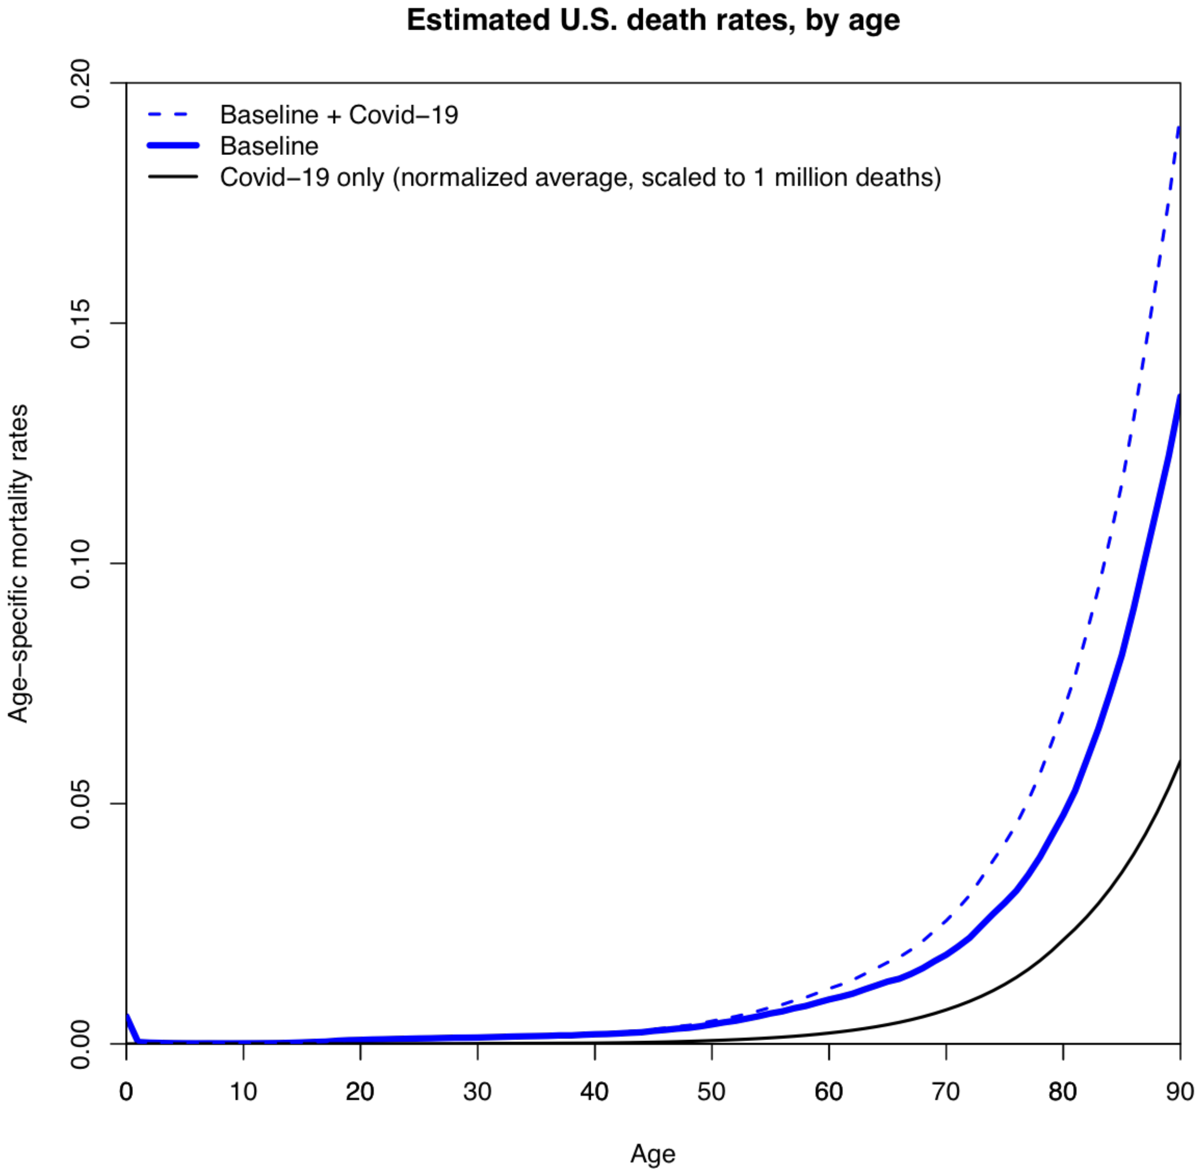

Figure 2: Estimated age-specific mortality in the United States in 2020 for the scenario of 1 million additional Covid-19 deaths. Baseline mortality in 2020 is assumed to equal that of 2017. Covid-19 mortality is estimated by averaging the normalized death rates in Figure 1 panel (b), then multiplying these average rates to result in 1 million additional deaths using the U.S. population age structure for 2020.

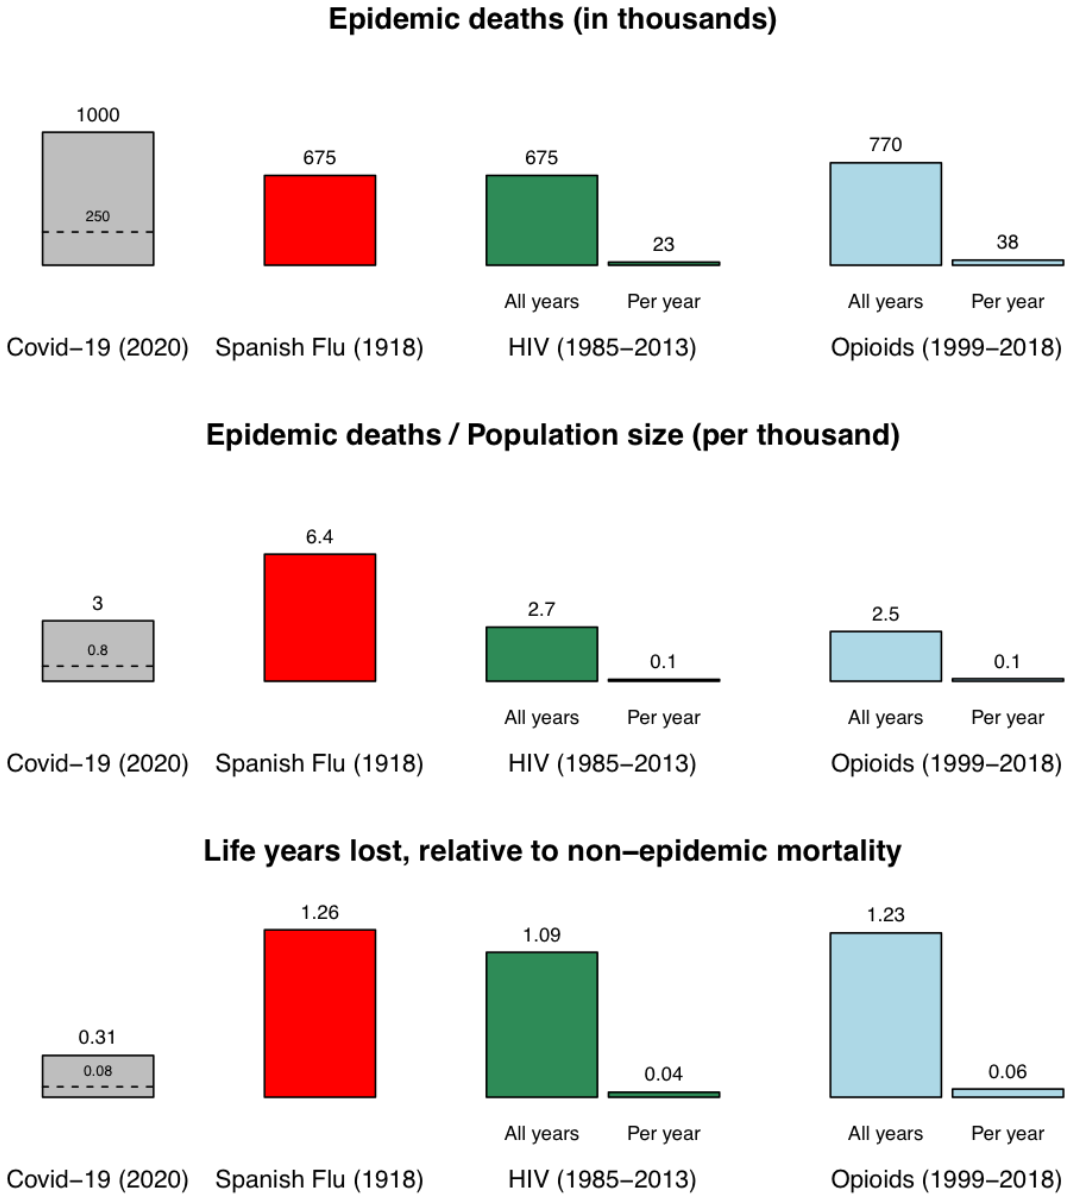

Figure 3. Mortality of Covid-19 scenario compared to past US epidemics according to different measures. In the scenario of 1 million Covid-19 deaths, the virus kills more Americans than past epidemics, but when population size is accounted for Spanish Flu is more deadly. Taking into account years of remaining life, we calculate that the Spanish Flu resulted in even larger losses. The scale of the HIV and opioid epidemics were much smaller each year, but over decades became comparable to Covid-10 in terms of per-capita deaths and Spanish Flu in terms of life-years lost. The scenario of 250 thousand Covid-19 deaths is also shown with a dashed line. (Source: authors' calculations detailed in appendix 1).