



City Research Online

City St George's, University of London

Citation: Cairns, A. J. G. & Blake, D. (2026). ADM's APPLE: The Accelerated Deaths Model with an Application to the Covid-19 Pandemic. Insurance: Mathematics and Economics, 128, 103231. doi: 10.1016/j.insmatheco.2026.103231

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/36845/>

Link to published version: <https://doi.org/10.1016/j.insmatheco.2026.103231>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

ADM's APPLE: The Accelerated Deaths Model with an Application to the Covid-19 Pandemic

Andrew J.G. Cairns¹

David Blake²

Abstract

The Accelerated Deaths Model (ADM) builds on the hypothesis that, within a given age cohort, those who are less healthy are more likely to die if infected with Covid-19 than healthier people, leaving a pool of on-average healthier survivors. We use the term ‘detrimental selection’ which has two complementary aspects: the lower years of life lost by those who experienced an accelerated death; and the higher average life expectancy of survivors which we call their ‘adjusted post-pandemic life expectancy’ (ADM's APPLE). Our model represents a novel synthesis of recent advances in our understanding of mortality heterogeneity and the development of the Proportionality Hypothesis – both of which have improved our understanding of the Covid-19 pandemic. In particular, we identify an important positive relationship between mortality heterogeneity and accelerated deaths.

We find, in the case of the Covid-19 pandemic in England, that the years of life lost by those who experienced an accelerated death, while significantly lower than pre-Covid life expectancy, was greater than reported in the media at the time. We also find that the increase in the mean life expectancy of survivors was very small. As a result, the impact on annuity providers (e.g., in terms of potentially higher annuity prices), pension schemes and life insurers was also very small. In contrast, we find that the impact on life expectancy of a general change in future mortality assumptions post-pandemic (i.e., the base mortality table and improvement rate) would be much greater.

The ADM has potentially wide application, e.g., to other types of contagion and to climate-related deaths, where we would expect there to be a positive correlation between deaths and all-cause mortality (consistent with the Proportionality Hypothesis), but where the degree of detrimental selection might be different.

Keywords

Accelerated Deaths Model; Adjusted Post-Pandemic Life Expectancy; Covid-19; Detrimental selection; Biological frailty; Proportionality Hypothesis; Mortality heterogeneity.

¹ Department of Actuarial Mathematics and Statistics, Heriot-Watt University, UK and the Maxwell Institute for Mathematical Sciences, UK. (Corresponding author: A.J.G.Cairns@hw.ac.uk)

² Bayes Business School, City St George's, University of London, UK.

We are very grateful for the extremely useful comments and suggestions made by two referees.

1. Introduction

This paper develops the Accelerated Deaths Model (ADM). The idea for the model was motivated by the Covid-19 pandemic and claims in the media, such as the one reported by Knapton (2020), that up to two-thirds of Covid victims would have died within nine months from other causes even if they had not caught Covid. Many individuals who died from Covid-19 (especially younger people) had ‘underlying conditions’, suggesting they had a greater number of co-morbidities than would be normal for their age. Such individuals would have died at a later date in the absence of Covid-19, but sadly died from or with Covid-19 at an earlier (i.e., accelerated) date. As a direct counterpart to this, the mean life expectancy of survivors increased. The ‘years of life’ (YLL) lost by those who experienced an accelerated death and the consequential higher mean life expectancy of survivors are two complementary aspects of what we call ‘detrimental selection’.³

Cairns et al. (2020) introduced the ADM during the early stages of the pandemic. At the time, the predominant view of some medical experts, whose opinions were being widely reported, was that the mean YLL by those who died from Covid was very low. However, alternative research concluded that those who died from Covid during the first wave would, in the absence of Covid, have lived for an average of 9 to 12 additional years, e.g., Hanlon et al. (2020). Hanlon et al.’s analysis took account of the link between comorbidities and the risk of death from Covid-19 and found that the mean YLL by age was lower than the mean remaining life expectancy (LE) of survivors of the same age; however, importantly, it was still much higher than the ‘nine months’ quoted in the press. Potential links between individual co-morbidities and Covid death rates have also been investigated by, e.g., Banerjee et al. (2020) and Carannante et al. (2022).

The underlying hypothesis of the ADM is that, within a given age group, those who are less healthy are more likely to die if infected with Covid than healthier people of the same age, leaving a pool of survivors that is, on average, healthier. The evidence gathered since 2020 allows us to get a more accurate fix on both the number of accelerated deaths and the health profile of those affected.

We also build on two recent quantitative developments:

- the Proportionality Hypothesis (Cairns et al., 2025): the finding that Covid-19 death rates are proportional to ‘biological frailty’, defined as all-cause death rates by age and (e.g., socio-economic) subgroup, in a ‘normal’ non-Covid year;⁴
- quantification of the extent of heterogeneity in mortality at a highly granular level within a specific age group (see, e.g., Cairns et al., 2024, and Hansen et al., 2025, and references therein), allowing us to specify a more accurate post-Covid deaths curve for each cohort.

We use these developments to redesign the original ADM and assess the degree of detrimental selection induced by Covid-19. In particular, we use the ADM to calculate ‘adjusted post-pandemic life expectancies’ (ADM’s APPLEs).

Applying the ADM to the Covid-19 pandemic in England, we find that:

- the mean years of life lost by those who died from Covid was below the mean pre-pandemic life expectancy, but was significantly above 0 (and closer to the of 9 to 12 years estimated by

³ As a concept, detrimental selection differs from ‘adverse’ selection: the latter involves an element of choice, whereas the former does not. An alternative term sometimes used in the actuarial literature is ‘anti-selection’ (see, e.g., Cairns et al., 2020).

⁴ ‘Biological frailty’ links directly to the better-known concept of ‘biological age’ (see, for example, Milevsky, 2020). People who were born in *different* years (so have a different ‘chronological age’) but who have the *same* individual prospective death rate (e.g., based on their state of health) are said to have the same biological age. This common prospective death rate is what Cairns et al. (2025) refer to as biological frailty.

Hanlon et al., 2020, than the nine months reported in the media near the start of the pandemic);

- the impact on the mean life expectancy of the surviving population turned out to be quite small (and smaller than that calculated by Cairns et al., 2020);
- empirically, the profile of Covid deaths turned out to have an inverted S shape which converges on the right to a strictly positive value above zero; this contrasts with the proposal in Cairns et al. (2020) of an exponential curve that decays to zero.

The ADM model is useful for a number of reasons. First, it allows providers of annuities and pensions, such as life insurers and pension schemes, to answer the following types of questions: ‘does the pandemic imply that the surviving population is, on average, healthier and longer lived than the average before the pandemic?’; and, if so, ‘what does it mean for the fair pricing of annuities – should prices be increased?’. Second, the ADM can be applied in other pandemics (especially those involving ‘novel’ viruses where the population has no prior exposure and, therefore, no prior immunity) where we would expect the Proportionality Hypothesis to hold, but where the magnitude of the pandemic might be different. The ADM might also be applicable in the case of accelerated deaths due to climate change (e.g., from extreme heatwaves).

The outline of the paper is as follows. Section 2 discusses the cohort deaths curve, both pre- and post-pandemic. Section 3 introduces the ADM. Section 4 considers some numerical illustrations of the ADM, including parameterisations consistent with the Covid-19 pandemic in England. Section 5 considers some implications of the ADM for life insurers. Section 6 looks at further potential applications of the ADM, such as seasonal flu and climate change. Section 7 concludes.

2. The cohort deaths curve pre- and post-pandemic

We focus initially on an age cohort’s aggregate deaths curve, $d(t, x)$, which shows the number of people in a cohort with an initial (or current) age x who die at age $x + t$ (i.e., t years later). A typical (stylised) example of a pre-pandemic deaths curve is presented by the solid black curve in Figure 1 for a cohort initially aged 75. As the cohort ages, its members become increasingly frail and the annual number of deaths rises, but the curve then peaks and begins to decline as the number of survivors in the cohort falls faster than mortality rates increase with age.

Members of the cohort are not all equal in terms their health, particularly, at the national level: some are in good health, while others are in poor health. A key idea underpinning our study is that people who are currently in poor health are likely to die earlier, on average, than people who are currently in good health. In the deaths curve illustrated in Figure 1, out of those currently aged 75 who die at each future age, the balance between those who are currently healthy and those who are currently unhealthy gradually shifts. Deaths in the near future (e.g., ages at death of 75 to 80) will be drawn more from those who are currently unhealthy, while deaths that occur far in the future (e.g., 95 to 100) will be mainly drawn from those who are currently in comparatively good health.




Figure 1: A stylised deaths curve for a cohort of individuals currently aged $x = 75$ in the absence of the Covid-19 pandemic, where x is the initial age, t is the number of years until death, and $x + t$ is the age at death.

Our key hypothesis (which we denote the ‘Accelerated Deaths Hypothesis’) is that, if they were to become infected, people who are in relatively poor health (relative to other people in the same age cohort) are more likely to die from Covid-19 compared with newly infected people who are in relatively better health. This contrasts with the alternative ‘Independence Hypothesis’ which states that the probability of dying from Covid, *if infected*, is independent of an individual’s state of health relative to others in their age cohort. The impact of these two hypotheses (in an exaggerated form to aid visual comprehension) is presented in Figure 2. In both panels of the Figure, the dashed curve is the deaths curve of the age-75 cohort in the absence of the Covid pandemic, while the solid curve is the deaths curve of the age-75 cohort who survived the pandemic. Taking a vertical slice at a given age x , the hatched region represents the number of people who were originally expected to die at age $x + t$ but who, instead, had an accelerated death due to Covid. The left plot is consistent with the Accelerated Deaths Hypothesis, under which, at each age $x + t$, Covid deaths are drawn predominantly from those who are less healthy – so that the hatched region, which represents Covid deaths, is more heavily weighted towards the left. The right plot corresponds to the Independence Hypothesis, under which Covid deaths are drawn at random from the cohort, independently of their current health status. The shape of the hatched regions up to each age $x + t$ shows the degree of detrimental selection up to that age under the two hypotheses. Because we are not interested in modelling the course of the pandemic itself, we will also assume, for convenience, that the whole pandemic occurs instantaneously (i.e., at $t = 0$) rather than spread out over the approximate two-year period of the actual pandemic.




Figure 2: Stylised deaths curves for a cohort of individuals currently aged $x = 75$ showing the impact of the Covid-19 pandemic. The hatched regions represent people who were expected to die at a later age but who die instead from Covid – the accelerated deaths. All of the accelerated deaths are then shifted to the thick blue bar on the left which represent the deaths at age 75 under the assumption that the whole pandemic occurs instantaneously. In each plot, we assume that 5% of the deaths are accelerated, i.e., they die immediately from Covid rather than die at the age at which they were originally expected. In the left panel, deaths are initially drawn from those who are less healthy, consistent with the Accelerated Deaths Hypothesis. In the right panel, deaths are drawn randomly, consistent with the Independence Hypothesis.

We will show that the data (for England) does support the Accelerated Deaths Hypothesis over the Independence Hypothesis. However, the modelling challenge is to determine how heavily the hatched region should be weighted towards the left and what shape it should take in the light of data collected in the early stages of the pandemic.

3. Building the Accelerated Deaths Model

In this section, we build the ADM. The particular ADM that we build can be thought of as a top-down macro-type model that uses disaggregated national data down to the highest level of disaggregation that is publicly available. An alternative approach is to construct a bottom-up micro-type model which uses, say, a Susceptible-Infected-Recovered-Deceased (SIRD) model of the individual and then aggregates up to the national level to estimate the number of accelerated deaths from the pandemic. Chen et al. (2025) used the latter model to design a catastrophe bond to hedge pandemic-induced mortality risks in the insurance sector.

There are two key factors that need to be taken into account in the light of recent modelling developments: the Proportionality Hypothesis and mortality heterogeneity within a specific age cohort. In addition, the ADM has a number of components or inputs that need to be calibrated (in some cases to be consistent with certain outputs): survivor curves in the absence of Covid-19; infection rates and infection fatality rates arising from the pandemic; and how survivor curves are modified following the pandemic. Since we are assuming that the whole pandemic occurs at time $t = 0$, the ADM also requires us to pre-specify the magnitude of pandemic in terms of total deaths: this is despite the fact that the actual magnitude will not be known with certainty until after the pandemic has

been declared over by the relevant health authority. The outputs of the ADM will be: pre- and post-pandemic life expectancies and years of life lost.

3.1 Two key factors: the Proportionality Hypothesis and cohort mortality heterogeneity

Cairns et al. (2025) proposed and investigated the Proportionality Hypothesis in the context of the Covid pandemic. Their conclusion was that, *given an individual has become infected*, their probability of dying from Covid is approximately proportional to their biological frailty, i.e., their all-cause death rate in a normal non-Covid year. They proposed subgroup and aggregate versions of the hypothesis.⁵

At the subgroup level, Covid-19 death rates can be decomposed as follows:

$$m_C(i, x) \equiv m(i, x)IR(i, x)RF(i, x), \quad (1)$$

where: $m_C(i, x)$ is the Covid-19 death rate in socio-economic subgroup i , at age x ; $m(i, x)$ is the corresponding all-cause death rate in a normal non-Covid year; $IR(i, x)$ is the infection rate; and $RF(i, x)$ is referred to as the ‘relative frailty’ (it also serves as the balancing term in the identity). Relative frailty measures the risk of death from Covid-19 by age and subgroup for a person newly infected with Covid-19 relative to the risk of death from all causes in a normal non-Covid year.⁶ (For a full summary of the notation and functions used in this paper, see Appendix A.)

Cairns et al. (2025)’s ‘Subgroup Proportionality Hypothesis’ states that relative frailty is approximately constant across all subgroups, i , and ages, x .⁷ We will exploit this hypothesis and its impact on the identity in (1) to model, in a precise way, what shape the hatched area in the left panel of Figure 2 should take.

We also need to take into account how much variation there is in the underlying all-cause death rates $m(i, x)$ between socio-economic subgroups and even between individuals. Heterogeneity in mortality rates (also known as mortality inequality) is a much-studied subject (see, e.g., Shkolnikov et al., 2012, Rashid et al., 2021, and Cairns et al., 2024, and references therein). In our paper, heterogeneity is assessed using the concept of ‘relative risk’, R , which measures the risk of death at a subgroup level (or even individual level) relative to national death rates at the same age. We will associate the ‘heterogeneity in mortality rates’ with the variability in R .

Mortality heterogeneity between subgroups occurs for a variety of reasons: variation in the prevalence of controllable risk factors (e.g., smoking and excess alcohol consumption; see Cairns and Redondo-Lourés, 2023); variation in the prevalence of non-controllable risk factors (e.g., genetic); and random good or bad luck in terms of the onset of specific diseases and co-morbidities.⁸ Using data for England, two studies by Wen et al. (2023) and Cairns et al. (2024)⁹ present detailed analyses of death rates by age at the level of small individual neighbourhoods, known as Lower Layer Super Output

⁵ We do not make use of the aggregate versions of the Proportionality Hypothesis in this study.

⁶ We will see below in Equation (8) that relative frailty is closely linked to the infection fatality rate.

⁷ There are three alternative, but equivalent, ways of expressing the Proportionality Hypothesis. To illustrate, the subgroup version can be expressed in the following ways:

- The Covid-19 infection fatality rate, $IFR(i, x)$, is approximately proportional to the all-cause death rate, $m(i, x)$, by both age and subgroup.
- The Covid-19 death rate, $m_C(i, x)$, is approximately proportional to the product of the all-cause death rate and the annualised infection rate by both age and subgroup, $m(i, x)IR(i, x)$.
- Relative frailty, $RF(i, x)$, is approximately constant across ages and subgroups.

⁸ For larger neighbourhoods or subgroups, random good or bad luck in terms of the onset of diseases will average out to some extent, leaving variation in the prevalence of controllable and non-controllable risk factors as the main drivers of heterogeneity.

⁹ See, also, Cairns et al. (2025).

Areas (LSOAs).¹⁰ They reveal considerable levels of mortality heterogeneity, especially in the 40-49 and 50-59 age groups. Other studies, using even more granular data (e.g., Savcisen et al., 2024, and Hansen et al., 2025, exploiting Danish data at the individual level), offer the potential to measure the mortality risks of individuals relative to national death rates. Often, though, these studies do not report the distribution of relative risks across the whole population (i.e., how much heterogeneity there is within a whole cohort), but it is reasonable to assume that, within each neighbourhood or socio-economic subgroup, individual mortality rates will be quite varied relative to the subgroup mean for a given age. Thus, we conjecture that variability in relative risks between individuals will be (potentially significantly) higher than the variability between subgroups for the following type of reasons:

- of two otherwise similar individuals within a subgroup, one is a smoker and the other is not, with the former having a higher relative risk;
- of two otherwise similar individuals within a subgroup, one has randomly, and through bad luck, developed breast cancer, while the other did not.

3.2 Components of the ADM I: Survivor curves in the absence of Covid-19

3.2.1 Survivor curves in a normal non-Covid year

A survivor curve represents the proportion of a cohort that survives from initial age x to a later age $x + t$. We make use of three types of survivor curve: aggregate; baseline; and subgroup. The aggregate survivor curve relates to a high-level group, such as the national population. The baseline curve refers to a hypothetical homogeneous subgroup of the aggregate population whose characteristics do not change as the population ages. While the aggregate survivor curve is observable, the baseline survivor curve is not, but it can be inferred (or backed out) from the aggregate survivor curve, given the distribution function for relative risk, R . The subgroups are subdivisions of the aggregate population, with either lighter or heavier mortality than the baseline, as measured by the subgroup relative risk, R .

At this point, and for the remainder of the paper, we switch from using the (discrete-time) annual death rate, $m(\cdot)$, to (its continuous-time equivalent) the instantaneous force of mortality, $\mu(\cdot)$. The purpose of this is to allow us to develop precise equations for the death and survivor curves and other quantities in continuous time.

We begin by defining the cohort baseline force-of-mortality curve, $\mu_B(t, x)$, where x is the initial age of a cohort at time 0, and $t > 0$ is the time. This is the force-of-mortality curve for the hypothetical homogeneous subgroup of the aggregate population with fixed health characteristics, against which all other subgroups and individuals are compared. The ‘baseline survivor curve’ is:

$$S_B(t, x) = \exp \left[- \int_0^t \mu_B(s, x) ds \right]. \quad (2)$$

We define relative risk as a random variable, $R > 0$, that differs across subgroups. For the purpose of this study, within each subgroup, R will be assumed to take a fixed value for all $t > 0$ and for all individuals in that subgroup. For such a subgroup, the force of mortality is $\mu(R, t, x) \equiv R\mu_B(t, x)$ and the ‘subgroup survivor curve’ is:

$$S(R, t, x) = \exp \left[- \int_0^t \mu(R, s, x) ds \right] = \exp \left[- \int_0^t R\mu_B(s, x) ds \right] = S_B(t, x)^R. \quad (3)$$

For consistency with Cairns et al. (2025), we refer to $\mu(R, t, x)$ as the ‘biological frailty’ of an individual within a subgroup with relative risk R : biological frailty is equal to the all-cause force of mortality in a normal non-Covid year. Our model assumes that all subgroups are initially of equal size

¹⁰ In England, there are 32,844 LSOAs with an average population of around 1,600 people of all ages.

at age 60, and that, empirically, the relative risks follow a specified probability distribution (with a mean of 1). We discuss the choice of distribution in Section 4.1.

The ‘aggregate survivor curve’ averages the subgroup survivor curves over all possible values of R . In theory this implies that:

$$S(t, x) = E[S(R, t, x)] = \int_0^\infty S(r, t, x) f(r) dr = \int_0^\infty S_B(t, x)^r f(r) dr, \quad (4)$$

where $f(r)$ is the density function for the relative risk R at age x across the whole population. In practice, however, we approximate the final term in Equation (4) by averaging over independent draws from the chosen distribution function for R (i.e., $R_1, \dots, R_N$, where N is the number of subgroups):

$$S(t, x) = N^{-1} \sum_{i=1}^N S(R_i, t, x) = N^{-1} \sum_{i=1}^N S_B(t, x)^{R_i}. \quad (5)$$

In short, we subdivide the general population into $i = 1, \dots, N$ subgroups, with each subgroup having its own relative risk R_i (in place of the generic relative risk r). Each subgroup i therefore represents a homogeneous subgroup with relative risk R_i . $S(t, x)$ is the average across the subgroups.

Consequently, the aggregate force of mortality averages the subgroup forces of mortality over all possible values of R :

$$\mu(0, x) = - \left. \frac{\partial S(t, x) / \partial t}{S(t, x)} \right|_{t=0} = \frac{-1}{NS(t, x)} \sum_{i=1}^N R_i S_B(t, x)^{R_i-1} \left. \frac{\partial S_B(t, x)}{\partial t} \right|_{t=0} = N^{-1} \sum_{i=1}^N R_i \mu_B(0, x). \quad (6)$$

For general $t > 0$, the formula for $\mu(t, x)$ is more complex as the N subgroups decline in size at different rates:

$$\mu(t, x) = \sum_{i=1}^N w_i(t) R_i \mu_B(t, x), \quad (7)$$

where the weights, $w_i(t)$, are proportional to the size of each subgroup at time t , $S(R_i, t, x)$.

An example of the relationship between the aggregate and baseline force of mortality can be seen in Figure 3. The relationship between the two curves can be captured by the following stylised facts:

- At younger ages, the two curves are quite close to each other.
- As the cohort ages, the aggregate curve starts to fall away from the baseline curve as the subgroup weights, $w_i(t)$, in Equation (7) shift towards healthier subgroups (i.e., those with lower relative risks, R_i).
- The gap between the two curves depends on the degree of mortality heterogeneity: the more dispersed is the distribution of R , the wider the gap between the baseline and aggregate force of mortality curves.




Figure 3: Baseline, $\mu_B(t, x)$, and aggregate, $\mu(t, x)$, force of mortality curves for Scenario A (in Table 2).

3.2.2 Survivor curve calibration

Since we are unable to observe the baseline force of mortality or survivor curves directly, we must infer these from what can be observed, namely the aggregate force of mortality, $\mu(t, x)$, and the corresponding aggregate survivor curves (see Equations (4) and (5)). The baseline survivor curve is then calibrated from these observable functions together with a specified theoretical distribution for R (or an empirical distribution of R over the N subgroups if data are available) to ensure that $N^{-1} \sum_{i=1}^N S_B(t, x)^{R_i} = S(t, x)$ (see Equation (5)). In the numerical examples later in this paper, we will use the Beard law of mortality (see Equation (31) below and, e.g., Richards, 2012) to model the aggregate survivor curve. This implies that the force of mortality curve is similar to the Gompertz curve at younger ages, but gradually flattens off at very high ages.

3.3 Components of the ADM II: Survivor curves in the presence of Covid-19

During the Covid-19 pandemic, two key variables of interest were the infection rate and the infection fatality rate (Cairns et al., 2025). By specifying assumptions about these variables, we can modify the survivor curves in the presence of Covid-19.

3.3.1 The infection rate

We will denote the infection rate, IR , by α . It is defined as the proportion of the population or subgroup that becomes infected with Covid-19 over a given period. At the most general level, $\alpha = \alpha(i, x)$, which allows for the possible dependence not only on a subgroup's current age, x , but also its relative risk, R_i , and other socio-economic or behavioural characteristics. A high value for the relative risk could potentially push infection rates up (e.g., if individuals with a higher relative risk are more susceptible to becoming infected, given they have been exposed to the virus¹¹) or down (if individuals

¹¹ For example, their immune systems are less able to fight off the initial exposure. This extrapolates (somewhat speculatively) the conclusions of Cairns et al. (2025) that individuals with a higher biological age (being a combination of chronological age and relative risk) are more likely to die following infection from Covid: once infected, their immune systems are less able to protect them from hospitalisation or death.

with higher relative risk behave in a way that reduces the chance of exposure to the virus). The same arguments can apply to infection rates by age.

In our numerical examples in Section 4, we will assume that $\alpha(i, x) = \alpha$ is constant. There are, so far, no publicly available data to analyse how infection rates depend on the relative risk and so we have little choice but to assume a constant infection rate (at least within a given age group) in the current version of the ADM. Everything else being equal, if α were twice as high at a given age, this would lead to a doubling of the number of deaths, and an approximate doubling of the required adjustment to the life expectancy from pre- to post-pandemic, consistent with the Proportionality Hypothesis. (2025).

3.3.2 The infection fatality rate

Building on the Proportionality Hypothesis, we now consider the infection fatality rate, *IFR*. Cairns et al. (2025) define the *IFR* as a *rate* from which we derive the *probability* of death, given that an individual is newly infected, that is:

$$\begin{aligned} &\Pr(\text{die due to or with Covid} \mid \text{newly infected with Covid, age } x, \text{ relative risk } R) \\ &= 1 - \exp(-IFR(R, x)). \end{aligned}$$

Note that if the *IFR* is small, then $1 - \exp(-IFR(R, x)) \approx IFR(R, x)$; this implies that the *IFR* can be considered to be an approximation to the probability of death for someone who is newly infected.

In the context of the Proportionality Hypothesis and, given our discussion above, the *IFR* for a subgroup with relative risk R_i is:

$$IFR(R_i, x) \equiv \phi \mu(R_i, 0, x) = \phi R_i \mu_B(0, x), \quad (8)$$

where we denote relative frailty, *RF*, by ϕ . The *IFR*, as defined in Equation (8), is an adaptation of Equation (2) in Cairns et al. (2025), where we replace the death rate $m(i, x)$ in subgroup i with the force of mortality $\mu(R_i, 0, x)$, where R_i is the relative risk in subgroup i and $t = 0$ in $\mu(R_i, 0, x)$.

To be consistent with the Proportionality Hypothesis, ϕ , in Equation (8), must be (approximately) constant across ages and subgroups (see Section 3.1). The *IFR* can be thought of as an intermediate quantity in the ADM, since it depends on a primary input parameter ϕ . If ϕ doubled then the *IFR* would also double.

3.3.3 Modifying the survivor curves in the presence of Covid-19

We now modify the survivor curves to account for a Covid pandemic. We are interested in the total impact of the pandemic, so we assume that all Covid deaths occur instantaneously at time $t = 0$. This assumption reflects the fact that *we are not attempting to model the progress of the pandemic itself: instead, we seek to model its impact on the surviving population.*

We assume a total of N subgroups at age x , with each subgroup, i , having its own relative risk, R_i , drawn from the specified distribution for relative risks. Immediately prior to the pandemic at time 0, all N subgroups have the same number of lives. The pandemic then occurs at time 0 and the survivor curve for subgroup i (denoted $S_C(R_i, t, x)$ in the post-pandemic scenario) immediately drops from 1 to

$$S_C(R_i, 0^+, x) = \alpha \exp(-\phi R_i \mu_B(0, x)) + (1 - \alpha) \approx 1 - \alpha \phi R_i \mu_B(0, x), \quad (9)$$

where 0^+ means at time 0 but immediately after the Covid deaths occurred.^{12,13} The dependence on R_i in Equation (9) ensures that the immediate drop from 1 to $S_C(R_i, 0^+, x)$ at time 0 is bigger for subgroups with a higher value of R_i .

Immediately after the pandemic, we initially assume that subgroup forces of mortality revert to pre-pandemic levels, i.e., $R_i\mu_B(t, x)$ at the future age of $x + t$. Thus, for $t > 0$,

$$S_C(R_i, t, x) = S_C(R_i, 0^+, x)S(R_i, t, x) = S_C(R_i, 0^+, x)S_B(t, x)^{R_i} \quad (10)$$

with an aggregate post-pandemic survivor curve found by averaging across all the subgroup post-pandemic survivor curves:

$$S_C(t, x) = N^{-1} \sum_{i=1}^N S_C(R_i, t, x). \quad (11)$$

Later in the paper (in section 4.5), we will modify the model to allow for post-pandemic changes in the base table for mortality (i.e., at time 0^+) and in future mortality improvement rates.

The differing values of $S_C(R_i, 0^+, x)$ in Equation (10) for different subgroups imply that the aggregate post-pandemic force of mortality

$$\mu_C(t, x) = -\frac{\partial S_C(t, x)/\partial t}{S_C(t, x)} \quad (12)$$

differs from that pre-pandemic (Equation (6)). In particular, for $t = 0^+$:

$$\mu_C(0^+, x) = \sum_{i=1}^N \tilde{w}_i(0^+) R_i \mu_B(0, x) \quad (13)$$

where $\tilde{w}_i(0^+) = S_C(R_i, 0^+, x)/NS_C(0^+, x)$ are the revised weights attached to each subgroup, contrasting with Equation (6), where the weights at time 0 are all equal to $1/N$ (since we assume that, just prior to the pandemic, all subgroups are of equal size).

3.4 Total deaths and magnitude

The overall impact of a pandemic can be measured in two ways: the absolute number of deaths; and the pandemic's 'magnitude'. The latter is defined as the total number of pandemic related deaths at age x (over the full duration of the pandemic) divided by expected deaths at age x in a normal non-pandemic year. In this study, we concentrate on the magnitude as this measure is independent of the size of a country's population. The total number of deaths can be determined by combining the country-specific populations at each age, forces of mortality by age, and magnitude by age.

The 'magnitude' of the pandemic is defined, for age x , as:

$$M_1(x) \equiv \frac{1 - S_C(0^+, x)}{\mu(0, x)}, \quad (14)$$

i.e., the ratio of those, across the whole population, who die from Covid-19 over the full duration of the pandemic to those who die in a normal non-Covid year, with the latter approximated by the aggregate force of mortality $\mu(0, x) = N^{-1} \sum_{i=1}^N R_i \mu_B(0, x)$ (defined in Equation (6)).

We also note that:

¹² The approximation in the final term on the RHS in Equation (9) (and those in Section 3.5 reliant on (9)) is provided as a convenient simplification which helps interpret the results (and is equal to $(1 - (IR \times IFR))$). In our numerical calculations, the exact expression (i.e., the first term on the RHS of the equation) is used.

¹³ Note that, everything else being equal, the proportion who die from Covid, $(1 - S_C(R_i, 0^+, x))$, is approximately proportional to the proportion infected, α , consistent with the Proportionality Hypothesis.

$$M_1(x) = \frac{\alpha \sum_{i=1}^N (1 - \exp(-\phi R_i \mu_B(0, x)))}{\sum_{i=1}^N R_i \mu_B(0, x)} \approx \frac{\alpha \sum_{i=1}^N \phi R_i \mu_B(0, x)}{\sum_{i=1}^N R_i \mu_B(0, x)} = \alpha \phi \equiv M_2(x), \quad (15)$$

where we define the ‘simplified magnitude’ $M_2(x) \equiv \alpha \phi$ as a useful approximation to $M_1(x)$. While $M_1(x)$ is constrained on the upside by the size of the population at age x , there is no such requirement for an upper bound on $M_2(x)$: whatever the value of $M_2(x) > 0$, $M_1(x)$ will automatically satisfy its upper population constraint.¹⁴ For this reason, we choose to focus on the simplified magnitude, $M_2(x)$, in our numerical analysis of the ADM.

As mentioned above, the true magnitude of the pandemic cannot be determined accurately until the pandemic is over. Nevertheless, $M_2(x)$ is an input parameter in the ADM and so a value for it has to be specified. Further, it is clear from (15) that, for a given age, x , the magnitude, $M_2(x)$, the infection rate, α , and relative frailty, ϕ , are interrelated. Since ϕ is unobservable, we will assume that it is determined as follows: $\phi \equiv M_2(x)/\alpha$.

3.5 Outputs of the ADM: Pre- and post-Covid life expectancies and mean years of life lost

The principal outputs of the ADM are pre-and post-Covid life expectancies and years of life lost.

We are interested in (remaining) life expectancies (LEs) for three populations: pre-pandemic LE for the whole population; mean years of life lost (YLL) by those who die from Covid; and adjusted post-pandemic life expectancy (APPLE) for the subset who survive the pandemic.

We begin by calculating the pre-pandemic life expectancy for each of the subgroups $i = 1, \dots, N$:

$$LE(R_i, x) = \int_0^\infty S(R_i, t, x) dt = \int_0^\infty S_B(t, x)^{R_i} dt. \quad (16)$$

Note from (16) that $LE(R_i, x)$ is a decreasing function of R_i :¹⁵ subgroups with a higher relative risk live shorter lifetimes on average. Now we calculate the (remaining) life expectancies for three populations in turn.

3.5.1 The mean pre-pandemic life expectancy (remaining)

The mean pre-pandemic life expectancy (remaining) is given by:

$$LE(x) = \int_0^\infty S(t, x) dt = \frac{1}{N} \sum_{i=1}^N LE(R_i, x), \quad (17)$$

with $S(t, x)$ defined in Equation (4). It would appear that $LE(x)$, through the second equality in Equation (17), depends on the distribution of R via the dependence of $LE(R_i, x)$ on both the individual values of R_i and the distribution of the R_i . For example, the $S(R_i, t, x)$ in Equation (16) do depend to a small extent on the distribution of R (via the R_i terms) because this distribution has a small impact on the baseline force of mortality, $\mu_B(t, x)$, and consequently the baseline survivor curve, $S_B(t, x)$ (see Equations (2) and (3) and Figure B1 in Appendix B.1). In Appendix B.3, we prove the second equality in (17). Specifically, we show that, with the decomposition of the aggregate population into subgroups, the aggregate-level $LE(x)$ (the first equality in (17)) is exactly equal to the subgroup mean of the $LE(R_i, x)$ in (17) (the second equality). So, even though the $LE(R_i, x)$ are a

¹⁴ As an analogy, it is the same as the difference between the mortality rate, $q(x)$, and the death rate, $m(x)$. The former must lie between 0 and 1, while the latter only needs to be positive. We can multiply $m(x)$ by 10 and the corresponding $q(x)$ will stay between 0 and 1.

¹⁵ This is demonstrated as follows:

$$\frac{\partial LE}{\partial R_i}(R_i, x) = \int_0^\infty \log(S_B(t, x)) \cdot S_B(t, x)^{R_i-1} dt < 0 \text{ since } \log(S_B(t, x)) < 0.$$

function of R_i and hence do have a small dependence on the distribution of R , the aggregate $LE(x)$ does not.

3.5.2 The mean years of life lost

The mean years of life lost (per person aged x , for consistency with LE and APPLE) by those who die from Covid is given by:

$$YLL(x) = \frac{\sum_{i=1}^N (1 - S_C(R_i, 0^+, x)) LE(R_i, x)}{\sum_{i=1}^N (1 - S_C(R_i, 0^+, x))} = \frac{N^{-1} \sum_{i=1}^N (1 - S_C(R_i, 0^+, x)) LE(R_i, x)}{1 - S_C(0^+, x)} \quad (18)$$

where $(1 - S_C(R_i, 0^+, x))$ is the proportion of subgroup i that dies from Covid. Equation (18) is a weighted average of the pre-pandemic LEs where the weights are proportional to the numbers in each subgroup that die from Covid.

We can use the approximation in Equation (9) to show that (see Appendix B.1):

$$YLL(x) \approx \frac{1}{N} \sum_{i=1}^N R_i LE(R_i, x). \quad (19)$$

YLL in Equation (19) is the mean of the pre-pandemic LEs for the subset who become infected and then die from Covid. Consequently, YLL measures the detrimental selection experienced by these individuals.

In contrast with the second equality for $LE(x)$ in Equation (17), which is an unweighted average of the $LE(R_i, x)$, $YLL(x)$ in Equation (19) can be seen to be a weighted average of the $LE(R_i, x)$ with the weights proportional to the R_i . $YLL(x)$, therefore, does depend on the distribution of R , both directly and via its interaction with $LE(R_i, x)$. We have also seen from our comments on Equation (16) that $LE(R_i, x)$ is a decreasing function of R_i . If we combine this with the weights in Equation (19), we can see that $YLL(x)$ will be less than $LE(x)$, the unweighted average. Finally, the more dispersed the R_i are, the more these weights shift towards low values of $LE(R_i, x)$, lowering $YLL(x)$ further. (This is confirmed by the approximations (B6) and (B9) in Appendix B.1.1.)

However, $YLL(x)$ does *not* depend on parameters α and ϕ , and hence the magnitude of the pandemic (since $M_2(x) \equiv \alpha\phi$). This can also be seen from Equation (18) which shows that the subgroup proportions dying from Covid, $(1 - S_C(R_i, 0^+, x))$ and the aggregate $S_C(0^+, x)$ are exactly proportional to α (and approximately proportional to M_2). So doubling of α (or M_2) has no (or only a small) effect on the YLL, since both the numerator and denominator in Equation (18) change by exactly (or approximately) the same percentage.

The intuition for this is clear. The average years of life lost (per person) by those who die from Covid should not depend (very much) on the infection rate or the scale of the pandemic. However, having been infected, dying from Covid will depend on the infection fatality rate which does depend on the R_i (via Equation (8)). This explains why $YLL(x)$ depends *only* on the distribution of R .

It follows from Equations (17) and (19) that (see Appendix B.1)

$$LE(x) - YLL(x) \approx \frac{1}{N} \sum_{i=1}^N (1 - R_i) LE(R_i, x) \quad (20)$$

measures the extent to which those who die from Covid are less healthy compared to the average for their age cohort. We can therefore define the impact of detrimental selection on YLL as:

$$100 \left(1 - \frac{YLL(x)}{LE(x)} \right). \quad (21)$$

As we have just seen, $YLL(x)$ decreases as the dispersion of the R_i terms increases, while $LE(x)$ is fixed. Hence, both $(LE(x) - YLL(x))$ in Equation (20) and the detrimental selection impact on YLL in Equation (21) both increase as the dispersion increases.

We can demonstrate this more clearly by further developing Equation (20) to show that (see Appendix B.1):

$$LE(x) - YLL(x) \approx \lambda_1(x)N^{-1} \sum_{i=1}^N (R_i - 1) \log(R_i). \quad (22)$$

The justification for this expression is that, at a given age x , $LE(R_i, x) \approx \lambda_0(x) - \lambda_1(x) \log(R_i)$, where $-\lambda_1(x)$ is the rate of change in $LE(R_i, x)$ with respect to $\log(R_i)$ evaluated at $R_i = 1$. This linear relationship can be seen in the top right panel of Figure B1 in Appendix B.1.

Since, in addition, $\log(R_i) \approx R_i - 1$ when R_i is close to 1, we have a further approximation:

$$LE(x) - YLL(x) \approx \lambda_1(x)N^{-1} \sum_{i=1}^N (R_i - 1)^2 = \lambda_1(x)\text{Var}(R). \quad (23)$$

This means that the difference between $LE(x)$ and $YLL(x)$ is approximately proportional to the variance of R . However, when R is not tightly distributed around 1, Equation (22) gives us a more accurate way to compare the impacts on $(LE(x) - YLL(x))$ of different distributions for R . For further details, see Appendix B.1. The key point to note is that $(LE(x) - YLL(x))$ increases and $YLL(x)$ decreases as the dispersion of R increases.

3.5.3 The mean adjusted post-pandemic life expectancy (remaining)

The mean adjusted post-pandemic life expectancy (remaining) is given by:

$$APPLE(x) = \frac{\sum_{i=1}^N S_C(R_i, 0^+, x) LE(R_i, x)}{\sum_{i=1}^N S_C(R_i, 0^+, x)} = \frac{N^{-1} \sum_{i=1}^N S_C(R_i, 0^+, x) LE(R_i, x)}{S_C(0^+, x)}. \quad (24)$$

We can again use the approximation in Equation (9) to show that (see Appendix B.1):

$$\begin{aligned} APPLE(x) &\approx \frac{LE(x) - M_2(x)\mu_B(0, x)YLL(x)}{1 - M_2(x)\mu_B(0, x)} \\ &\approx LE(x) + M_2(x)\mu_B(0, x)(LE(x) - YLL(x)), \end{aligned} \quad (25)$$

and hence

$$APPLE(x) - LE(x) \approx M_2(x)\mu_B(0, x)(LE(x) - YLL(x)). \quad (26)$$

Equation (26) shows that the difference between $APPLE(x)$ and $LE(x)$ depends on the difference between $LE(x)$ and $YLL(x)$. We know from Equations (20), (22) and (23) that $(LE(x) - YLL(x))$ increases as the dispersion of R increases. Hence, $APPLE(x)$ and $(APPLE(x) - LE(x))$ both increase as the dispersion of R increases. But these increases are a direct consequence of the increase in $(LE(x) - YLL(x))$ caused by the increased dispersion of R .

However, unlike $YLL(x)$, $(APPLE(x) - LE(x))$ and $APPLE(x)$ do depend on $M_2(x)$: a higher number of Covid deaths amongst less healthy members of each age cohort will increase the average life expectancy of the healthier group of post-pandemic survivors. Doubling $M_2(x)$ (through a combination of increases in α and/or ϕ) will approximately double $(APPLE(x) - LE(x))$. However, $M_2(x)$ is further scaled by $\mu_B(0, x)$ which will take a relatively low value (except when x gets very large), so $M_2(x)\mu_B(0, x)$ will typically be much less than 1 even if $M_2(x) \gg 1$. Hence, we would expect to find that, following the Covid pandemic, the increase in $(APPLE(x) - LE(x))$ is much less

than the increase in $(LE(x) - YLL(x))$ and the reduction in $YLL(x)$. This is clearly demonstrated in Figure 9 below.

The percentage increase from $LE(x)$ to $APPLE(x)$ gives a good indication of the impact of detrimental selection on the life expectancy of survivors:

$$100 \left(\frac{APPLE(x)}{LE(x)} - 1 \right) \approx 100 \frac{M_2(x) \mu_B(0, x) (LE(x) - YLL(x))}{LE(x)}. \quad (27)$$

This shows that the detrimental selection impact is directly proportional to the scale of the pandemic (it doubles if M_2 doubles). It also *increases* as YLL *decreases*. But, for reasons explained in the previous paragraph, the increase is much less than the decrease in $YLL(x)$. Equation (27) is the direct complement to Equation (21).

3.5.4 The relationship between $LE(x) - YLL(x)$ and $APPLE(x) - LE(x)$

From equations (17), (18) and (24), there is an exact relationship between $LE(x)$, $YLL(x)$ and $APPLE(x)$:

$$LE(x) = (1 - S_C(0^+, x)) YLL(x) + S_C(0^+, x) APPLE(x). \quad (28)$$

Equation (28) can be rearranged as follows:

$$APPLE(x) - LE(x) = (LE(x) - YLL(x)) \frac{(1 - S_C(0^+, x))}{S_C(0^+, x)}. \quad (29)$$

Since $S_C(0^+, x)$ is very close to 1 and since, correspondingly, $(1 - S_C(0^+, x)) \ll 1$, Equation (29) shows that an increase in $(APPLE(x) - LE(x))$ is directly proportional to, but much smaller than, any change in $(LE(x) - YLL(x))$. This confirms what we have demonstrated in Sections 3.5.2 and 3.5.3.

Finally, note that as the dispersion of R tends to 0, $APPLE(x)$ and $YLL(x)$ both tend to $LE(x)$ in Equation (28) (i.e., $APPLE(x) = LE(x) = YLL(x)$; see Appendix B.1.3). As the dispersion of R increases from 0, $YLL(x)$ begins to fall and $APPLE(x)$ begins to rise in a way that ensures that $LE(x)$ in Equation (28) does not change.

3.6 Inputs vs outputs in the ADM

We summarise the inputs and outputs of the ADM as follows:

- Inputs: simplified magnitude, M_2 ; the infection rate, α ; relative frailty which is defined as $\phi \equiv M_2/\alpha$; and the distribution function for relative risk, R , and, in particular, its dispersion.
- Outputs: pre-pandemic LE for the whole population; mean years of life lost (YLL) by those who die from Covid; and adjusted post-pandemic life expectancy (APPLE) for the subset who survive the pandemic.

Table 1 summarises the impact on outputs of changes in input parameters in qualitative terms. In the next section, we present the results of a numerical analysis which quantifies the impacts in Table 1 under different scenarios.

Table1: Impact on ADM outputs of an increase in an ADM input parameter

Inputs	Outputs		
	LE	YLL	APPLE
Simplified magnitude, M_2	$\times$	≈ 0	+
Infection rate, α	$\times$	0	+
Relative frailty, $\phi \equiv M_2/\alpha$	$\times$	≈ 0	+
Dispersion of R	0	–	+

Note: $\times$: output variable does not depend on the input parameter; 0: output variable is not affected by changes in the input parameter; ≈ 0 : output variable is approximately unaffected by changes in the input parameter; +: output variable increases when an input parameter increases; and –: output variable decreases when an input parameter increases. In the case of the log-normal distribution for R , which is what we assume in Sections 4.3-4.5, the relevant dispersion measure is the standard deviation, σ , of $\log R$.

4. Numerical analysis of the Accelerated Deaths Model

We now illustrate the impact of the ADM on life expectancies and other quantities of interest using a range of scenarios. In doing so, we highlight how the results are sensitive to changes in key parameters in the model – building on Table 1. The section is structured around two values for the magnitude, $M_2 = 2$ and 0.4. The higher value of $M_2 = 2$ was chosen to make it easier to see in the following Figures the sensitivity of various model outputs to changes in the distribution of the relative risk, R (mortality heterogeneity) and α (proportion infected): it allows us to gain a good understanding of the role and importance of each parameter in the model. The smaller value of $M_2 = 0.4$ is based on the observed magnitude of the Covid pandemic in England. For this value, we focus on the impact of the pandemic on life expectancies. In combination, the two values of M_2 plus sensitivity of the various outputs to other parameters in the model give a good idea of how potential future pandemics (including ones that are more lethal than Covid) and in countries other than England might impact life expectancies.

As far as possible, we use realistic parameter values that are consistent with male and female mortality and related statistics in a variety of developed countries. We wish to focus attention on the model itself, rather than a specific application of the model to a particular pandemic, country and sex.

As discussed below, the degree of heterogeneity in mortality rates is currently not well understood. Nevertheless, we provide a lower bound for the degree of heterogeneity based on neighbourhood data for England; we also show that the degree of heterogeneity must have an upper limit. We use two values of the proportion infected, α : $\alpha = 0.8$ is based on English data, while $\alpha = 0.2$ is used as an alternative lower value in order to assess the sensitivity of the model to this parameter. Lastly, our aggregate force of mortality, $\mu(t, x)$, is chosen to be broadly appropriate for a range of developed countries and for both sexes without being specific to any.

4.1 Heterogeneity in mortality rates and the log-normal distribution for R

The distribution of $R > 0$ is a key determinant of difference between the pre-pandemic $LE(x)$ and the $APPLE(x)$. Of particular importance is the dispersion of the distribution of R .¹⁶ In England and many other countries, data are not publicly available at the individual level. However, for England, we do have data at the small neighbourhood level of LSOAs. We recognise that there will be further heterogeneity in mortality rates within each neighbourhood (i.e., at the individual level), and, so, the true level of heterogeneity in mortality might be noticeably higher than the between-neighbourhood

¹⁶ Note that R is the only stochastic variable in the model.

heterogeneity; the latter therefore provides a lower bound to the true heterogeneity in mortality in England.

Given the type of data available for England, we can decompose relative risk, R , as follows:

$$R \equiv R^{(1)}R^{(2)}, \quad (30)$$

where $R^{(1)}$ captures the between-neighbourhood (LSOA) heterogeneity, and $R^{(2)}$ captures the within-neighbourhood heterogeneity.

The distribution of $R^{(1)}$ has been estimated by Wen et al. (2023) and Cairns et al. (2024) for English males and females. Histograms derived from Cairns et al. (2024) are plotted in Figure 4 for the relative risks of English males aged 70-79. The blue histogram gives each LSOA equal weight. The pink histogram weights each LSOA by its average population for that age group over the period 2001 to 2018. Within this age group, LSOAs with a high relative risk tend to have smaller populations compared to LSOAs with low relative risk. Hence, when we weight by population size (pink bars), we see a shift to the left in the distribution. The skewed shape of the pink histogram is approximately what we would expect if the relative risk, $R^{(1)}$, has a log-normal distribution or other similar skewed distribution (where the underlying variable takes only positive values). We assume that the relative risk, $R^{(1)}$, has a mean of 1 (this follows from the fact that the aggregated LSOA data is the same as the total population of England). When we combine this with a specified Beard curve (see Equation (31) below) for the aggregate (national) force of mortality, the mean of 1 determines the level of the otherwise unobserved underlying baseline force of mortality, $\mu_B(t, x)$.¹⁷




Figure 4: Histogram of the relative risk, $R^{(1)}$, estimated for 32,844 LSOAs for English males in the 70-79 age group. Blue bars: each LSOA has equal weight. Pink bars: LSOAs are weighted by the size of the population in 70-79 age group.

The choice of distribution for $R^{(1)}$ is investigated further in Figure 5, where we compare the empirical cumulative distribution function (CDF) of the population-weighted LSOA observations (solid black

¹⁷ There is an identifiability issue which needs to be resolved. If we double all values of $R^{(1)}$ and halve the baseline force of mortality curve, $\mu_B(t, x)$, then the aggregate force of mortality curve, $\mu(t, x)$ is unchanged. In other words, there is an arbitrariness in the level of $\mu_B(t, x)$. By setting the mean of the relative risk, $R^{(1)}$, to 1, the identifiability problem is removed, and $\mu_B(t, x)$ can be estimated with precision.

curve) with the log-normal (blue curve) and an alternative candidate distribution, the translated Gamma¹⁸ (red curve) (see, for example, Klugman et al., 1998). It can be seen that the log-normal gives a reasonable approximation, while the translated Gamma (with one additional parameter) gives a slightly better fit especially in the left tail. However, the latter distribution has a certain disadvantage which we will highlight shortly.




Figure 5: Cumulative distribution functions for the relative risk by LSOA, $R^{(1)}$, for English males aged 70-79. Solid black curve: empirical CDF weighted by population size. Red and blue curves: fitted log-normal and translated Gamma distributions.

$R^{(2)}$ in Equation (30) captures the additional relative risk at the individual level within each LSOA. Unfortunately, we do not have any data for England that allows us to calibrate the distribution of $R^{(2)}$, so will need to make certain assumptions. One assumption we make is that $R^{(2)}$ is independent of $R^{(1)}$ and also has a mean of 1.¹⁹ Others will be discussed immediately below.

Now, if both $R^{(1)}$ and $R^{(2)}$ have a log-normal distribution, then so does R . (We will assume in Sections 4.3-4.5 below that the relevant dispersion measure for the distribution of R is the standard deviation, σ , of $\log R$.) But if $R^{(1)}$ has a translated Gamma distribution, there is no standard distribution for $R^{(2)}$ that gives a convenient form for R . Further, if $\log R^{(j)} \sim N(-\frac{1}{2}\sigma_j^2, \sigma_j^2)$ for $j = 1, 2$, then $\log R \sim N(-\frac{1}{2}(\sigma_1^2 + \sigma_2^2), (\sigma_1^2 + \sigma_2^2))$. Since this relies only on the sum $(\sigma_1^2 + \sigma_2^2)$ rather than the individual variances, we propose to use the log-normal for $R^{(1)}$, $R^{(2)}$ and R rather than a potentially more accurate distribution for $R^{(1)}$ (such as the translated Gamma). This means we only need to specify a single value for $\sigma^2 = (\sigma_1^2 + \sigma_2^2)$. From the distribution of R for English males in the 70-79 age group shown in Figures 3 and 4, $\sigma_1 \approx 0.3$, so the combined σ will be something in excess of 0.3. Various values of σ will be considered in different scenarios below, in light of the fact that,

¹⁸ If X has a Gamma distribution, then $Y = X + k$ for some constant k has a ‘translated’ Gamma distribution.

¹⁹ $R^{(2)}$ might have some dependence on $R^{(1)}$ but still have mean 1: i.e., independence might be dropped in favour of $R^{(1)}$ and $R^{(2)}$ being uncorrelated. For example, its standard deviation might depend on $R^{(1)}$. But, without individual level data, we are unable to pursue this further.

with no publicly available data on the distribution of $R^{(2)}$, we are unable to get a reliable estimate of σ_2^2 (for England).

The between-LSOA variance, σ_1^2 , will be largely explained by variation in the prevalence of underlying controllable risk factors (such as smoking), possibly augmented by some uncontrollable risk factors (e.g., genetic), where these vary significantly in terms of prevalence from neighbourhood to neighbourhood. The within-LSOA variance, σ_2^2 , will depend on:

- individual controllable risk factors, where these deviate from the LSOA average (e.g., an individual is or is not a smoker relative to their LSOA average smoking prevalence of, say, 20%);
- individual variation in preventable and non-preventable risk factors, where these vary from the LSOA average;
- random good or bad luck in terms of the onset of individual diseases or morbidities.

4.2 The aggregate national force of mortality

In our analysis, we choose to use the Beard curve for the aggregate force of mortality (see, e.g., Richards, 2012):

$$\mu(t, x) = \frac{ae^{b(x+t-x_0)}}{1 + \frac{a}{c}e^{b(x+t-x_0)}}. \quad (31)$$

For lower ages, this approximates to the Gompertz curve $ae^{b(x+t-x_0)}$, but as $x + t$ gets to high ages $\mu(t, x)$ plateaus at the constant c . The form of the curve is, therefore, consistent with what we typically observe at the national level in many countries, although there is some uncertainty over the level of the plateau and the nature of a transition from the Gompertz growth to the plateau. For further discussion of the plateau and its level, see Barbi et al. (2018). In terms of the results that follow, the Gompertz part of the curve (i.e., parameters a and b) is much more important than the plateau, so the precise model for $\mu(t, x)$ at very high ages is not critical to the outcome of our analysis of the impact of the pandemic.

The curve is parameterised so that the force of mortality at the base age x_0 is approximately a . We choose the following values for the pre-pandemic national mortality as being typical for mortality in a developed country: $x_0 = 60$, $a = 0.006$, $b = 0.09$, and $c = 0.65$ (see, e.g., Barbi et al., 2018). The Gompertz parameter b is composed of two parts: a period-mortality-table Gompertz rate of increase of 0.1 *minus* an annual improvement rate assumed to be 0.01. This results in a pre-pandemic remaining cohort life expectancy at age 60 of just over 27 years. To repeat, this parameterisation is not intended to be specific to any country or sex, but is appropriate for a variety of developed countries and for males and females.

4.3 Examples of the ADM I: A high-magnitude pandemic

We now consider the ADM under a variety of scenarios listed in Table 2. The first four scenarios use a high magnitude of $M_2(x) = M_2 = 2$ or 200%. A high value for M_2 acts as a ‘magnifying glass’ to show detail that is not easily visible at lower values that might be more appropriate for the Covid-19 pandemic, as experienced in England. Given $M_2 = 2$, the next key parameter is the level of mortality heterogeneity in relative risk, measured by the standard deviation, σ , of $\log R$. We experiment with $\sigma = 0.4$ (a low value just above the minimum value of 0.3 suggested in Section 4.1; Scenario A) and $\sigma = 0.8$ (a high value; Scenario B). We also experiment with two values of the proportion infected, $\alpha = 0.2$ (a low value in Scenarios A and B) and $\alpha = 0.8$ (a high value in Scenarios C and D, based on ONS estimates for England: see Massie, 2023). Relative frailty, ϕ , is derived from the values of M and α : i.e., $\phi \equiv M_2/\alpha$. The parameters in the last four columns are either standard Gompertz

parameters introduced in Section 4.2 (a and b) or alternative, post-pandemic values that capture a shift in the base table and a change in the mortality improvement rate (θ and β) (see Equations (31) and (35)).

Table 2: Parameter values used in Scenarios A to J

Scenario	M_2	σ	α	ϕ	a	b	θ	β
A	2	0.4	0.2	10	0.006	0.09	0	0
B	2	0.8	0.2	10	0.006	0.09	0	0
C	2	0.4	0.8	2.5	0.006	0.09	0	0
D	2	0.8	0.8	2.5	0.006	0.09	0	0
E	0.4	0.6	0.2	2	0.006	0.09	0	0
F	0.4	0.8	0.2	2	0.006	0.09	0	0
G	0.4	0.6	0.8	0.5	0.006	0.09	0	0
H	0.4	0.8	0.8	0.5	0.006	0.09	0	0
I	0.4	0.8	0.8	0.5	0.006	0.09	0	0.0025
J	0.4	0.8	0.8	0.5	0.006	0.09	0.05	0

Notes: M_2 is the ‘simplified magnitude’ of the pandemic (Equation (15)). σ is the level of mortality heterogeneity in relative risk, R . α is the infection rate. Relative frailty, ϕ , is derived from the values of M_2 and α : i.e., $\phi \equiv M_2/\alpha$. a and b represent the standard Gompertz parameters introduced in Equation (31); $c = 0.65$. θ and β are post-pandemic parameter values that capture a permanent proportional shift (θ) in the baseline force of mortality, and a reduction, β , in the baseline force-of-mortality improvement rate (Equation (35)).

In Figure 6, we plot (for $M_2 = 2$) the aggregate deaths curve for a cohort initially aged 75 pre-pandemic, $d(t, x)$ (dashed curve):

$$d(t, x) = -\partial S(t, x)/\partial t \quad (32)$$

and post-pandemic, $d_C(t, x)$ (solid curve):

$$d_C(t, x) = -\partial S_C(t, x)/\partial t. \quad (33)$$




Figure 6: Deaths curves under Scenarios A and B for a cohort initially aged 75 pre-Covid (dashed curve) and post-Covid pandemic survivors only (solid curve). Hatched region: accelerated deaths.

As in Figure 2, the hatched region between the two curves represents accelerated deaths due to Covid.²⁰ The two hatched regions look quite similar, but closer inspection reveals that the hatched region for Scenario B in the right plot is slightly fatter on the left (around age 75) and thinner on the right (e.g., ages 95 to 100) than Scenario A. This indicates that greater heterogeneity in the health of the underlying population (higher σ) leads to a greater left skewness in Covid-19 deaths, consistent with the Accelerated Deaths Hypothesis. This states that Covid-19 deaths are initially drawn from those who are less healthy and hence have higher relative risk as measured by their higher R value. The distribution of R measures the degree of mortality heterogeneity which increases with σ in the ADM. *Our analysis of the link between greater mortality heterogeneity and the pattern of accelerated deaths is one of the key contributions of the paper.*

In Figure 7, we plot the corresponding pre- and post-pandemic aggregate force-of-mortality curves, $\mu(t, x)$ (black curve) and $\mu_C(t, x)$ (red dashed curve) (see Equations (31) and (12), respectively) for the scenarios A and B. Even with a high value for M_2 , it is difficult to see the difference between the black and red curves, but it is clear that the difference is greater for Scenario B (right plot) where σ is higher.




Figure 7: Aggregate force of mortality curves for a cohort currently age 75 pre-Covid (black curve) and for survivors post-Covid (red dashed curve). Level of mortality heterogeneity in relative risk: $\sigma = 0.4$ (left panel) and $\sigma = 0.8$ (right panel).

However, the differences become visually clearer when we look at the proportion of accelerated deaths from Covid-19:

$$\pi_C(t, x) = 1 - \frac{d_C(t, x)}{d(t, x)}. \quad (34)$$

²⁰ Note that $d_C(t, x)$ (solid curve in Figure 6) is the deaths curve of those who have survived the pandemic and will include both people who have had Covid and those who have not. So, the solid deaths curve is for a smaller pool of survivors than the pre-pandemic pool (dashed curve). The hatched region shows those who experienced an accelerated death.

Equation (34) characterises the shape of the hatched region (accelerated deaths) in Figure 6 relative to the pre-Covid deaths curve. In Figure 8, each curve illustrates how $\pi_C(t, x)$ varies with the age of the cohort, $x + t$, at the time of death, with the different curves illustrating sensitivity to the initial cohort age, mortality heterogeneity, σ , and proportion infected, α , all with a fixed magnitude, $M_2 = 2$.




Figure 8: The proportion of accelerated deaths from Covid-19, $\pi_C(t, x)$, against age at death ($x + t$) for Scenarios A, B and C, where the magnitude $M_2 = 2$. The left panel shows the effect of different levels of mortality heterogeneity in relative risk, σ (0.4 vs 0.8, with $\alpha = 0.2$). The right panel shows the effect of different values of the proportion infected, α (0.2 vs 0.8, with $\sigma = 0.4$). Each plot has three pairs of lines corresponding to three cohorts starting at time $t = 0$ at different initial ages $x = 65, 75$ and 85 .

The first observation is that $\pi_C(t, x)$ is greater the higher the initial age, x . This simply reflects the fact that, as age increases, a greater proportion of those still alive at age x will die from Covid.

Other features of Figure 8 need more interpretation. The curves are calculated empirically (and so do not have an explicit functional form), but they all have a similar shape to the right section of an inverted logistic or sigmoid curve (which both have an inverted S shape) which converges on the right to a *strictly positive value above zero*. This positive value differs from scenario to scenario as well as by age. The left panel considers the sensitivity of $\pi_C(t, x)$ to changes in σ by comparing Scenarios A and B, while the right panel considers the sensitivity of $\pi_C(t, x)$ to changes in α by comparing Scenarios A and Scenario C: in each case, the higher the value of the parameter, the steeper the curve declines towards a positive limit on the right. This type of shape and positive right limit contrasts with the earlier proposal in Cairns et al. (2020) of an exponential curve that decays to zero. The latter appeared to be a reasonable (as well as pragmatic) choice made in the early days of the Covid-19

pandemic – without the benefit of the subsequent evidence that accumulated during the course of the pandemic. By contrast, the inverted S-shape curves in Figure 8 are the result of a more rigorous analysis that combines the Proportionality Hypothesis with our recent increased understanding of mortality heterogeneity.

The reason for convergence of $\pi_C(t, x)$ on the right of Figure 8 to a positive limit is connected to the shape of the left tail of the distribution of relative risk, R (see Figure 4). Convergence of $\pi_C(t, x)$ to zero on the right would require the density function of R to be strictly positive at $R = 0$ (contrasting with the log-normal distribution where the density converges to zero at $R = 0$). But this would imply that each cohort would have some individuals who were in exceptional health relative to the average for their current age. As an example, if the cohort had a chronological age of 75, these exceptionally healthy individuals would have a biological age of, say, 20 to 30. But this level of health for a 75-year-old seems quite implausible and is certainly not backed up by the LSOA-level data (Figure 4). On the other hand, describing a 75-year-old as ‘looking’ 20 years younger is certainly plausible. This would be consistent with $\sigma = 0.8$. With this value of σ , about 0.6% of 75 year olds would have a biological age of 55 or less (and, therefore, might ‘look’ 20 years younger). With similar reasoning, for a 75-year-old to look 50 years younger with the same probability, we would need to have $\sigma = 2$. Since this is highly unlikely, it suggests that σ must be well below 2.

For Scenarios A and B, Figure 9 shows (remaining) life expectancies for the three populations discussed earlier:

- life expectancies for the whole of the population pre-pandemic, i.e., in the absence of Covid-19;
- mean years of life lost (YLL) for those who die from Covid-19; and
- adjusted post-pandemic life expectancies (APPLEs) for those who survive Covid-19 or are never infected.

Equation (28) shows the relationship between the three life expectancies.




Figure 9: Mean remaining life expectancies (pre-pandemic LEs (solid black curve) and APPLEs (dashed green curve)) and mean years of life lost (YLL) by those who die from Covid (red curve) in Scenarios A and B, where magnitude $M_2 = 2$ (high) and proportion infected $\alpha = 0.2$. Level of mortality heterogeneity in the relative risk: $\sigma = 0.4$ (left panel) and $\sigma = 0.8$ (right panel).

Remaining life expectancies are plotted in Figure 9 for fixed values of $M_2 = 2$ and $\alpha = 0.2$, and low and high values of σ (0.4 and 0.8: Scenarios A and B). In both cases, it is hard to see the difference between the pre-pandemic LE (black curve) and the APPLE (green dashed curve). In contrast, the YLL (red curve) for those who die from Covid is noticeably lower than the pre-pandemic LE. Additionally, a comparison of the left and right panels for Scenarios A and B, respectively, shows that YLL (red line) has a considerable dependence on σ . For example, at age 60, mean YLL is around 25 when $\sigma = 0.4$ and 22 when $\sigma = 0.8$. As we proved theoretically in Section 3.5.2, a higher σ (implying greater heterogeneity in individual states of health within an age cohort) is associated with a larger number of people in the poorest health, and hence fewer average years of life lost. Again, this is consistent with the Accelerated Deaths Hypothesis and Covid's 'detrimental selection' of those in the poorest health.

Nevertheless, even when σ takes a high value of 0.8, the YLL for age 60 is still 22 years. Although this is well below the population-wide pre-pandemic LE of 27 years, it clearly helps to rebut suggestions made in the early stages of the pandemic that those who died from Covid-19 were already at 'death's door' (to paraphrase comments quoted in Knapton, 2020). Countering these media comments, Hanlon et al. (2020) estimated that the mean years of life lost (averaging over all ages) was about 9 to 12 years.²¹ Hanlon et al.'s estimates are consistent with the values displayed in Figure 9 (averaging across the red curves), even though we use a different methodology and data set. We also observe from Figure 9 that YLL is highly sensitive to changes in σ . By contrast, YLL is insensitive to changes in the magnitude, M_2 , and the proportion infected, α – again as proved in Section 3.5.2.

The relationship between the APPLE and the pre-pandemic LE is illustrated more clearly in Figure 10. This Figure considers the same two scenarios, A and B, illustrated in Figure 9 plus a third scenario, D, with a higher infection rate, α , but the same magnitude and high level of mortality heterogeneity ($M_2 = 2$, $\sigma = 0.8$, and $\alpha = 0.8$). At age 60, there is an increase of only 0.1% to 0.25% (depending on the scenario) from the pre-pandemic LE to the APPLE (Equation (27)), reflecting the fact that, even in these high-magnitude scenarios ($M_2 = 2$), only a very small proportion of people at age 60 die from Covid-19. The plot also shows that the percentage increase rises steadily with the initial age. At higher ages, a greater proportion of the cohort will die from Covid-19 (in line with the Proportionality Hypothesis), removing a greater proportion of the less healthy.

Figure 10 also explores the sensitivity of the percentage increase to changes in σ and α (while adjusting ϕ to keep $M_2 = 2$ constant). We can see that the pattern of sensitivity to these parameters is different. Doubling σ more than doubles the percentage increase at all ages. But a fourfold increase in α from 0.2 to 0.8, accompanied by a corresponding reduction in ϕ from 10 to 2.5 (so that M_2 stays constant) has only a minor effect at age 60, although a much bigger effect at higher ages. The relatively minor impact of a change in α is connected to the approximation in Equation (9): if $M_2 = \alpha\phi$ stays constant but the balance between α and ϕ changes, then there is very little impact on the $S_C(R, 0^+, x)$ post-Covid survival curves. Overall, even with a high-magnitude pandemic, Figure 10 shows that the impact on the APPLE is relatively low, below 1% at younger ages and still below 3% at age 85. To repeat, the increase in APPLE, in response to an increase in σ , is the counterpart to the reduction in YLL (see Equation (25) or (26)). This again illustrates the two dimensions of detrimental selection.

²¹ Hanlon et al. (2020) also find that the YLL vary considerably in a way that is consistent with what we would expect: by age (YLL falls as age increases); and by the number of long-term adverse conditions (YLL falls as the number of conditions rises).




Figure 10: Impact of detrimental selection on the mean life expectancy of survivors – percentage change by initial age from the pre-pandemic remaining life expectancy (LE) to the APPLE for Scenarios A, B and D (Equation (27)). The three curves show the sensitivity to changes in mortality heterogeneity in relative risk, σ , and proportion infected, α , of the percentage change in LE in the three scenarios, when the magnitude of the pandemic $M_2 = 2$ (high).

4.4 Examples of the ADM II: A low-magnitude pandemic consistent with the Covid-19 pandemic in England

We now apply the ADM to the Covid-19 pandemic over the period 2020-2024 and introduce some (more realistic) scenarios that are consistent with the observed scale of Covid deaths in England. Specifically, we now reduce the magnitude to $M_2 = 0.4$ or 40%, significantly lower than the high value of $M_2 = 2$ in the previous section. The lower magnitude is consistent with total deaths involving Covid for ages 60 and above for England over the full course of the pandemic as a proportion of total annual deaths in 2019 (see, e.g., Office for National Statistics, 2023a,b, and Worldometers, 2024).^{22,23} The lower value for M_2 is intended to make the ADM more realistic in terms of parameterisation, but this section should still be regarded as illustrative of the model in general, rather being a focus on English pandemic mortality in particular.

In Figure 11, we show the percentage increase in the APPLE by age with different values of α (0.2 and 0.8) and σ (0.6 and 0.8), when the magnitude $M_2 = 0.4$ (Scenarios E, F, G and H). With the lower magnitude, results are less sensitive than the scenarios illustrated in Figure 10 to the choice of α , but just as sensitive to the choice of σ as before. In more general terms, we note that the percentage increase in the APPLE is very low in Figure 11, implying that the detrimental selection effect on LEs

²² Strictly, these data are for England and Wales, but, since the population of Wales is only 5% of the total, the use of $M_2 = 0.4$ applies equally to England as to the slightly larger population of England and Wales.

²³ In the UK, there are around 600,000 deaths annually across all ages in a normal non-Covid year, so, $M_2 = 2$ would imply around 1.2 million deaths from Covid (over a period of 2 years). As of April 2024, there had been around 230,000 Covid deaths across the whole UK since the start of the pandemic, implying a magnitude of around $M_2 = 0.4$ for the pandemic in England.

(that the survivors are healthier and live longer than the pre-Covid population) is very small in aggregate (see Equation (27)).²⁴




Figure 11: Impact of detrimental selection on the mean life expectancy of survivors – percentage change by initial age from the pre-pandemic life expectancy (LE) to the APPLE for Scenarios E, F, G and H (Equation (27)). The four curves show the sensitivity to changes in mortality heterogeneity in relative risk, σ , and proportion infected, α , of the percentage change in LE in the four scenarios, when the magnitude of the pandemic $M_2 = 0.4$ (low, consistent with experience in England).

Comparison of Figures 10 and 11 (e.g., when $\sigma = 0.8$ and $\alpha = 0.8$, scenarios D and H) allows us to assess the impact of reducing M_2 from 2 to 0.4. The difference is a factor of 5 and the difference between the LE and APPLE is now approximately a factor of 5 lower (see Equation (25)). Overall, Figure 11 shows that the impact on APPLE is very low, below 1% across all scenarios when $M_2 = 0.4$.

4.5 Examples of the ADM III: Changing the base mortality table and the post-Covid improvement rate

In the calculations above, it is assumed that baseline all-cause mortality and the rate of improvement in mortality would be unaffected by the pandemic: the pandemic happens (at time $t = 0$) and then mortality rates by subgroup revert to their trajectories as projected before the pandemic. We refer to the case $M_2 = 0.4$, $\sigma = 0.8$, and $\alpha = 0.8$ as Scenario H and now introduce two further scenarios that reflect changes in circumstances and practices post-pandemic:

²⁴ As Equation (27) shows, detrimental selection effect is small because a small $M_2(x)$ is scaled by typically an even smaller $\mu_B(0, x)$.

- Scenario I: no change in the post-pandemic base mortality table, but the assumed post-pandemic mortality improvement rate is reduced by 0.25% per annum ($\beta = 0.0025$ in Table 2).²⁵
- Scenario J: the base mortality table post pandemic is 5% higher than previously assumed ($\theta = 0.05$ in Table 2),²⁶ but there is no change in the assumed mortality improvement rate compared to pre-pandemic assumptions. In combination, this means the 5% shift persists in projected forces of mortality. This might be due to a deferral of mortality improvements²⁷ that would otherwise have happened during the pandemic, or it might be linked to persistent additional deaths due to ‘long-Covid’ (see, e.g., NHS, 2025). See Chemaitelly et al. (2023) for a discussion.

The two adjustments are applied directly to the subgroup forces of mortality, $\mu(R, t, x)$. Specifically,

$$\tilde{\mu}(R, t, x) = \mu(R, t, x)\exp(\theta + \beta t). \quad (35)$$

Scenarios I and J both incorporate the pandemic and the resulting detrimental selection analysed above. Additionally, we assume that the distribution of the relative risk, R , at each initial age, x , at time $t = 0$ is the same in all three scenarios. The impact on the APPLE under Scenarios H, I and J is illustrated in Figure 12. Both I and J result in higher future forces of mortality than in scenario H, and so the APPLES are lower. We can also see that, at ages 60 to 70, scenarios I and J have a much bigger impact on APPLES. For I relative to H, the gap gradually closes as the change in the improvement rate has less time to take effect before the elderly cohort dies. For J relative to H, we see a slowly widening gap (with APPLE decreasing) at higher initial ages.

4.6 Sensitivity to the distribution function for R

For completeness, we considered if our results are sensitive to the use of the log-normal distribution for R . Appendix B.2 conducts the sensitivity analysis: we compare the log-normal, Gamma and translated Gamma distributions. Our main conclusion is that the results do have a small dependence on the choice of distribution. Nevertheless, the differences are quite small provided we calibrate the distributions on the basis of Equation (22) rather than on the basis of matching variances.

²⁵ A reduction of 0.25% in the improvement rate is consistent with the range of small adjustments in the valuation assumption setting that have been observed in practice in the UK actuarial profession (see, e.g., Isio, 2023). It can be argued, however, that this adjustment is linked less to the Covid-19 pandemic than to lower economic growth rates in the UK (i.e., the continuation of so-called ‘austerity’ in public health spending which began after the Global Financial Crisis in 2007-09 and which started to lower life expectancy trend improvements in 2011).

²⁶ A 5% increase in the base table is consistent with what has been observed in 2023 (when Covid deaths were at a very low level compared to 2020) relative to what was predicted at the end of 2019 for 2023. See also Continuous Mortality Investigation, (2024).

²⁷ For example, efforts to develop vaccines and treatments for Covid-19 might have involved the redirection of resources which delayed ongoing medical research and the ensuing mortality improvements by a few years.




Figure 12: Impact of detrimental selection on the mean life expectancy of survivors – percentage change by initial age from the pre-pandemic life expectancy (LE) to the APPLE for Scenarios H, I, and J, which all assume $M = 0.4$, $\sigma = 0.8$, and $\alpha = 0.8$ (Equation (27)). Scenario H: Covid only. Scenario I: Covid plus a post-pandemic reduction in the assumed mortality improvement rate. Scenario J: Covid plus a post-pandemic parallel upward shift in the base mortality table.

4.7 Summary

Summarising the analysis above:

- Key model inputs are: the simplified magnitude of the pandemic, M_2 (although the true value of this parameter will only be known once the pandemic is over); the infection rate, α ; and mortality heterogeneity, measured by the distribution function for relative risk, R , and, in particular, its dispersion (σ).
- Relative frailty, ϕ , is another key input, but is not observable and has to be inferred from two variables which are observable, namely, M_2 and α , giving $\phi \equiv M_2/\alpha$. Despite being an important factor in the development of a pandemic, its value can only be determined accurately after the pandemic is over and the true value of M_2 is known.
- Key model outputs are: the pre-pandemic life expectancy (LE), the mean years of life lost (YLL), and the adjusted post-pandemic life expectancy (APPLE).
- YLL and APPLE are particularly sensitive to changes in σ , the former negatively and the latter positively. Doubling σ more than doubles the differences between (a) the pre-pandemic LE and the YLL (due entirely to the reduction in YLL, one dimension of detrimental selection), and (b) the APPLE and the pre-pandemic LE (being the complement of $LE - YLL$).
- YLL is not sensitive to changes in M_2 , α or ϕ . We show that the years of life lost by an individual dying of Covid depends only on their individual health status and not on the scale

of the pandemic or the infection rate or their relative frailty²⁸ (for reasons explained in Section 3.5.2).

- The percentage increase in the APPLE is approximately proportional to the percentage increase M_2 , α or ϕ . However, once these input parameters are scaled by $\mu_B(0, x)$, the overall impact on APPLE is low (Equation (27)).
- To illustrate this point in the case of England, the change from the pre-pandemic LE to the APPLE was found to be very small (less than 1%), implying that the impact of detrimental selection on the average life expectancy of survivors was very small in aggregate (Figure 11).
- We found a much larger impact on the APPLE, especially at lower ages, if the base table for mortality or the future mortality improvement rate changed from their pre-pandemic values either as a direct or indirect result of the pandemic.
- Greater mortality heterogeneity in the health of the underlying population (a higher value for σ in the case of the log-normal distribution for R) leads to a greater left skewness in Covid-19 deaths, consistent with the Accelerated Deaths Hypothesis and, in turn, reflecting the strength of the detrimental selection effect. However, we found that empirically, the profile of Covid deaths had an inverted S shape which converges on the right to a strictly positive value above zero (Figure 8); this contrasts with the proposal in Cairns et al. (2020) of an exponential curve that decays to zero.
- On the basis of evidence from England, the higher value for σ in the case of the log-normal distribution for R must lie above 0.3 (representing observable heterogeneity between subgroups). There is no clear upper bound to σ , but on the basis that a 75-year-old cannot reasonably have a biological age of 25, σ cannot be larger than 2. It should probably be significantly less than 2 but we cannot be specific about this upper boundary in the way than we can be specific about the lower bound of 0.3.

5. Implications for annuity providers, pension schemes and life insurers

In this section, we briefly consider some implications of our analysis for annuity providers, pension schemes and life insurers.

5.1 Implications I: The price of life annuities

Annuity prices (and the value of pension scheme liabilities) following the Covid pandemic will depend on the relationship between the pre- and post-pandemic $LE(x)$ and $APPLE(x)$. A life expectancy measure can be considered to be the fair value of a life annuity value from the same age with 0% interest and no expenses. The pandemic resulted in the surviving population being, on average, healthier and longer lived than the average before the pandemic and that, as a result, annuity prices might be expected to rise. However, from the above analysis, it is clear that, in the case of England, the impact of detrimental selection on the fair value of annuities for positive rates of interest would be very small, implying negligible changes to annuity prices. This echoes the conclusions of

²⁸ Recall, that the Subgroup Proportionality Hypothesis states that relative frailty (ϕ) is not dependent on the subgroup or individual (i.e., the R_i).

Cairns et al. (2020).²⁹ Those conclusions, together with the numerous online presentations³⁰ given by its authors to global audiences of actuaries working on pension risk transfers, helped to maintain the bulk annuity and buyout markets during the early stages of the pandemic (Blake and Cairns, 2021) and our new results here reinforce this confidence in the face of the worst global pandemic since Spanish Flu in 1918-19.

5.2 Implications II: Base mortality table and the mortality improvement rate

The results in Figure 12 also implicitly reflect the impact of post-pandemic changes in mortality assumptions on annuity prices by age. A typical annuity portfolio of lives and amounts will be weighted towards the 60 to 70 age group. Figure 12 suggests that annuity providers need to pay considerable attention to both the base mortality table (and how this differs from the 2019-based projection) and the future mortality improvement rate post-pandemic. They need to answer the following questions: ‘do either need adjustment relative to pre-pandemic assumptions?’; and ‘can each change be properly justified and then be quantified with a high degree of confidence?’

5.3 Implications III: The impacts on death rates

The various impacts of the force of mortality can also be investigated from a life insurance perspective. Like annuities, the short-term impacts of the force of mortality (similar to Figure 7, but with $M_2 = 0.4$) are very small: for pandemic survivors, equivalent to less than a one-month reduction in biological age and then reverting to zero. However, there might be a bigger impact on life insurance pricing and reserving if there is a significant change in the base table or improvement rates (with the latter being less important for life insurers). In particular, the impact would be greatest for term assurance business as a result of a change in the base table.

5.4 Implications IV: The assessment of future pandemics

As alluded to earlier, the ADM is just as relevant in insurers’ assessments of future pandemic risk. A key point underpinning the ADM is the Proportionality Hypothesis (which requires relative frailty (ϕ) to be constant in Equation (8)). The Proportionality Hypothesis is likely to apply when a new virus or agent is of a novel variety where individuals do not harbour protective antibodies. If everyone is equally exposed, then it is their underlying biological frailty that governs their ability to fight off the novel virus. However, we need to exercise some caution, as each pandemic will be different and have its own complexities. An example is Spanish Flu where older people carried protective antibodies (so it was not, therefore, a ‘novel’ virus) from a previous flu pandemic. Consequently, younger people who had no prior exposure were much more badly affected.

6. Further potential applications of the ADM: Seasonal flu and climate change

The ADM has the potential to be used to investigate the impact of ‘regular’ seasonality, such as that associated with seasonal flu, and climate-related extreme events.

²⁹ However, some of the conjectures in Cairns et al. (2020) are not consistent with the findings in the current study. For example, the increase in the APPLE in Cairns et al. (2020) was greater than reported here because the hatched region in Figure 6 was assumed to be much more left skewed, implying much greater mortality heterogeneity than turned out to be the case. This is despite the fact that the 2020 paper made the assumption that there would be 80-120,000 deaths, whereas the actual number of deaths was higher at 230,000.

³⁰ The first was on 15 May 2020. See https://www.youtube.com/watch?v=m_muHnwch-g&feature=youtu.be or <https://www.prudential.com/risk-transfer/impact-of-Covid-19>.




Figure 13: Weekly death registrations in England per 100,000 by deprivation level and sex over the period 2005 to 2018 for the age group 75 to 84. Source: Cairns et al. (2020).

As a first example, the ADM might be used to model the short-term impact of seasonal flu outbreaks (mostly outbreaks with a magnitude M_2 that is somewhat smaller than the values used in Section 4). However, we need to be more cautious here because of the potential impact of existing antibodies and annual flu vaccination campaigns, again adding to the complexity of the analysis.

Nevertheless, Figure 13, showing death registrations from all causes per 100,000 population in England,³¹ allows us to assess the applicability of the Proportionality Hypothesis between subgroups with reference to two subgroups (IMD deciles 1 to 5, and IMD deciles 6 to 10).³² Weekly rates are plotted on a log scale. With a log scale, the gap between successive peaks and troughs seems to be

³¹ Registration data from Office for National Statistics (2019).

³² The Index of Multiple Deprivation (IMD) measures relative deprivation (ranked from most to least deprived) in English Lower Layer Super Output Areas (UK Government, 2019).

constant regardless of whether excess winter mortality is high (e.g., 2008-09) or low (e.g., 2005-06). This observation is consistent with the Subgroup Proportionality Hypothesis:

$$m_w(i, t, w, x) = m(i, t, x)s(t, w, x) \quad (36)$$

where: $m_w(i, t, w, x)$ is the week w death rate within year t , subgroup i , age x ; $m(i, t, x)$ is the annual, all-cause death rate; and the seasonality component $s(t, w, x)$ is not specific to the subgroup. This seasonal component would incorporate both general cold weather seasonality as well as the highly variable influenza outbreaks.

Another area that merits further investigation is modelling the impact of excess deaths during a period of extreme heat, with events of this type becoming more common due to climate change. One hypothesis is that the probability of death caused by extreme heat is proportional to an individual's biological frailty (and hence consistent with the Proportionality Hypothesis). Other factors might need to be taken into account, such as geographical variations in extreme temperatures and an individual's ability to take protective action against extreme heat (e.g., staying in an air-conditioned building). Under this hypothesis, the mean years of life lost would be similar to those calculated in Figure 9, as the YLL is not overly sensitive to the duration or magnitude of the underlying event.

Guibert et al. (2025) model relative risk using French data as a continuous function of temperature and by age group. At the national level, and using all-cause daily death counts, their results suggest that there are only small differences between age groups. This is, therefore, consistent with the Aggregate Proportionality Hypothesis (i.e., by age) and, accordingly, with the use of the ADM. However, on the basis of current evidence, it falls short of proving that the Subgroup Proportionality Hypothesis and ADM are appropriate in this context, since Guibert et al. (2025) did not (or were not able to) incorporate heterogeneity within an age group. Their study contrasts with Qiao et al. (2015) whose analysis of heat-related deaths is built on the concept of 'mortality displacement' (also known as 'harvesting') which generates a stronger detrimental selection effect than we find here; they project a much lower YLL than we do. So more work needs to be done.

7. Conclusions

In this study, we have developed the Accelerated Deaths Model using a novel synthesis of two distinct recent lines of research into mortality modelling. The ADM builds on the Proportionality Hypothesis of Cairns et al (2025) which demonstrates that Covid-19 death rates are proportional to biological frailty, defined as all-cause death rates in a normal non-Covid year. It also builds on recent developments in our understanding of the heterogeneity in mortality in different age cohorts, allowing us to specify a more accurate deaths curve for each cohort. In particular, the profile of Covid deaths appears to have an inverted S shape which converges on the right to a strictly positive value above zero; this contrasts with the proposal in Cairns et al. (2020) of an exponential curve that decays to zero.

The ADM has a small number of key parameters and just a single stochastic variable, i.e., relative risk. The ADM places strong constraints on the size of and the relationship between certain parameters, e.g., the level of mortality heterogeneity or the link between infection rates, pandemic magnitude and relative frailty. In turn, this places strong constraints on the range of feasible output variables, such as mean years of life lost (YLL) by those who die from Covid and the adjusted post-pandemic life expectancy (APPLE).

Our main conclusion is that the impact on the surviving population in England of detrimental selection (i.e., the extent to which people who died from Covid-19 were less healthy compared with

their age cohort as a whole) was very small and much smaller than the impact of a general change in future mortality assumptions post-pandemic (i.e., the base table and improvement rate). As a result, the impact on annuity prices of an increase in APPLE was also very small.

The ADM model developed here can be applied to other pandemics, where the degree of detrimental selection might be different. It also has potential applications in the cases of seasonal flu and climate-related extreme events.

We end by proposing some suggestions for further work:

- Refine the ADM to allow for variations in the relative frailty (and hence the infection fatality rate) over time (due to vaccinations, etc.) and due to multiple infections.
- Improve data collection on infection rates by, e.g., age, sex and deprivation level.
- Measuring mortality heterogeneity (or relative risk) more accurately at the individual level. Heterogeneity between neighbourhoods gives us a lower bound for the standard deviation of the relative risk distribution, but within-neighbourhood heterogeneity might be just as significant, if not more so. Analysis of individual level mortality is now possible in some countries at the national level (e.g., Denmark, see Savcicens et al., 2024, and Hansen et al., 2025) or where there is partial coverage of the national population in other countries (e.g., England, see Bhaskaran et al., 2022). But more needs to be done to turn the output from these studies into an explicit distribution for the relative risk.
- Consideration of changes to the future mortality improvement rate following the pandemic. How much of an adjustment is required? Should such an adjustment be pandemic related, or should it be linked to bigger-picture economic changes that are occurring at the same time?
- Developing an early warning system for predicting the magnitude of a pandemic. A key issue in the current version of the ADM is that we need to specify the magnitude of the pandemic, M_2 . However, this is not known for certain until the pandemic has transitioned to the endemic phase which might be around two years after it started. It would be valuable if the magnitude of a pandemic could be predicted much earlier, taking account of the early data collected on the trajectory of infection and infection fatality rates by age and sex.

References

- Banerjee, A., Pasea, L., Harris, S., Gonzalez-Izquierdo, A., Torralbo, A., Shallcross, L., Noursadeghi, M., Pillay, D., Sebire, N., Holmes, C., Pagel, C., Wong, W.K., Langenberg, C., Williams, B., Denaxas, S., and Hemingway, H. (2020) Estimating excess 1-year mortality associated with the Covid-19 pandemic according to underlying conditions and age: A population-based cohort study. *Lancet*, 395:1715-1725. DOI: doi.org/10.1016/S0140-6736(20)30854-0.
- Bhaskaran, K., Rentsch, C. T., Hickman, G, Hulme, W. J., Schultze, A., Curtis, H.J, Wing, K., Warren-Gash, C., Tomlinson, L., Bates, C. J., Mathur, R., MacKenna, B., Mahalingasivam, V., Wong, A., Walker, A. J., Morton, C. E., Grint, D., Mehrkar, A., Eggo, R. M., Inglesby, P., Douglas, I. J., McDonald, H. I., Cockburn, J., Williamson, E. J., Evans, D., Parry, J., Hester, F., Harper, S., Evans, S. J. W., Bacon, S., Smeeth, L., and Goldacre, B. (2022) Overall and cause-specific hospitalisation and death after Covid-19 hospitalisation in England: A cohort study using linked primary care, secondary care, and death registration data in the OpenSAFELY platform. *PLoS Med* 19(1): e1003871. DOI: doi.org/10.1371/journal.pmed.1003871

Barbi, E., Lagona, F., Marsili, M., Vaupel, J.W., and Wachter, K.W. (2018) The plateau of human mortality: Demography of longevity pioneers. *Science*, 360(6396):1459-1461.

Blake, D., and Cairns, A.J.G. (2021) Longevity risk and capital markets: The 2019-20 update. *Insurance: Mathematics and Economics*, 99: 395-439.

Cairns, A.J.G., Blake, D., Kessler, A., and Kessler, M. (2020) The impact of Covid-19 on future higher-age mortality. Working paper. DOI: https://www.researchgate.net/publication/342079076_The_Impact_of_Covid-19_on_Future_Higher-Age_Mortality (A version of this paper was subsequently published as: Cairns, A.J.G., Blake, D., Kessler, A., and Kessler, M. (2025) The impact of Covid-19 on higher-age mortality. *European Society of Medicine: Medical Research Archives*, 13(1). DOI: doi.org/10.18103/mra.v13i1.6186)

Cairns, A.J.G., Blake, D., Kessler, A., Kessler, M., and Mathur, R. (2025) Covid-19 mortality: The Proportionality Hypothesis. *European Actuarial Journal*, 15: 509-554. DOI: doi.org/10.1007/s13385-024-00400-9

Cairns, A.J.G., and Redondo Lourés, C. (2023) Higher-age US mortality by education and cause of death: Trends, inequality and controllable risk factors. In *2023 Living to 100 Compendium*, Society of Actuaries. DOI: <https://www.soa.org/resources/essays-monographs/2023-living-to-100-compendium/>

Cairns, A.J.G., Wen, J., and Kleinow, T. (2024) Drivers of mortality: Risk factors and inequality. *Journal of the Royal Statistical Society, Series A*, 187: 989-1012. DOI: doi.org/10.1093/jrssa/qnae017

Carannante, M., D'Amato, V., and Iaccarino, G. (2022) The future evolution of the mortality acceleration due to the Covid-19: The Charlson Comorbidity Index in stochastic setting. *Frontiers in Cardiovascular Medicine*, 9. DOI: [10.3389/fcvm.2022.938086](https://doi.org/10.3389/fcvm.2022.938086)

Chemaitelly, H., Faust, J.S., Krumholz, H.M., Ayoub, H.H., Tang, P., Coyle, P., Yassine, H.M., Al Thani, A.A., Al-Khatib, H.A., Hasan, M.R., Al-Kanaani, Z., Al-Kuwari, E., Jeremijenko, A., Kaleeckal, A.H., Latif, A.N., Shaik, R.M., Abdul-Rahim, H.F., Nasrallah, G.K., Al-Kuwari, M.G., Butt, A.A., Al-Romaihi, H.E., Al-Thani, M.H., Al-Khal, A., Bertollini, R., and Abu-Raddad, L.J. (2023) Short- and longer-term all-cause mortality among SARS-CoV-2- infected individuals and the pull-forward phenomenon in Qatar: A national cohort study. *International Journal of Infectious Diseases*, 136(Nov): 81-90. DOI: [10.1016/j.ijid.2023.09.005](https://doi.org/10.1016/j.ijid.2023.09.005).

Chen, Z., Li, H., Mao, Y., and Zhou, K. Q. (2025) Learning from Covid-19: A catastrophe mortality bond solution in the post-pandemic era. *Insurance: Mathematics and Economics*, 123: 103-113. DOI: doi.org/10.1016/j.insmatheco.2025.103113

Continuous Mortality Investigation (2024) England and Wales mortality monitor – end of 2023. Institute and Faculty of Actuaries. DOI: <https://www.actuaries.org.uk/documents/england-wales-mortality-monitor-Covid-19-update-q4-2023> (Accessed 21/8/2024.)

Guibert, Q., Pincemin, G., and Planchet, F. (2025) Impacts of climate change on mortality: An extrapolation of temperature effects based on time series data in France. *International Journal of Forecasting* (in press). DOI: doi.org/10.1016/j.ijforecast.2025.07.004

Hanlon, P., Chadwick, F., Shah, A., Wood, R., Minton, J., McCartney, G., Fischbacher, C., Mair, F.S., Husmeier, D., Matthiopoulos, J., and McAllister, D.A., (2020) Covid-19 – exploring the implications of long-term condition type and extent of multimorbidity on years of life lost: A modelling study. *Wellcome Open Research*, 5:75. DOI: <https://wellcomeopenresearch.org/articles/5-75>

Hansen, N.U., Ergemen, Y.E., and Kallestrup-Lamb, M. (2025) Individual health indices via register-based health records and machine learning. *European Actuarial Journal*, 15: 607–632.

Isio (2023) Mortality update. DOI: <https://www.isio.com/app/uploads/2024/01/7.-Isio-mortality-update-August-2023-2.pdf> (Accessed 21/8/2024.)

Klugman, S., Panjer, H.H., and Willmot, G. (1998) *Loss Models: From Data to Decisions*. New York: Wiley.

Knapton, S. (2020) Two thirds of coronavirus victims may have died this year anyway, government adviser says. *Daily Telegraph*, 25 March. DOI: <https://www.telegraph.co.uk/news/2020/03/25/two-thirds-patients-die-coronavirus-would-have-died-year-anyway> (Accessed 25/10/2024.)

Massie, L. (2023) How many people have had Covid-19? The challenge of reinfections. *National Statistical*, 15 February. DOI: <https://blog.ons.gov.uk/2023/02/15/how-many-people-have-had-Covid-19-the-challenge-of-reinfections/> (Accessed 6/11/2025.)

Milevsky, M.A. (2020) Calibrating Gompertz in reverse: What is your longevity-risk-adjusted global age? *Insurance: Mathematics and Economics*, 92: 147-161.

NHS (2025) Long Covid. DOI: <https://www.nhs.uk/conditions/long-Covid/> (Accessed 6/11/2025.)

Office for National Statistics (2019) Weekly deaths registrations by IMD, sex and age group, England and Wales, 2005 to 2018. DOI: <https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/deaths/adhocs/10929weeklydeathsregistrationsbyimdsexandagegroupenglandandwales2005to2018> (Accessed 6/11/2025.)

Office for National Statistics (2023a) Deaths registered in England and Wales – 21st century mortality. DOI: <https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/deaths/datasets/the21stcenturymortalityfilesdeathsdataset> (Accessed 6/11/2025.)

Office for National Statistics (2023b) Single year of age and average age of death of people whose death was due to or involved coronavirus Covid-19). DOI: <https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/deaths/datasets/singleyearofageandaverageageofdeathofpeoplewhosedeathwasduetoorinvolvedCovid19> (Accessed 6/11/2025.)

Qiao, Z., Guo, Y., Yu, W., and Tong, S. (2015) Assessment of short- and long-term mortality displacement in heat-related deaths in Brisbane, Australia, 1996-2004. *Environmental Health Perspectives*, 123:766-72.

Rashid, T., Bennett, J. E., Paciorek, C. J., Doyle, Y., Pearson-Stuttard, J., Flaxman, S., Fecht, D., Toledano, M. B., Li, G., Daby, H. I., Johnson, E., Davies, B., and Ezzati, M. (2021) Life expectancy and risk of death in 6791 communities in England from 2002 to 2019: High-resolution spatiotemporal analysis of civil registration data. *Lancet Public Health*. 6(11):e805–e816. DOI: [doi.org/10.1016/S2468-2667\(21\)00205-X](https://doi.org/10.1016/S2468-2667(21)00205-X)

Richards, S.J. (2012) A handbook of parametric survival models for actuarial use. *Scandinavian Actuarial Journal*, 2012: 233-257.

Savcicens, G., Eliassi-Rad, T., Hansen, L.K., Mortensen, L.H., Lilleholt, L., Rogers, A., Zettler, I., and Lehmann, S. (2024) Using sequences of life-events to predict human lives. *Nature Computer Science*, 4: 43–56.

Shkolnikov, V.M., Andreev, E.M., Jdanov, D.A., Jasilionis, D., Kravdal, Ø., Vågerö, D., and Valkonen, T. (2012) Increasing absolute mortality disparities by education in Finland, Norway and Sweden, 1971-2000. *Journal of Epidemiology and Community Health*, 66: 372-378.

UK Government (2019) English indices of deprivation 2019. DOI: <https://www.gov.uk/government/statistics/english-indices-of-deprivation-2019>

Wen, J., Cairns, A.J.G., and Kleinow, T. (2023) Modelling socio-economic mortality at neighbourhood level. *ASTIN Bulletin*, 53: 285-310.

Worldometers (2024) United Kingdom. DOI: <https://www.worldometers.info/coronavirus/country/uk/> (Accessed 6/11/2025.)

Appendix A: Notation

Pre-Covid functions	Definition
$m(i, x)$	<i>Annual</i> all-cause death rate in a non-Covid year at age x in subgroup i (deaths between exact ages x and $x + 1$ in a given calendar year)
$m_C(i, x)$	<i>Annualised</i> death rate from Covid (deaths between exact ages x and $x + 1$ over a defined period and then annualised)
$\mu_B(t, x)$	Baseline instantaneous force of mortality at time t for an individual aged x at time 0.
$\mu(t, x)$	Aggregate force of mortality across all subgroups
$\mu(R, t, x) = R\mu_B(t, x)$	Force of mortality for individuals in a subgroup with relative risk R
$S_B(t, x)$	Baseline survivor curve corresponding to $\mu_B(t, x)$
$S(R, t, x) = S_B(t, x)^R$	Survivor curve for individuals in a subgroup with relative risk R
$S(t, x)$	Aggregate survivor curve taking account of the distribution function for R
$d(t, x)$	Aggregate curve of deaths corresponding to $S(t, x)$
$LE(R, x)$	Pre-Covid life expectancy of an individual aged x in a subgroup with relative risk R
$LE(x)$	Aggregate pre-Covid life expectancy
Covid-related functions	
$\alpha(i, x)$	Infection rate at age x in subgroup i : mostly assumed to be a constant, α , in this study
ϕ	Relative frailty: assumed to be constant in this study
$IFR(R, x) = \phi R\mu_B(0, x)$	Infection fatality rate for individuals initially aged x in a subgroup with relative risk R
$1 - \exp(-IFR(R, x))$	Probability of death from Covid given an individual has been infected
$S_C(R, t, x)$	Survivor curve with Covid present for individuals initially aged x in a subgroup with relative risk R
$S_C(R, 0^+, x)$	As immediately above, but just after deducting Covid deaths at time $t = 0$
$S_C(t, x)$	Aggregate survivor curve with Covid present, taking account of the distribution function for R
$d_C(t, x)$	Aggregate deaths curve for Covid survivors only
$\pi_C(t, x) = 1 - d_C(t, x)/d(t, x)$	Proportion of Covid deaths that are accelerated
$\mu_C(t, x)$	Force of mortality for Covid survivors corresponding to $S_C(t, x)$
$YLL(x)$	Mean years of life lost per person aged x for those who die from Covid
$APPLE(x)$	Mean adjusted post-pandemic life expectancy per person aged x for those who survive the pandemic
$M_1(x)$	Pandemic magnitude at age x : aggregate probability of death from Covid relative to the pre-Covid force of mortality, $\mu(x)$
$M_2(x)$	Simplified magnitude, defined as $\alpha\phi$ and approximately equal to $M_1(x)$

Appendix B: Technical issues

B.1 Useful approximations

In this appendix, we develop some useful approximations linking the relationships between YLL, APPLE and LE. In the equations below, the weights, w_i , represents the conditional probability that a given person who dies from Covid-19 belongs to subgroup i .

B.1.1 The relationship between YLL and LE

Consider first the relationship between YLL and LE:

$$YLL(x) = \sum_{i=1}^N w_i LE(R_i, x) \quad (B1)$$

where

$$\begin{aligned} w_i &= \frac{1 - S_C(R_i, 0^+, x)}{\sum_{j=1}^N (1 - S_C(R_j, 0^+, x))} \\ &\approx \frac{\alpha \phi R_i \mu_B(0, x)}{\sum_{j=1}^N \alpha \phi R_j \mu_B(0, x)} \quad (\text{from Equation (9)}) \\ &= \frac{R_i}{\sum_{j=1}^N R_j} \\ &\approx N^{-1} R_i \quad (\text{since } N^{-1} \sum_j R_j \approx 1). \end{aligned} \quad (B2)$$

Hence,

$$YLL(x) \approx \frac{1}{N} \sum_{i=1}^N R_i LE(R_i, x). \quad (B3)$$

We also then have

$$LE(x) - YLL(x) \approx \frac{1}{N} \sum_{i=1}^N (1 - R_i) LE(R_i, x) \quad (B4)$$

and the impact of detrimental selection on YLL is given by:

$$100 \left(1 - \frac{YLL(x)}{LE(x)} \right). \quad (B5)$$

We can develop (B4) through two further approximations that help considerably with our interpretation of the results. First, note that $LE(R_i, x)$ can be approximated as a linear function of $\log R_i$:

$$LE(R_i, x) = \lambda_0(x) - \lambda_1(x) \log R_i. \quad (B6)$$

Figure B1 (top right panel) shows this linearity graphically. The logic behind this approximation is as follows:

- At very low ages, and assuming that the baseline force of mortality follows the Gompertz model, $\mu_B(t, x) = a \exp(b(x + t - x_0))$, we have, for $R = 1$, $LE(R, x - 1) \approx LE(R, x) + 1$.
- If R is different from 1, then the force of mortality becomes

$$\mu(R, t, x) = a R \exp(b(x + t - x_0)) = a \exp\left(b\left(x + t + \frac{\log R}{b} - x_0\right)\right). \quad (B7)$$

In (B7), $(\log R)/b$ is equivalent to the difference between biological and chronological age in the Gompertz model, which is a good approximation in our model to the baseline force of mortality. So, at these young ages, $LE(R, x) \approx LE(1, x) - (\log R)/b$.

- At higher ages $LE(R, x)$ begins to flatten out (although it has to stay positive), but empirically the linear approximation in $\log R$ remains a good one. For example, in Figure B1 (top right

panel) the slope at about -8.5 is flatter than the slope of $-1/b = -1/0.09 = -11.1$, reflecting this gradual flattening out at higher ages.

- From Figure B1 (top right panel), it is clear that this slope, $-\lambda_1(x)$, is almost, but not exactly, independent of the distribution for R .




Figure B.1: Comparison of various intermediate and final outputs under different distributional assumptions (Scenarios G, G1, G2 and G3) for R . Top left: the baseline force of mortality, $\mu_B(t, x)$, is derived from the aggregate $\mu(t, x)$ and the chosen distribution function for R . Top right: Scatterplot showing the relationship between simulated values of R and $LE(R, x)$ for $x = 60$ under the four scenarios. Bottom left: aggregate $LE(x)$ (dashed line; same under all four scenarios) and the $YLL(x)$ for the four scenarios. Bottom right: the impact of detrimental selection on years of life lost (Equation (B5)).

Using Equation (B6), we then find that

$$LE(x) - YLL(x) \approx N^{-1} \lambda_1(x) \sum_{i=1}^N (R_i - 1) \log(R_i). \quad (B8)$$

If the R_i are all close to 1 (i.e., R has a low variance around the mean of 1), then we can use $\log(R_i) \approx R_i - 1$ to derive our second approximation:

$$LE(x) - YLL(x) \approx N^{-1} \lambda_1(x) \sum_{i=1}^N (R_i - 1)^2 = \lambda_1(x) \text{Var}(R). \quad (B9)$$

B.1.2 The relationship between APPLE and LE

Next, consider the relationship between APPLE and LE:

$$APPLE(x) = \frac{\sum_{i=1}^N S_C(R_i, 0^+, x) LE(R_i, x)}{\sum_{i=1}^N S_C(R_i, 0^+, x)} \quad (B10)$$

$$\approx \frac{N^{-1} \sum_{i=1}^N (1 - \alpha \phi R_i \mu_B(0, x)) LE(R_i, x)}{N^{-1} \sum_{i=1}^N (1 - \alpha \phi R_i \mu_B(0, x))} \quad (B11)$$

$$\approx \frac{LE(x) - \alpha \phi \mu_B(0, x) YLL(x)}{1 - \alpha \phi \mu_B(0, x)} \quad (B12)$$

$$= \frac{LE(x) - M_2(x) \mu_B(0, x) YLL(x)}{1 - M_2(x) \mu_B(0, x)} \quad (\text{since } M_2(x) \equiv \alpha \phi) \quad (B13)$$

$$\approx LE(x) + M_2(x) \mu_B(0, x) (LE(x) - YLL(x)) \quad (B14)$$

$$\Rightarrow APPLE(x) - LE(x) \approx M_2(x) \mu_B(0, x) (LE(x) - YLL(x)). \quad (B15)$$

It follows from (B14) that the impact of detrimental selection on the life expectancy of survivors is given by:

$$100 \left(\frac{APPLE(x)}{LE(x)} - 1 \right) \approx 100 \frac{M_2(x) \mu_B(0, x) (LE(x) - YLL(x))}{LE(x)}. \quad (B16)$$

B.1.3 The relationship between YLL, APPLE and LE

Finally, note that as $\text{Var}(R) \rightarrow 0$, $APPLE(x)$ and $YLL(x)$ both converge to $LE(x)$. To demonstrate this, if $\text{Var}(R) = 0$, then $R_i = 1$ ($\forall i$) and, using (B3) and (B14), it follows that YLL, APPLE and LE are all equal to each other:

$$YLL(x) = \frac{1}{N} \sum_{i=1}^N LE(1, x) = LE(x) = APPLE(x). \quad (B17)$$

B.2 Sensitivity of the results to the distribution function for R

The main results in this paper are illustrated on the basis of the assumption that R has a log-normal distribution. But it is of interest to know how the results might change if the distribution of R changes. To test this, we considered three alternatives to Scenario G (see Table 2 in the main text):

- Scenario G: $R \sim \text{Log-Normal}(-0.018, 0.6^2)$
- Scenario G1: $R \sim \text{Gamma}(2.31, 2.31)$
- Scenario G2: $R \sim \text{Gamma}(2.76, 2.76)$
- Scenario G3: $R - 0.2 \sim \text{Gamma}(1.58, 1.975)$.

The parameters of G1 are chosen so that R has the same mean and variance as scenario G (based on Equation (B9)). Scenarios G2 and G3 are calibrated so that we match $N^{-1} \sum_{i=1}^N (R_i - 1) \log(R_i)$ (based on Equation (B8)). The impact of these changes from the log-normal are illustrated in Figure B1.

In the top left panel, we see the impact of the change in the baseline force of mortality (recall that the aggregate force of mortality is fixed and equal to the Beard curve (Equation (31))). All four curves are very close up to about age 85 before drifting apart slowly, with Scenario G1 deviating the most.

In the lower left panel, we plot the pre-pandemic mean LE by age (dashed curve), which is the same for all four scenarios, against $YLL(x)$. All four YLL curves are quite close together, but we can see that the red curve for scenario G1 is a bit different from the other three. This is the first indication that

approximation (B8) is better than using the variance of R to calibrate models to the original log-normal model for the distribution of R . The relationships between scenario G1 and the other three scenarios are clearer in the lower right panel, where we plot $100(1 - YLL(x)/LE(x))$ (Equation (B5)). Here we can see that scenarios G, G2 and G3 are generally much closer to each other than G1, but nevertheless drift apart slowly after age 70.

The top right panel plots $LE(R_i, 60)$ against $\log(R_i)$ for the four scenarios. If all four scenarios had exactly the same baseline force of mortality, the $LE(R_i, 60)$ curves would be identical. Hence, the differences that we see are due to differences in the baseline force of mortality that we observe in the top left panel. More importantly, the approximate linearity in the top right panel that forms the basis of approximation (B8) is clearly observable.

B.3: Aggregate life expectancy

The ADM begins with a given aggregate force of mortality curve, $\mu(t, x)$. The model has two expressions for pre-Covid aggregate life expectancy: $LE_1(x) = \int_0^\infty S(t, x) dt$, where $S(t, x) = \exp\left[-\int_0^t \mu(s, x) ds\right]$ is the aggregate survival function; and $LE_2(x) = \int_0^\infty f(r)LE(R, x)dr$, where $f(r)$ is the density function of the random variable R (see Equation (4)). Note, we are utilising the whole assumed distribution for R below, rather than N draws from this distribution, $R_1, \dots, R_N$, as is used in Section 3.5.1; this means that we will need to use modified versions of some of the equations in the main text.

We can show that these two expressions $LE_1(x)$ and $LE_2(x)$ are equal:

$$\begin{aligned}
LE_2(x) &= \int_0^\infty f(r)LE(R, x)dr \text{ (a modified form of Equation (17))} \\
&= \int_0^\infty f(r) \int_0^\infty S_B(t, x)^r dt dr \\
&= \int_0^\infty \int_0^\infty f(r)S_B(t, x)^r dr dt \\
&= \int_0^\infty E[S_B(t, x)^R] dt \\
&= \int_0^\infty S(t, x) dt \text{ (using Equation (4))} \\
&= LE_1(x).
\end{aligned} \tag{B18}$$

This proves that the distributional effects of R in the calculation of pre-pandemic life expectancy, $LE(x)$, cancel out, implying that $LE(x)$ does not depend on R .