

Covid-19 Mortality: The Proportionality Hypothesis

Andrew J.G. Cairns¹

David Blake²

Amy Kessler³

Marsha Kessler⁴

Rohit Mathur⁵

21 August 2024

Abstract

We introduce and provide evidence to support the Proportionality Hypothesis which states that Covid-19 infection fatality rates are approximately proportional to all-cause death rates by age and subgroup (e.g., socio-economic class). We also show that vaccination played a very significant role in preventing people infected with Covid-19 from needing to be hospitalised, since it reduced the average severity of an infection. Death rates involving Covid-19 were very significantly lower for people in the fully vaccinated group compared to the unvaccinated group. During the pandemic, death rates from other causes were in some cases reduced (e.g., flu and pneumonia), in some cases unchanged (e.g., lung cancer) and in some cases elevated (e.g., heart disease). We discuss the implications of our findings both for potential adjustments to extrapolative mortality models which allow for future pandemics in a way that is consistent with the Proportionality Hypothesis and for insurance companies in terms of both modelling extreme scenarios and the design of mortality catastrophe bonds.

Keywords

Covid-19 mortality; the Proportionality Hypothesis; infection fatality rate; all-cause mortality; biological frailty; relative frailty; age; deprivation

Acknowledgements

The authors thank numerous webinar audiences for their invaluable questions and feedback during the development of this paper. We are particularly grateful for the very constructive comments of the journal's assigned editor and two anonymous referees.

¹ Heriot-Watt University and the Maxwell Institute for Mathematical Sciences, A.J.G.Cairns@hw.ac.uk (Corresponding author)

² Bayes Business School, City St George's, University of London, D.Blake@city.ac.uk (Corresponding author)

³ Prismic Life, Amy.Kessler@prismiclif.com

⁴ Commonwealth Care Alliance, mkessler@commonwealthcare.org

⁵ Prudential Financial Inc., Rohit.Mathur@prudential.com

1. Introduction

We conduct a detailed examination of the development of Covid-19 mortality in England over the period from the beginning of 2020 to Spring 2023. Our analysis allows us to propose the Proportionality Hypothesis which states that Covid-19 infection fatality rates are approximately proportional to all-cause death rates (1) by age (the ‘aggregate’ version) and (2) by age and subgroup, e.g., socio-economic class (the ‘subgroup’ version). The Covid-19 ‘infection fatality rate’ (IFR) is defined as the proportion of newly infected individuals who subsequently die from Covid-19.¹ We show that the hypothesis holds up well in the phase of the pandemic prior to the rollout of the vaccination programme (denoted Phase 1 below) and in the phase after the vast majority of the population had either had the full programme of vaccines or had Covid-19 (denoted Phase 3).

1.1 The three phases of the pandemic

The pandemic in England fell into three phases:

- Phase 1: the year 2020
- Phase 2: from the beginning of 2021 to mid-2022
- Phase 3: from mid-2022 onwards.

Phase 1 covers the initial period of the pandemic when there was a relatively high degree of infection synchronisation (as shown in Figure 2 below) across regions and socio-economic subgroups, and, more importantly, vaccination rollouts had not yet commenced. Phase 2 is the transition period involving the first vaccine rollout, the second and subsequent booster vaccines, together with the sequence of new variants of the virus that brought new waves of infection. Phase 3 (which, at the time of writing, was still ongoing) marks the transition to an endemic state for the virus, with a steady flow of new infections and relatively low numbers of deaths, similar to influenza (flu). The relatively low number of deaths in Phase 3 is also partly explained by the fact that a high percentage of the population now had antibodies, as a result of previous infection or vaccination.

1.2 Data sources and definitions

We utilise a range of datasets that are listed in Appendix A. Much of the data come from the UK’s Office for National Statistics (ONS)² with some additional data from the Human Mortality Database and the UK’s National Health Service (NHS). Wherever possible, we try to use data in a consistent way through time. However, there are inconsistencies between different datasets. For example, infection rates, hospitalisations and death rates are published using different age groupings across the different datasets.

¹ IFR is the ratio of the number of individuals who die from Covid-19 to the number of individuals who have been infected with Covid-19. In the paper, we also use the term Covid-19 ‘death rate’. This is defined as the proportion of the total population who die of Covid-19 during a specified time period. The Covid-19 death rate is equal to the infection rate multiplied by the IFR (by age and subgroup) – see Equation (3) below. In our primary statement of the Proportionality Hypothesis – given in Section 3.2 below – we decided to focus on the IFR. Our reason for this is that the Covid-19 death rate depends on infection rates as well as the IFR. But infection rates vary considerably through time and they can also vary considerably between age groups and socio-economic subgroups at any given point in time. So a statement of the Proportionality Hypothesis framed around the Covid-19 death rate – which we also present below – is more intricate than the statement expressed in terms of the IFR.

² Most of which are freely available on the ONS website.

There are also different definitions. We use several distinct measures of Covid-19 deaths that depend on data availability for specific tasks in the paper. For most purposes in this paper, we count a death as a Covid-19 death when Covid-19 is *mentioned* on the death certificate. We also, in places, make use of death counts where Covid-19 is listed as the *underlying or main cause* of death. Alongside these differences, death counts can be based on *registered* deaths during a specified period, or on actual *occurrences*. While death registrations might not be perfect, they have the advantage of being consistent over time and accurate (for the given definition), apart from the first few weeks of the pandemic when there was some misreporting of cause of death. Mostly, we also use *week of registration*, but we use *week of occurrence* as an alternative in one place in the paper.

1.3 The concepts of biological and relative frailty

Frailty is an important concept in this paper, but it can have different meanings to different professional groups. In an actuarial and demographic context, actuarial frailty relates to all the factors (some of which are unobserved) that affect human mortality other than age. It generally involves an adjustment up or down to an individual's or subgroup's death rate relative to a standard rate at the same age and sex (see Vaupel et al., 1979, and discussions in Carannante et al., 2023a). In a medical context, the concept of frailty is quite different: it is an absolute measure of poor health, rather than a measure that is relative to an age-dependent benchmark. As such, it represents an accumulation of visible comorbidities, disabilities, and chronic and acute illnesses, combined with the invisible but gradual deterioration of cells and organs.^{3,4}

In this paper, we employ two distinct concepts of frailty and, since these differ in definition from actuarial frailty, we give each a precise name. The first concept is 'biological frailty'. This is measured in absolute terms (and therefore is closer to the medical concept of frailty) and is defined as the one-year death rate from all causes (in a 'normal' non-Covid year), excluding deaths from external causes (e.g., road and other accidents, suicide, and accidental poisoning).⁵ Biological frailty depends on both the impact of age (i.e., the baseline mortality curve) and the variation within an age group (e.g., due to variation in death rates at the individual or subgroup level relative to the average for that age group). We use the term 'biological frailty' to match the sister concept of 'biological age'; where the latter can differ from 'chronological age' (see, e.g., Huang et al, 2017, and Milevsky, 2020).⁶ In our study, biological frailty is based on all-cause mortality rates in a typical pre-Covid year and by socio-economic subgroup.

The second concept is Covid-specific 'relative frailty'. This measures the risk of death from Covid-19 by age and subgroup (if data for the latter are available) for a person newly infected with Covid-19 *relative* to the risk of death from all causes in a normal non-Covid year.

³ See, for example, <https://www.uptodate.com/contents/frailty>.

⁴ For other interpretations of frailty and an analysis of the factors which influence frailty at both the national and subgroup level, see, e.g., Carannante et al., 2023b, Savcisens et al., 2023, Wen et al, 2021, Cairns et al., 2024a, and Sherris and Wei, 2021 (and the references within).

⁵ In subsequent mentions of the all-cause death rate, it should be assumed that this always excludes external causes.

⁶ Chronological age measures how many calendar years an individual has lived. However, two individuals of the same chronological age might have different prospective death rates (due to differences in the cumulative damage to their bodies, disability, ill-health, behaviours, etc.). In each case, we measure their current prospective death rate and match this rate to a specific point on the national mortality table. The age in the national table that matches the individual's prospective death rate is referred to as their biological age. Biological age can therefore be higher or lower than chronological age.

1.4. Key findings

The key findings are as follows. The main theme running through the paper is the Proportionality Hypothesis, which focuses on the Covid-19 infection fatality rate (IFR). There are two versions depending on the granularity of the data published during the pandemic. The ‘subgroup’ Proportionality Hypothesis states that the IFR is approximately proportional to all-cause death rates (i.e., biological frailty) by both age and subgroup. The ‘aggregate’ Proportionality Hypothesis focuses on age only and states that the IFR is approximately proportional to all-cause death rates by age. As a consequence, and contrary to much media commentary at the time, death rates from Covid-19 *did not* disproportionately affect the more deprived groups in society. Where we did observe differences in mortality, these were due both to *pre-existing* health inequalities at the all-cause level and to differences in infection rates between sub-groups.

We find strong evidence showing the benefits of vaccination. First, infection rates fell significantly following vaccination of older groups in advance of younger age groups. Second, for those vaccinated who did become infected with Covid-19, vaccination led, on average, to less severe cases and to fewer deaths (i.e., a lower infection fatality rate). By the end of 2022, the infection fatality rate was 1/20th to 1/30th of its value at the end of 2020: the result of the combination of vaccination and/or prior infection alongside improved treatments.

The impact of Covid-19 on other causes of death was quite varied. For some causes, particularly cancers, there was no observable impact *in the short term*. For others, death rates went up (e.g., heart disease). And, in a few cases, most notably influenza and pneumonia, death rates dropped substantially, probably as a result of lockdowns and social distancing.

Lastly, we discuss some takeaways for insurers. First, insurers should revisit their extreme pandemic mortality scenarios. In particular, models should allow for significant random variation between different geographical areas, age groups and socio-economic subgroups. Second, we conclude that mortality catastrophe bonds might not be as effective as a hedging instrument as had previously been believed, unless they incorporate a linkage to sub-population mortality indices.

1.5 Plan for the rest of the paper

The rest of this paper is structured as follows. In Section 2, we review the headline data for England, covering the different pandemic waves and regional variations, and the emergence of Covid-19 variants. Section 3 discusses Phase 1, covering the year 2020, in detail, introduces the Proportionality Hypothesis and then assesses it by examining infection and death rates by deprivation decile and region. Section 4 analyses Phase 2, covering the period between the beginning of 2021 and mid-2022, in detail and assesses how well the Proportionality Hypothesis holds up. Section 5 considers Phase 3, the transition to the endemic state of the virus which began in mid-2022 and also reassesses the hypothesis. The impact of Covid-19 on other causes of death is discussed in Section 6, while Section 7 shows how extrapolative mortality models can be adjusted to allow for future pandemics in a way that is consistent with the Proportionality Hypothesis. Section 8 considers the implications for insurers and Section 9 concludes. The data sources are listed in Appendix A. Appendix B summarises the findings of the paper.

2. What do the data reveal? The headlines

2.1 Pandemic waves and regional variations

In this section, we review the key features of the pandemic between the beginning of 2020 and the Spring of 2023, covering Phases 1-3.

We look, first, at weekly mortality for England and Wales using data from the Short-Term Mortality Fluctuations database available from the Human Mortality Database (see Appendix A).⁷ Figure 1 shows, for the period since 2014, weekly death rates for males and females in the 75-84 age group, the group that was most badly affected by Covid-19. The pre-pandemic years 2014-19 are represented by grey lines, 2020 in black, 2021 in red and 2022 in green. The grey lines give a good indication of the seasonality in mortality in England and Wales in pre-pandemic ‘normal’ years. These typically peak in January and bottom out in the late summer. Winter rates can be seen to be much more variable than summer rates, due, for example, to the severity of seasonal flu. The year 2020 (black line) started off as a relatively benign year, but we can then see the very significant first wave of excess deaths around weeks 15 to 20 which was much more severe than any of the recent bad flu years. A second wave of excess deaths appeared towards the end of 2020 and peaked in January 2021 (red line). Spring 2021 was characterised by a period of lower-than-normal mortality. This, perhaps, reflected the possibility that many who were already in very poor health and likely to die around this time had died in the previous months with or due to Covid-19.⁸ After that, the weekly mortality curves for males and females fluctuated above and below the pre-pandemic curves as new waves of Covid-19 came and went, partly in response to the emergence of new variants of the virus. But Covid-19-related deaths from mid-2021 onwards can be seen to be relatively modest compared to the first wave in 2021.

⁷ This database combines England and Wales, so does not allow the data for England to be separately identified. Although Wales is more rural than England, any bias will be quite small, since the population of England is around 18 times larger than Wales (56.5 million vs 3.1 million in 2021; <https://www.ons.gov.uk/peoplepopulationandcommunity/populationandmigration/populationestimates>; accessed 6/8/23). The remainder of this study uses data from England only.

⁸ We denote this phenomenon ‘accelerated deaths’ due to Covid-19.

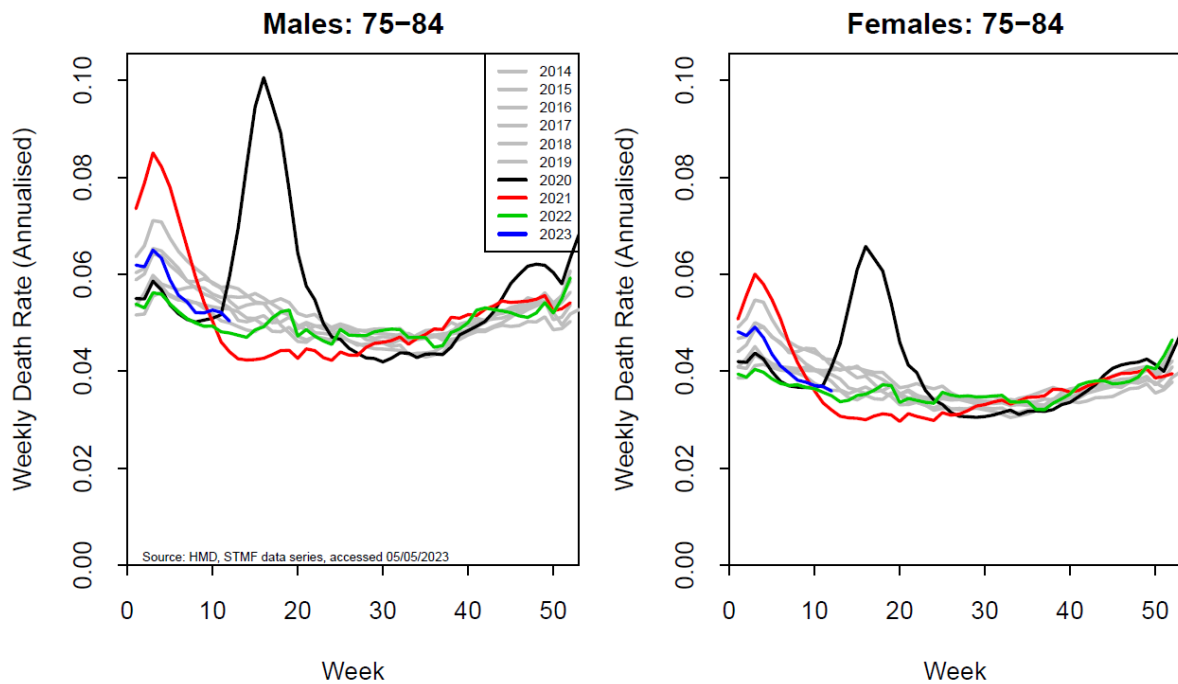


Figure 1: Weekly death rates in England and Wales for males and females aged 75 to 84 during a given year. Grey lines show rates for 2014 to 2019. Coloured lines show the impact of the pandemic in 2020, 2021, 2022 and Spring 2023.

In Figure 2, we use weekly deaths data provided by the ONS which allows analysis across a number of characteristics, including region. In particular, we look at Covid-19-related deaths (where Covid-19 is *mentioned* on the death certificate and so might not be the underlying cause of death) as a proportion of the 5-year average of all deaths over the five years 2015-2019 for the same week of the year (males and females combined, all causes, all ages).

This figure allows us to observe how much variation there was between different parts of England and how this changed during the course of the pandemic. In the upper plot, the first wave of the pandemic was closely synchronised across the nine official regions of England, with the exception that London was a week or two ahead. But, in spite of this synchronisation (and the fact that the first national lockdown was introduced simultaneously across the whole country on 26 March 2020), the magnitude of the first wave varied considerably between regions. London and other regions with dense urban areas experienced significantly higher Covid-19 death rates. Summer 2020 saw very low death rates from Covid-19 (and also very low infection rates).

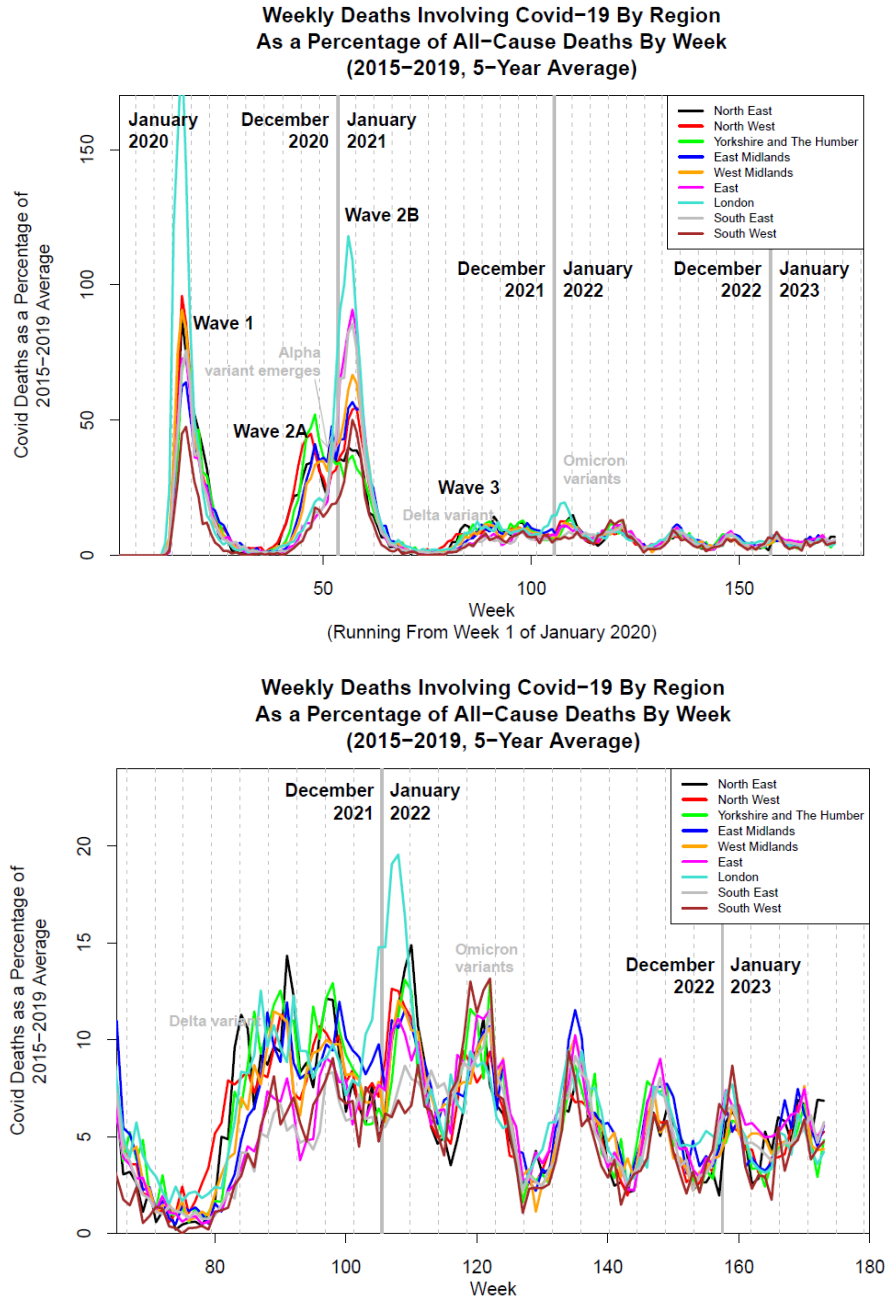


Figure 2: Upper: weekly Covid-19 deaths (males and females combined) by region as a percentage of all-cause deaths (average by week over the period 2015-2019) from week 1 of 2020 through to Spring 2023. Covid-19 deaths count deaths where Covid-19 was mentioned on the death certificate. Lower: zoomed in to allow analysis of the smaller individual waves from Spring 2021 to Spring 2023.

The second wave consisted of two parts.⁹ The first part (labelled Wave 2A) occurred in the Autumn of 2020 and was largely focused in the northern regions (accompanied by some local lockdowns). But in December 2020, a new variant of the virus, Alpha, emerged, generating much higher peak mortality (with a second national lockdown) in January 2021 (labelled Wave 2B) – with much higher death rates in London and lower death rates in the more-rural South West. This was followed by a period in early summer with very low death rates.

⁹ Arguably there were two separate waves, but we have chosen to classify them as two parts of the same wave.

For the remainder of 2021 and through the first half of 2022 (Figure 2, lower plot), there was a long and irregular third wave which involved the Delta and Omicron variants. Covid-19 death rates were much lower than for Waves 1 and 2.¹⁰ Individual regions had distinct patterns, and January 2022 witnessed the last significant peak, with London death rates clearly higher than other regions.

For the rest of 2022 and into the Spring 2023, the pandemic moved into the endemic phase with relatively regular small waves (but with declining amplitude) approximately every three months. Further, the waves reached their troughs at higher levels (indicating more Covid-19 deaths) than those of the Summers of 2020 and 2021.¹¹ In addition, the waves became much more synchronised, with relatively little variation by region. This synchronisation suggests that, as people's lives returned to normal, new Covid variants could circulate between regions much more quickly, contrasting with Autumn 2020 when travel restrictions were much more severe and this prevented Wave 2A from spreading from the northern regions to the south.

Other cause-of-death data from the ONS (but not publicly available) provides information for the whole of 2020 at the sub-regional level of health boards (Clinical Commissioning Groups; CCGs¹²). Over the data period, there were 106 CCGs across England with an average of around 500,000 people in each. The data allow us to compare deaths from Covid-19 (as the *underlying* cause of death) for ages 40-89 during 2020 as a percentage of deaths from all causes in 2019. Results are plotted in Figure 3. As with Figure 2, the chosen metric implicitly adjusts for socio-economic variation between CCGs. Each CCG is represented by a single dot and shows Covid-19 deaths as a percentage of 2019 deaths for males (x-axis) versus females (y-axis). Dots are coloured according to the urban-rural mix of the CCG: bright green for the most rural CCGs and bright red for the most urban. As might be expected, there is a strong correlation between male and female death rates from Covid-19. But there are other important findings. First, there is much more variation at the sub-regional CCG level than at the regional level (compare with Figure 2). Second, Figure 3 reveals the strong association between Covid-19 death rates and the urban-rural mix. The implication of this, again as might be expected, is that Covid-19 spreads more easily in more dense urban areas with consequently higher death rates. This has implications for the development of future (extreme) pandemic scenarios in an insurance context and indicates the importance of knowing a victim's geographic location, as discussed further in Section 8.

¹⁰ The Spanish flu pandemic of 1918–19 also had three waves with the second wave having the highest peak; <https://www.britannica.com/event/influenza-pandemic-of-1918-1919>

¹¹ For example, the trough in week 140 is at a higher level than that in week 75 (lower panel of Figure 2).

¹² Clinical Commissioning Groups were created following the Health and Social Care Act in 2012, and replaced Primary Care Trusts on 1 April 2013. They were clinically-led statutory NHS bodies responsible for the planning and commissioning of health care services for their local area. As of 1 April 2021, following a series of mergers, there were 106 CCGs in England. They were dissolved in July 2022 and their duties taken over by the new Integrated Care Systems (ICSs); <https://www.nhsconfed.org/articles/what-are-clinical-commissioning-groups> (accessed 6/8/23).

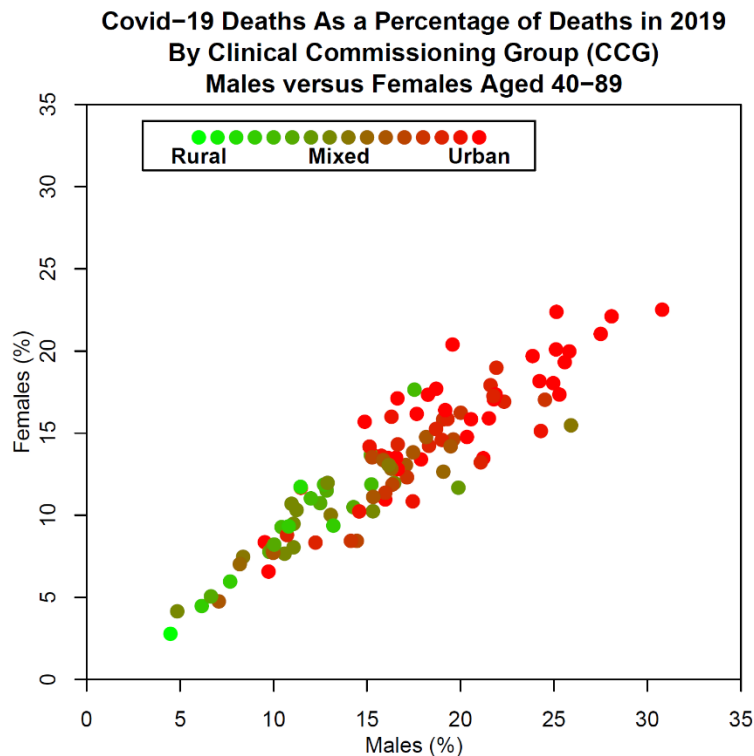


Figure 3: Number of deaths registered in 2020 by Clinical Commissioning Group (CCG) where Covid-19 is recorded as the underlying cause of death as a percentage of deaths due to all causes in 2019, both over ages 40 to 89. Each dot represents one CCG and shows the percentage of Covid-19 deaths for males versus females.

2.2 The emergence of variants

Figure 4 shows the dominant variants of Covid-19 over time from the end of 2020. It can be seen that each new variant (e.g., the Delta variant from late Spring 2021) typically takes hold quite rapidly, displacing the previous variant in one to two months. There are a number of possible explanations: higher reproductive numbers for new variants (also known as the R number: the average number of new infections caused by each existing infected person); a gradually falling R number for the previous variant, as the proportion of the population still susceptible falls and the population moves towards herd immunity against that variant; and whether or not prior infections give strong or only limited protection against new variants.

The emergence of a new dominant variant is typically associated with a fresh wave of infections and deaths, e.g., Wave 2B being linked to the emerging Alpha variant, and Wave 3 to the Delta variant (Figure 2).

The period from the end of 2022 was characterised by a sequence of Omicron variants. In some cases, single subvariants took over rapidly (e.g., BA.2), while BA.4, BA.5, BE and BF coexisted for several months. Each time a new dominant subvariant emerges, there was a small wave of new cases and associated deaths (see Figure 2). However, Nyberg et al. (2022) found that the Omicron variant was much less severe and lethal than the Delta variant.

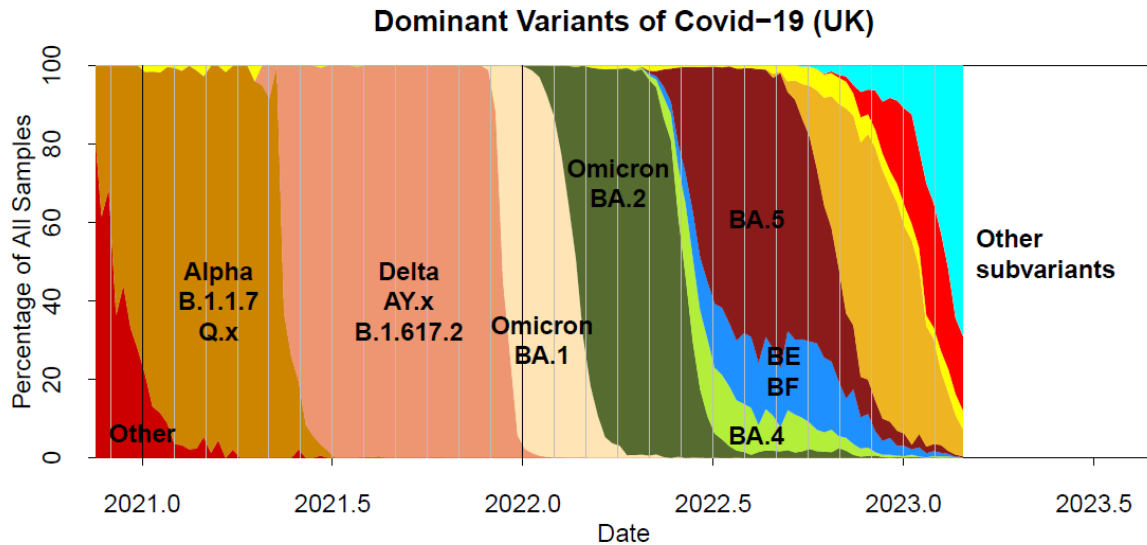


Figure 4: Dominant variants of Covid-19 in the UK between the end of 2020 and Spring 2023.

3. Phase 1: the year 2020

Section 2 gave an overview of the whole of the pandemic and its main features. In this section, we focus on Phase 1, a phase that is characterised as the period before vaccinations commenced in December 2020. After this date, the vaccination rollout affected infection and Covid-19 death rates in complex ways, and we discuss this in the following section. In contrast, Phase 1, while much worse in terms of death rates, is much more straightforward to analyse.

3.1 What do the data reveal in detail?

Our starting point in this investigation is to examine Covid-19 death rates by age. These are shown in Figure 5, where Covid-19 death rates are plotted separately against age for the first and second waves. Wave 1 (solid red line) runs for 26 weeks from March to August 2020 and Wave 2 (dashed red line) a further 26 weeks from September 2020 to February 2021¹³ – by which time vaccination was just starting to reduce death rates in the highest age groups. For comparison, Figure 5 also includes two further curves: all-cause death rates by age in a pre-pandemic year, 2018 (dotted lines), for which we have both accurate exposures and detailed cause-of-death data; and all-cause death rates, with deaths from external causes excluded (black line). All-cause death rates excluding external causes for 2015, 2016 and 2017 are also included for comparison (orange, blue and green lines).

¹³ Hence, Waves 1 and 2 cover a complete year.

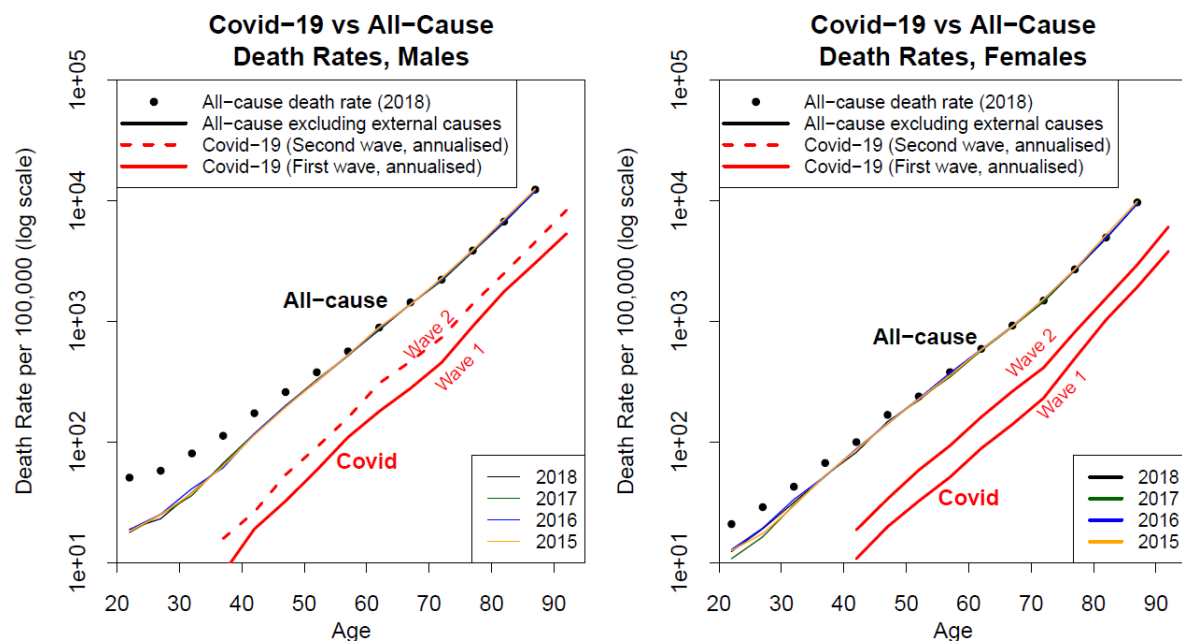


Figure 5: Death rates for males (left) and females (right) by age. Average Covid-19 death rates during Wave 1 (26 weeks from March to August 2020) and Wave 2 (26 weeks from September 2020 to February 2021). All-cause death rates and all-cause death rates excluding external causes are included for comparison and are based on 2018 values. All-cause death rates excluding external causes for 2015, 2016 and 2017 are also included for comparison.

Once external causes have been excluded, we can see a striking, near-parallel relationship between (log) death rates from all causes in a normal year and Covid-19 death rates in both Waves 1 and 2.¹⁴ It can be seen that the pre-Covid all-cause death rate curves for 2015 to 2018 are almost identical meaning that the near-parallel relationship is not sensitive to the choice of ‘normal’ pre-Covid year. The key implication of this is that a birth cohort’s ability to fight a Covid-19 infection is linked to that cohort’s ‘biological frailty’ (defined earlier as its one-year all-cause death rate).

We now build on this initial link between Covid-19 death rates and biological frailty in two ways. First, at any given age, there is evidence of significant heterogeneity in death rates. For example, within the same age group, some socio-economic groups experience higher death rates than other groups (see, e.g., Wen et al., 2021, 2023¹⁵, and Cairns et al., 2024a, and references therein). This reflects, for example, differences between subgroups in the prevalence of various health-related risk factors. For instance, more deprived areas will typically have a higher proportion of people who smoke. In turn, this leads to higher death rates from smoking-related diseases. So, we extend what we infer at the aggregate level from Figure 5 to include subgroups. Thus, we postulate that Covid-19 death rates will be proportional to all-cause death rates by age and by socio-economic or other subgroup.

¹⁴ The near-parallel relationship between Covid-19 death rates by age and all-cause death rates was first noted in an early-pandemic blog by Professor David Spiegelhalter (<https://medium.com/wintoncentre/what-are-the-risks-of-covid-and-what-is-meant-by-the-risks-of-covid-c828695aea69>) and was also hinted at by Professor Moshe Milevsky in a blog showing two parallel lines (<https://moshemilevsky.com/covid-19-and-longevity-risk/>). See, also, Figure 12.1 of Spiegelhalter and Masters (2021) (not split between waves 1 and 2) which indicates a continuation of the near parallelity to even younger ages.

¹⁵ See, also, the LIFE App at Cairns et al. (2024b).

Second, we know that Covid-19 infection rates vary significantly between socio-economic subgroups and between age groups. Everything else being equal, if the infection rate is doubled, then we would anticipate that the Covid-19 death rate would also double approximately.

Building on this, we can decompose Covid-19 death rates (annualised) as:

$$m_C(i, x) \equiv m_A(i, x)\widetilde{IR}(i, x)RF(i, x) \quad (1)$$

where, for subgroup i at age x ,

- $m_C(i, x)$ is the Covid-19 death rate (annualised);
- $m_A(i, x)$ is the corresponding all-cause death rate (annualised) in a normal pre-Covid-19 year (i.e., biological frailty);
- $\widetilde{IR}(i, x) = IR(i, x)/\tau$ is the annualised infection rate (where $IR(i, x)$ is the infection rate, the proportion who have become newly infected (i.e., the incidence of new infections) over a defined period of time of length τ years; this is annualised by dividing by the measurement period, τ); and
- $RF(i, x)$ is a balancing term that we define as ‘relative frailty’ (RF). This measures the infection fatality rate due to Covid-19 (as defined in equation (2) below) of an infected individual *relative* to annual all-cause mortality within the same subgroup, i , and at the same age x .

We need to be careful about interpreting the various components of Equation (1) in relation to the measurement period, τ . In particular, the proportion of the total population who die as a result of an infection during this time period is $1 - \exp(-m_C(i, x))$. Annualisation (e.g., as shown in Figure 5) implies that we assume that new infections continue to occur for a full year at the same rate.

Subgroups might be based on a number of different geodemographic measures, e.g., the Index of Multiple Deprivation (IMD),¹⁶ ethnic group, employment group, region and urban-rural location. It is well known (see, e.g., Wen et al, 2023 and references therein) that, at a given fixed age x , average all-cause mortality in these subgroups varies considerably. Cairns et al. (2024a) find that most of this variation in pre-Covid *all-cause mortality* was due to socio-economics factors, implying that region or other locational identifiers had relatively little *additional* impact. This is a very important finding which we utilise below.

The infection rate, $IR(i, x)$, also varies in potentially different ways between subgroups. The pattern of variation in infection rates between socio-economic subgroups might be quite different from the pattern of variation in all-cause death rates in the same subgroups: for example, infection rates show significant dependence on regional and other factors in a way that is not evident in all-cause death rates (e.g., Ward et al., 2020, and, by inference, regional variation in Figure 2 and urban-rural variation in Figure 3).

Next, we define the ‘infection fatality rate’, a quantity that has received considerable attention since the early stages of the pandemic (see, e.g., Verity et al., 2020):

$$IFR(i, x) \equiv m_A(i, x)RF(i, x) \quad (2)$$

which measures the (approximate) proportion of people in group i at age x who are newly infected and who subsequently die from Covid-19. This implies that Equation (1) can be rewritten

¹⁶ The Index of Multiple Deprivation measures relative deprivation (ranked from most to least deprived) in small areas in England called lower-layer super output areas. For further details, see <https://www.gov.uk/government/statistics/english-indices-of-deprivation-2019>

$$m_c(i, x) \equiv \widetilde{IR}(i, x)IFR(i, x). \quad (3)$$

Further, note that the IFR, as defined in Equation (2), is a ‘rate’, from which we can derive the (exact) probability of death from Covid-19, given an individual has just been infected, as $1 - \exp(-IFR(i, x))$.

3.2 The Proportionality Hypothesis

We are now in a position to make a formal statement of the Proportionality Hypothesis. There are two versions (subgroup and aggregate) depending on the granularity of the data that are available.

The subgroup Proportionality Hypothesis states that:

PH_{SG}: The Covid-19 infection fatality rate, $IFR(i, x)$, is approximately proportional to the all-cause death rate, $m_A(i, x)$, by both age and subgroup.

PH_{SG} can also be expressed in two alternative, but equivalent, ways:

- 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, $m_A(i, x)\widetilde{IR}(i, x)$.
- Relative frailty, $RF(i, x)$, is approximately constant across ages and subgroups.

In Sections 4 and 5, we will consider the aggregate Proportionality Hypothesis.¹⁷ This version of the hypothesis is of interest and relevance when suitable data by subgroup are not available. In this case, we can still model Covid-19 death rates by age, but at the national rather than subgroup level. In this case, Equation (1) simplifies to

$$m_c(x) \equiv m_A(x)\widetilde{IR}(x)RF(x), \quad (4)$$

where $\widetilde{IR}(x)$ is the annualised infection rate by age, $IR(x)/\tau$.

The aggregate Proportionality Hypothesis is then formally stated as:

PH_A: The Covid-19 infection fatality rate, $IFR(x)$, is approximately proportional to the all-cause death rate, $m_A(x)$, by age.

PH_A can also be expressed in two alternative, but equivalent, ways:

- The Covid-19 death rate, $m_c(x)$, is approximately proportional to the product of the all-cause death rate and the annualised infection rate by age at the national level, $m_A(x)\widetilde{IR}(x)$.
- Relative frailty, $RF(x)$, is approximately constant across ages at the national level.

Further, if the annualised infection rate, $\widetilde{IR}(x)$, is approximately constant across ages, then the Covid-19 death rate, $m_c(x)$, will be proportional to $m_A(x)$ – which is what Figure 5 shows for ages above 40. However, there is no guarantee that the infection rate will always be constant across ages. For this reason, we decided to express the primary version of the Proportionality Hypothesis in terms of the proportionality between the infection fatality rate and the all-cause death rate, since this does not require the constancy of the infection rate across ages to hold.

¹⁷ In particular, we will discuss the IFR in the context of the Proportionality Hypothesis in more detail in Sections 4.2 and 5.2.

3.3 Assessing the Proportionality Hypothesis

Assessing the Proportionality Hypothesis is cleanest in Phase 1, since vaccinations had not yet commenced and there were few, if any, reinfections. We now look at two sources of data that, in combination, demonstrate consistency with the subgroup Proportionality Hypothesis, PH_{SG} in Phase 1.

Infection rates

The first large-scale analysis of infection rates was published as part of the Imperial College REACT study (Ward et al., 2020). This study established that around 6% of the population had been infected by Covid-19 during the first wave of the pandemic (based on antibody prevalence). It then looked at which subgroups had higher or lower antibody prevalence. The following is a summary of the conclusions of this study concerning infection rates:

- There was no significant difference in infection rates between males and females.
- There were only relatively modest differences between deprivation quintiles (the most deprived had about 10% higher infection rates than the least deprived).
- Similarly, there were only relatively modest differences between age groups, especially at higher ages. Compared with REACT's 35-44 reference age group, younger age groups (e.g., the 18-24 age group) had infection rates that were about 40% higher. However, older age groups had at most 20% higher infection rates than the reference age group. This level of variation would not significantly alter our conclusion that the lines in Figure 5 are almost parallel above age 40.¹⁸
- There were much higher levels of variation by:
 - region (consistent with Figure 2; e.g., infection rates in London were 200% higher than the more-rural South West);¹⁹
 - ethnic group (Blacks had 100% higher infection rates than Whites);
 - employment group (patient-facing healthcare and care-home workers had infection rates up to 200% higher than 'other workers');
 - household size (seven-or-more-person households had a 60% higher infection rate than single-person households).

Infection rates on their own do not prove the Proportionality Hypothesis, but we use the information above (specifically the relatively modest variation in the first wave between deprivation groups) in the mortality analysis that follows.

Age-standardised mortality rates by deprivation decile and region

Figure 6 plots age-standardised mortality rates (ASMRs)²⁰ relative to the least deprived decile 10. The grey bars show ASMRs for all-cause mortality with values that are consistent with ASMRs in pre-pandemic years.

In the summer months of 2020, ONS (2020) published analysis of Covid-19 death rates by IMD decile and region. A key graphic focused on the impact of deprivation (redrawn as the

¹⁸ Because of the log scale in Figure 5 and the fact that all-cause mortality varies significantly, perturbing infection rates by +/- 20% at one or other end of the age range is only just noticeable and certainly does not alter the conclusion that the lines are nearly parallel above age 40.

¹⁹ This contrasts with the findings discussed above of Cairns et al. (2024a) who conclude that region is not significant in a normal non-Covid year.

²⁰ Despite being called a mortality rate, the ASMR is calculated using death rates.

dark blue bars in Figure 6 for the period March to August 2020) with the conclusion that the more deprived groups had been disproportionately affected by Covid-19. Relative to the grey bars, it appeared that Covid-mortality exhibited increasing levels of mortality as deprivation increased. However, we know that, compared with other parts of the country, London (a) had much higher Covid-19 infection rates and (b) has disproportionately higher levels of deprivation. The combination of the two in a city hosting a significant percentage of the total population of England (around 16%²¹), was potentially distorting the Covid-19 ASMRs for the most-deprived deciles during the first wave of the pandemic.

The ONS simultaneously published ASMRs by region, allowing us to assess this potential regional distortion. We used a simple log-bilinear model for the Covid-19 ASMR in region r and decile i :

$$ASMR(r, i) = \exp(\beta(r) + \gamma(i))$$

with $\gamma(10) = 0$ to ensure identifiability. This model allows us to separate the regional effects (which the REACT study, reported above, identified as significant) from the deprivation effects. The orange bars in Figure 6 show Covid-19 death rates by deprivation decile relative to the least-deprived decile 10, with regional effects filtered out: that is, the orange bars are equal to $100\exp(\gamma(i))$, for deciles $i=1, \dots, 10$. They are much closer in height to the grey bars for all-cause mortality. Thus, after controlling for regional effects, we can conclude that Covid-19 did not, in fact, disproportionately affect more deprived subgroups. Instead, it reflected (and, indeed, highlighted) pre-existing health inequalities.

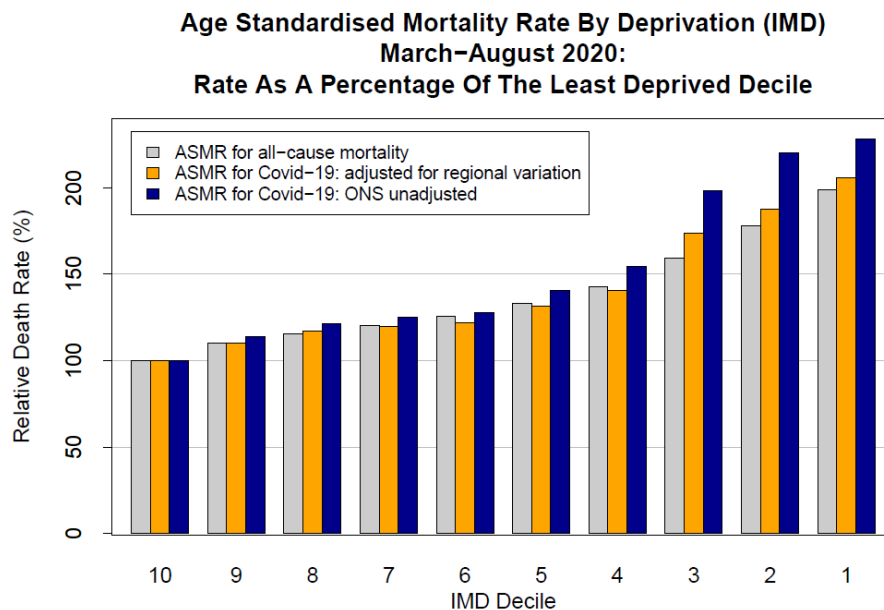


Figure 6: Death rates (%) relative to IMD decile 10 mortality for all-cause mortality (grey bars) and for Covid-19 mortality (orange and blue bars), March-August 2020. Blue bars: unadjusted Covid-19 mortality by decile. Orange bars: Covid-19 mortality, controlling for regional effects. IMD decile 1 is the most deprived, while IMD decile 10 is the least deprived.

²¹ <https://www.statista.com/statistics/294729/uk-population-by-region/>

What does the evidence indicate about the Proportionality Hypothesis?

Our broad conclusion from the evidence above is as follows. During the pre-vaccination Phase 1 of the pandemic, data by region and IMD decile are consistent with the subgroup Proportionality Hypothesis. In particular, relative frailty does not depend on socio-economic subgroup (specifically, data support the hypothesis that relative frailty does not depend on deprivation decile).

However, a lack of granularity in the data also meant that our ability to assess the full hypothesis (specifically jointly by subgroup and age) was not possible. Hence, the differences we observe in Covid-19 death rates between groups are most likely to be due to:

- existing differences in the average biological frailty of different socio-economic groups (which we measure as the average all-cause death rate in the coming year for these groups) as well as by age;
- differences in Covid-19 infection rates between groups as well as by age.

As we progress, in the next section, to a discussion of Phase 2, we should be mindful that the model expressed in Equation (1) needs to be adjusted to take account of vaccinations (and which vaccine) and prior infections of the virus (and which variant).

4. Phase 2: from the beginning of 2021 to mid-2022

4.1 What do the data reveal in detail?

Phase 2 covers the statistically much more complex period during which a variety of new vaccines were developed and rolled out across the population, new variants emerged (see Figure 4 above), improved treatments for severe cases were developed and many in the population caught Covid-19 for a second or third time.

Vaccinations and antibody prevalence

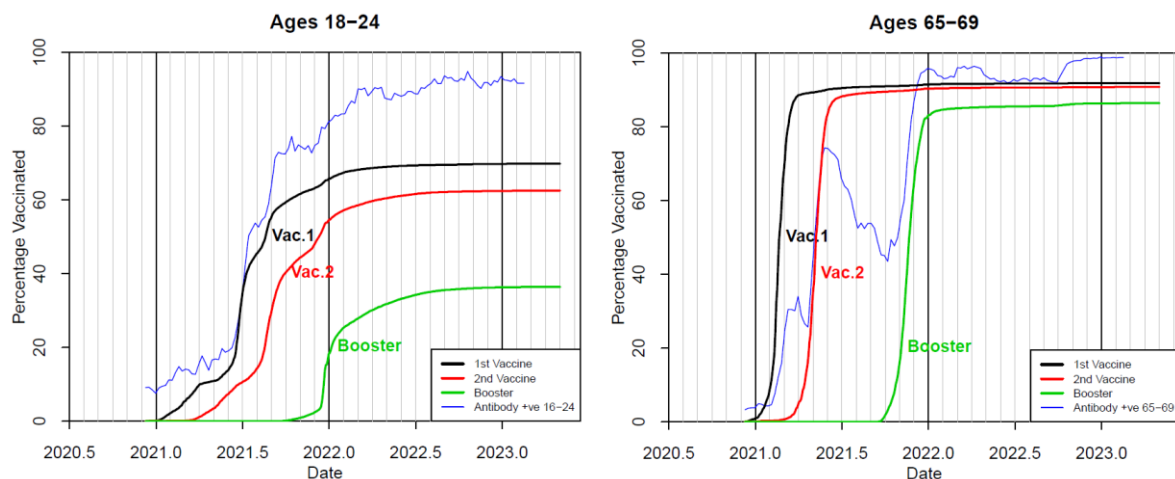


Figure 7: Vaccination and antibody status of specific age groups (18-24 (left) and 65-69 (right); males and females combined), December 2020 to April 2023. Vaccination status shows the proportion of the age group who have received 1, 2 or 3 (booster) vaccinations. Antibody status shows prevalence: the percentage of the age group who have significant levels of antibodies.

In Figure 7, we plot (by way of illustration) the proportion of people in the 18-24 and 65-69 age groups who had received 1, 2 or 3 vaccines (the latter known as the booster vaccine). Vaccines in the UK were given by age (eldest first) as well as to people who were clinically vulnerable due to pre-existing health conditions. Thus, we can see that the first main wave of vaccinations in the 18-24 age group lags the 65-69 age group by about 4 to 5 months. But, in the 18-24 age group, while most vaccinations took place in June 2021, around 15% had been vaccinated at an earlier stage due to clinical vulnerability. Second vaccinations took place about 8 to 12 weeks after the first, and then the third or so-called booster vaccination occurred in November and December 2021. Take-up of the three vaccines was lower in the younger than the older age groups. Potential reasons for this include (see, for example, ONS, 2021b): prior infections leading to a perception that there was less need for a vaccination; vaccine hesitancy related to social media discussions or general distrust; and personal judgements concerning the benefits of vaccination versus the potential side effects.

Figure 7 also shows the percentage of the population who were antibody positive. A positive test would reflect either prior infection or vaccination. For the 65-69 group, it can be seen that the first vaccine was only partially successful in building up antibody counts to a high enough level to produce a positive result. However, the second vaccine pushed antibody levels much higher. But we also see in the older age groups that antibody prevalence gradually declined after the first and second vaccinations. This might be a general age-related effect (older people are less able to maintain antibody levels), but it could also be that younger age groups were mixing more during 2021 and the consequent repeated exposure to infected persons might have kept their antibody levels higher. However, following the Autumn 2021 booster, antibody prevalence in the 65-69 and other older age groups was much more persistent. This might simply be because a third vaccine was required to keep antibody levels high, but it might also be due to a change to what was perceived to be a more effective vaccine for the booster.

By around the middle of 2022, antibody levels had reached very high levels in all age groups. This indicates that almost everyone by that time had either received one or more vaccinations or they had been infected with Covid-19. Specifically, it meant that most people who had chosen not to get vaccinated had now had at least one dose of Covid-19. We discuss this further in Section 5.1.

Infection rates (prevalence)

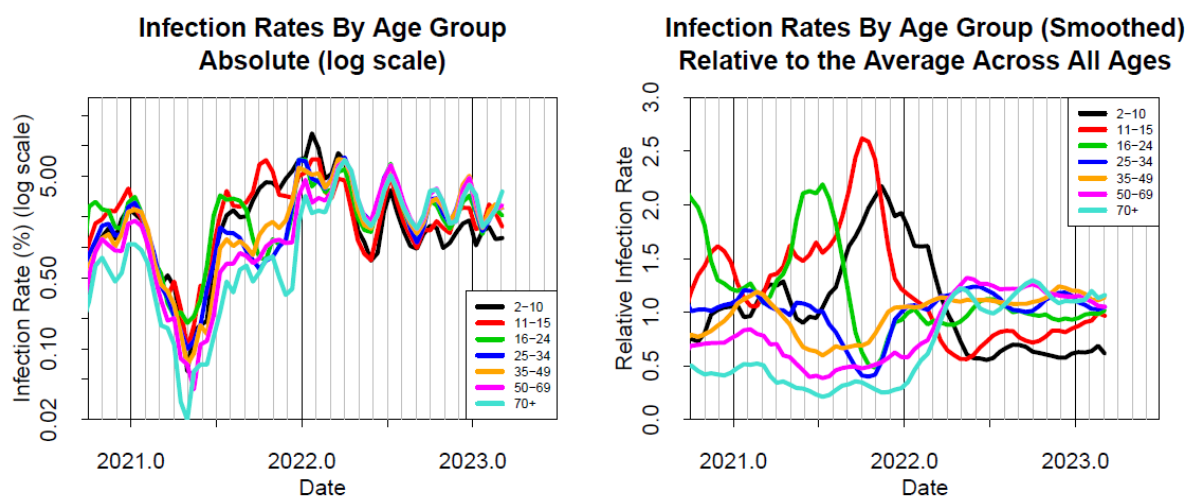


Figure 8: Left: infection rates (prevalence) over time (males and females combined; log scale) by age group, October 2020 to February 2023. Right: infection rates by age group relative to the arithmetic average over all age groups.

Figure 8 shows the prevalence of infection rates by age group over time using data available from the weekly ONS infection survey. The survey takes a large random sample each week and measures the *prevalence* of Covid-19.²² The nature of the survey means that the data can be considered to be reasonably accurate, as well as being consistent through time. An alternative measure could be based on self-reporting of new cases, for example, following home testing. However, the latter is prone to mis-reporting and growing levels of under-reporting, so results would not be consistent through time.

The left-hand plot of Figure 8 shows the raw infection rates with a Winter 2020-21 peak followed by low levels of infection in the Spring of 2021, and then rising again through the remainder of 2021 as the Delta and then Omicron variants took hold. However, the form of the left-hand plot makes it difficult to discern the relationship between different age groups during 2021. The right-hand plot shows the prevalence in each age group relative to the arithmetic average of the prevalences across the seven age groups. This reveals considerable swings between different age groups during 2021, being the result of at least two drivers. First, the vaccination rollout targeted older people first. This coincided with a drop in infection rates in the older age groups (70+ and 50-69) *relative* to the younger age groups in late 2020 and the early part of 2021: the older groups had been vaccinated; the younger groups had not. In the middle of 2021, infection rates for the middle-aged groups of adults dropped as they, in turn, got vaccinated, while relative infection rates amongst children continued to rise towards the autumn before dropping back once they too were vaccinated. Second, the benefits of and timing of vaccination were augmented by behavioural changes. During 2021, younger adults might have become less vigilant in terms of avoiding social contact (pushing up infection rates), while older adults continued to practise social distancing.

The first half of 2021 suggests that vaccination did, indeed, provide significant protection against infection (e.g., in the 50-69 group relative to the then unvaccinated 18-24 age group), although this fell well short of complete protection from infection.

Hospital admissions

Figure 9 shows daily hospital admission rates (counting admissions of persons with an existing Covid-19 diagnosis plus persons diagnosed with Covid-19 once they were in hospital). The left-hand plot shows the raw admission rates by age group, and the right-hand plot shows admission rates for each age group as a percentage of the numbers of admissions across all ages.

²² *Prevalence* measures the number of people who are infected with Covid-19 at any point in time. It is different from the *incidence* of new infections which requires further assumptions about the duration of infection to estimate accurately the number of new infections from the infection survey data.

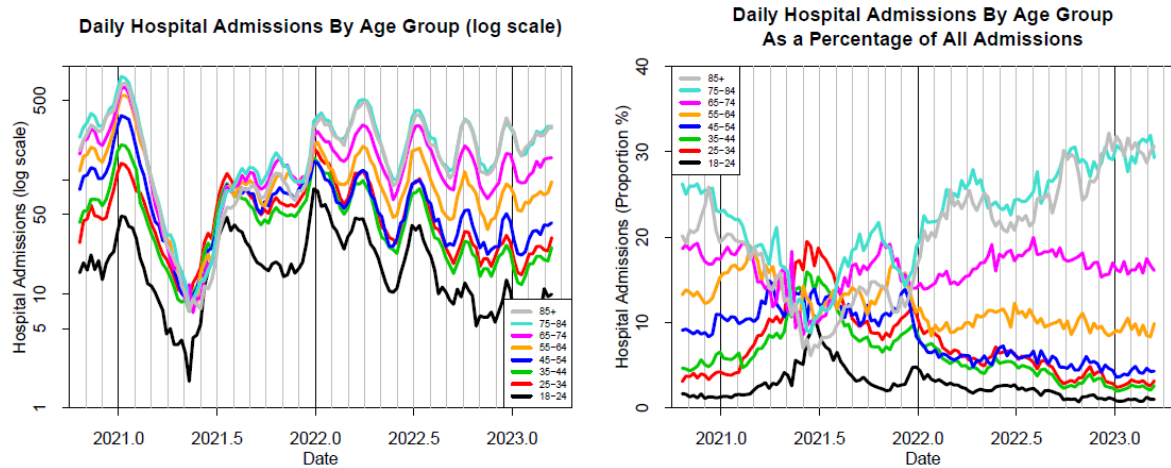


Figure 9: Covid-related hospital admissions by day (males and females combined) by age group over time, October 2020 to February 2023. Left: daily numbers of admissions. Right: daily admissions by age group as a proportion of daily admissions across all age groups.

The (generally) positive age gradient²³ on the far-left side of both plots (i.e., before the vaccination rollout in 2021) gives an indication that, given someone had become infected, age had a significant positive influence on subsequent hospitalisation rates (supporting the link to biological frailty by age in the Proportionality Hypothesis). However, during 2021, we see significant disruption to the age gradient which, most likely, reflected the rollout of vaccinations to older people first. Specifically, in the right-hand plot, the grey and light-blue lines (85+ and 75-84 age groups) declined rapidly in the early months of 2021, while the proportions for younger age groups (e.g., red and green lines) increased. Then, as the summer progressed and younger adults received two vaccines, the age gradient was restored. This strongly supports the conjecture that vaccination had a much more important effect than behavioural changes in all age groups. We conclude that vaccination did play a very significant role in preventing people infected with Covid-19 from needing to be hospitalised, since it reduced the average severity of an infection.

The age gradient in the Autumn of 2021 was not as steep as it was in December 2020. The most likely explanations for this are: the higher vaccination rates in the older age groups (Figure 7); behavioural changes; and a transition from the Alpha to the milder Delta variant. We then see that the age gradient changed again between the Autumn of 2021 and Spring 2022, narrowing first before widening again to near December 2020 levels. In the intervening 6 months, there was the booster rollout (initially narrowing the gap between age groups), a rise in antibody levels to the high 90%^s (Figure 7), and a transition to the even milder Omicron variant.

Average ages

²³ In this case, the age gradient refers to the generally positive, although not necessarily a completely monotonic, relationship between hospital admissions and age. In Figure 9, it is measured by the vertical spreads between the lines for any given date. The wider the vertical spreads, the steeper the gradient.

Average Age of People Affected By Covid-19

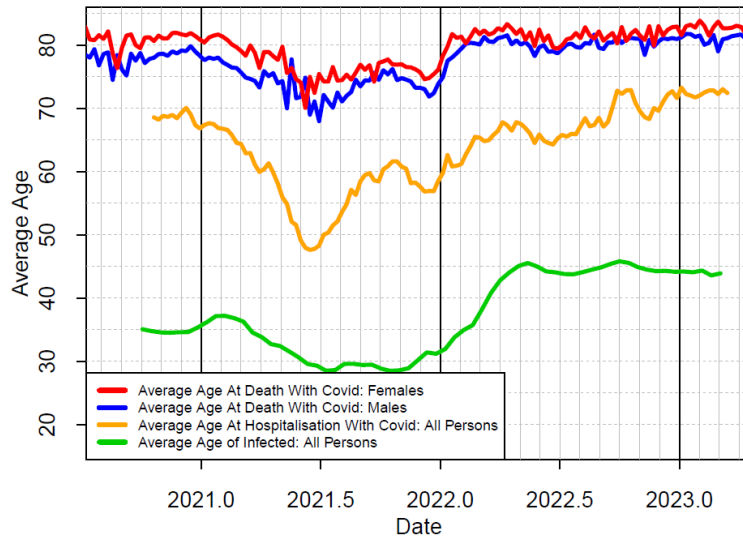


Figure 10: The average age of people affected by Covid-19 at various levels of severity of outcome, July 2020 – April 2023. Green line: average age of persons infected based on the ONS infection survey. Orange: average age of persons with a Covid-19-related hospital admission. Blue (males) and red (females): average age of those who died with Covid-19 (i.e., Covid-19 mentioned on the death certificate).

The complex age-related information about infection rates, hospitalisation and deaths can be summarised into average ages. All data sets use age groups rather than single ages, and we use the mid-point of each age range in the calculation of the average ages. Data are available for deaths of males and females separately, so we present average ages for both in Figure 10 over time. Data for hospitalisations and infections are for males and females combined. Average ages at death for females are about 2 years higher than males, but, otherwise, males and females moved up and down in synchronisation with each other.

The average age of those infected with Covid-19 was relatively low, although this was consistent with the general age profile of the population, allowing for higher infection rates at younger ages as a consequence of closer social contacts. The average age of those hospitalised was significantly higher, reflecting the observation that, given someone was infected, the older they were, the more likely they were to be hospitalised (in line with the age gradient seen in Figure 9). The average age at death was higher still, implying that the death rate of those hospitalised was higher for older people.

Now consider how the average ages changed as the pandemic progressed. In the early part of 2021, the average age at infection, hospitalisation and death dropped significantly, reflecting the impact of vaccination. The average age at hospitalisation dropped by as much as 20 years by the middle of 2021, providing strong support for the proposition that vaccination had a very strong impact on the severity of the infection, by keeping more people out of hospital.²⁴ Average ages then rose as the rollout of vaccines 1 and 2 covered all age groups, but then dipped temporarily again at the end of 2021 due to the booster vaccines starting with the elderly. Most boosters had been delivered by the end of January 2022.

²⁴ We think it is possible to go further than this and argue that the size of the drop is so big and the timing so closely linked to the timing of vaccination by age group that the big dip in Figure 10 could only be down to vaccination. Variation in infection rates could make a small contribution to the dip, but it would have to be a very big coincidence that infection rates tipped one way and then went into reverse in the middle of June 2021.

A final observation is that, after March 2022, the average age of those infected rose above its December 2020 level, while, in contrast the average ages at hospitalisation and death were similar to December 2020. The reasons for this are not clear, but it could be a complex interaction of age effects related to vaccination, type of vaccine, behaviour, prior infection with Covid-19 and the variant. For example, Nyberg et al. (2022) indicate that the difference between Omicron and Delta in terms of rates of hospitalisation was bigger at higher ages. Thus, everything else being equal, the transition from Delta to Omicron reduced the spread between the average ages of infection and hospitalisation. However, it is difficult to see this in the data as the booster rollout was happening at the same time.

Severity and lethality

Individual cases of Covid will lie somewhere on a scale from mild to life-threatening, and, ultimately, death in some cases. We refer to this sliding scale as ‘severity’. As a way of measuring severity at the aggregate level in an objective way through time, we consider the ‘infection hospitalisation rate’ (IHR) which we define as the number of hospitalisations (as defined above in subsection ‘Hospital admissions’) divided by the number of new infections for the same age group.

‘Lethality’ refers to the probability that a newly infected person subsequently dies with (or due to) Covid-19. This equates lethality with the ‘infection fatality rate’ (IFR; Equation (2)). We calculate IFRs by 10-year age group: 35-44, 45-54, etc.

The IHRs and IFRs for each age group are plotted in Figure 11. Each line has a certain amount of volatility due to sampling variation in the data, but the trends for each age group in each plot are clear. Each has a downward trend (before stabilising in mid-2022), reflecting a combination of the benefits from vaccination (lowering both the IHR and IFR), better treatments (lowering the IFR), reinfections (mainly arising in 2022) and changes in the severity and lethality of new variants. But the key takeaway from this plot is that, by the middle of 2022, the IFR was about 1/20th to 1/30th of what it was around the end of 2020: a remarkable reduction in such a short period of time. In comparison, the IHR had fallen to around 1/10th of what it was over the same period, indicating that the proportion of severe cases had also fallen significantly. In both panels, there is a strong and well-defined age gradient.

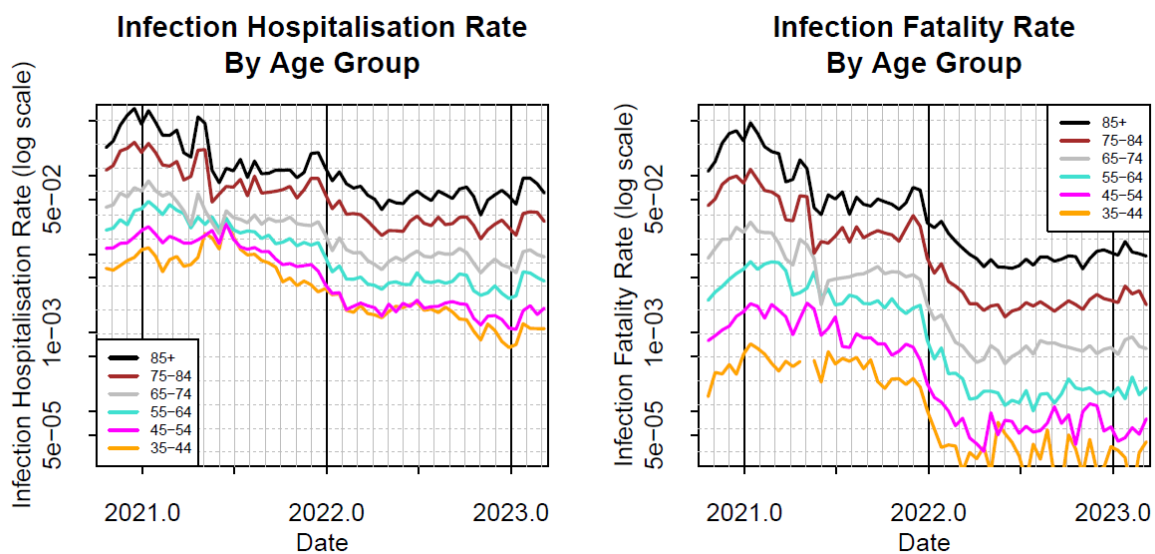


Figure 11: Infection hospitalisation rates by age group (left), and infection fatality rates by age group (right), October 2020 – February 2023.

In each of the plots in Figure 11, we can see that the lines for the six age groups bunch up somewhat in the middle of 2021, reflecting the age-related vaccine rollout. We can also compare Autumn 2021 (all adults vaccinated who chose vaccination) with late Spring 2022 (all adults having received the booster). The spread between the age groups in the IHR plot widened over this period to December 2020 levels. The reason for this widening is not clear, and is, in fact, not consistent with the analysis of Nyberg et al. (2022) which suggested that the transition from Delta to Omicron benefitted older adults more than younger adults, nor with the lower uptake of the booster vaccine amongst the younger adults. But the widening spread might be consistent with a greater proportion of infections in younger adults being reinfections compared to older adults. Everything else being equal, an individual who has recovered fully from a prior infection might have a higher degree of protection from future infections than a fully vaccinated individual who has avoided infection.

Our findings about the IHR are, however, consistent with those of Ferdinands et al. (2022). They focused on vaccine effectiveness – how good vaccines are at preventing hospitalisation relative to the unvaccinated group – and found that this depended on which vaccine has been received, the variant of Covid, and the time elapsed since vaccination. Vaccine effectiveness was found to wane over time at a rate that supports the case for annual boosters – as in the case of flu.

Death rates by vaccination status

The ONS also publishes monthly death rates by vaccination status (ONS, 2023). Figure 12 shows the results for males in their 70s, a particularly high risk group. In general, the data show that death rates involving Covid-19 were very significantly lower for people in the fully vaccinated group compared with the unvaccinated group (Figure 12, right panel). However, these rates are not adjusted for geodemographic differences between the vaccinated and unvaccinated groups, so it is less clear how much of the spread in Covid-19 death rates is due to vaccination and how much due to other differences. We can gain some insight into this by comparing the Covid-19 mortality spread (Figure 12, right panel) with the all-cause mortality spread (Figure 12, left panel) for the same month. Two observations are notable. First, all-cause mortality was also significantly lower for the fully vaccinated group.²⁵ Second, the spread between the fully vaccinated and unvaccinated groups was much wider for Covid-19 mortality than for all-cause mortality. Both these observations indicate that there were clear benefits from a programme of full vaccination. Patterns for other age groups are similar.

²⁵ The vertical spread in all-cause mortality by vaccination status is consistent with vaccine uptake being heavily dependent on socio-economic group.

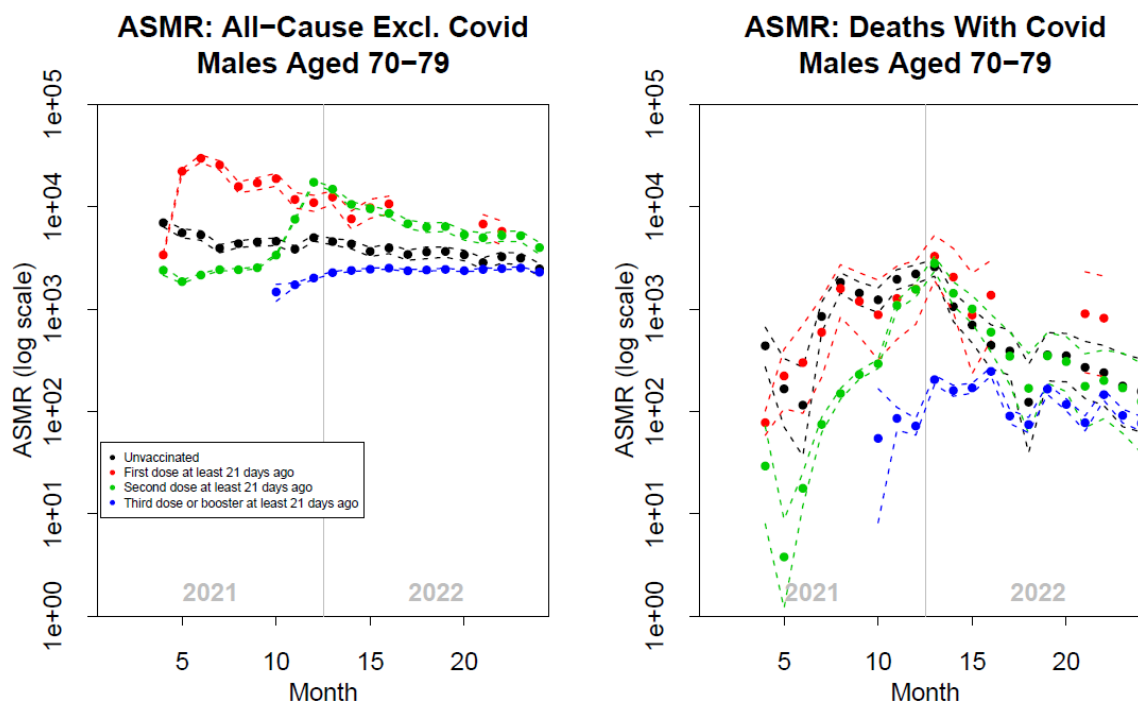


Figure 12: Age-standardised mortality rates (ASMRs) per 100,000 by month for males aged 70-79 by vaccination status, April 2021 to December 2022. Left: ASMRs for all-cause mortality excluding deaths with Covid-19 mentioned on the death certificate. Right: ASMRs for deaths with Covid-19. Dots show the ONS best estimates (source: ONS 2023). Dashed lines: 95% confidence limits (ONS 2023). Dots and confidence intervals are omitted when there is insufficient data. The fully vaccinated group jumps from the red to green dots in month 4 and from the green to blue dots in month 10. Month 1 is January 2021.

As we move towards the middle of 2022, however, the spread between the fully vaccinated and unvaccinated groups narrowed, although the vertical ranking is broadly maintained. Most likely this is because, as remarked earlier, almost all of the unvaccinated group had had at least one dose of Covid-19 and so their levels of immunity might be similar to or even stronger than those in the fully vaccinated group. Some studies have investigated the relative benefits of prior infection over a full programme of vaccination. For example, Altarawneh et al. (2022) suggest that the future protection provided by each is about the same, but a combination of prior infection plus full vaccination (i.e., hybrid immunity) is stronger. The meta-analysis of Bobrovitz et al. (2022) draws similar conclusions. The numbers of people with hybrid immunity only began to rise steeply after the Omicron variant began to dominate. In due course, it will be possible to investigate how the protection enjoyed by people with hybrid immunity persists relative to vaccination only or prior infection only.

The other notable observation from Figure 12 is that those who dropped out of the vaccination programme after 1 or 2 vaccines ended up with higher all-cause mortality than the unvaccinated group. This is a feature that has been discussed in ONS (2021a) where two potential reasons are offered for the higher rate in this group. First, the demographics of the group that dropped out might be quite different from the general population. Specifically, people from more deprived groups and from certain ethnic groups were more likely to drop out: these are groups that also have significantly higher death rates than the population average. Second, people who were in very poor health and who, for example, might be undergoing significant medical treatment, were much more likely to defer getting their second (or booster) vaccination. For the 70-79 male age group, the spread that we observe

between the unvaccinated and those who dropped out after one vaccine (from June 2021 onwards) was much wider than can be explained by general socio-economic differences (see, e.g., Wen et al., 2021), so the wide spread in this case was likely to be driven by drop-outs due to very poor health.

Part of the spread in all-cause mortality between the unvaccinated and the fully vaccinated groups (Figure 12, left panel) might be due to socio-economic differences between those who declined to be vaccinated versus those who were fully vaccinated. If there were no benefits to vaccination, then the all-cause mortality spread between the two groups would be the same as for deaths with Covid-19 (Figure 12, right panel). The fact that the spread is significantly wider for Covid-19 deaths provides further evidence that vaccination reduced the lethality of Covid-19.

4.2 Assessing the Proportionality Hypothesis

We can now assess the Proportionality Hypothesis in Phase 2. As discussed earlier (Section 3.3), in Phase 1 of the pandemic, it was possible to establish that death rates by deprivation level were consistent with the subgroup Proportionality Hypothesis, due to the relative simplicity of the data (i.e., no vaccinations and one dominant variant). In Phases 2 and 3, the data are much more complex.

Let us start with the infection fatality rate (IFR) at the level of the individual. It will depend on several factors:

- the all-cause death rate for the individual's age and socio-economic subgroup (see Equation (2));
- new factors including vaccination status (unvaccinated or 1, 2 or 3 vaccines), which vaccines have been administered and when, prior infection status, and with which variant an individual is currently infected.

This would then allow us to generalise our definition of relative frailty (by generalising the formulation in Equations (1) and (2)) to

$$RF(i, \theta, t, x) \equiv IFR(i, \theta, t, x)/m_A(i, x). \quad (5)$$

In this equation, we have introduced a time element, t , to the IFR that is evident at the aggregate (national) level in Figure 11. Additionally, the socio-economic subgroups, i , are further subdivided to reflect the additional factors or characteristics listed in the second bullet point above and these are represented in Equation (5) by the vector θ .

For this generalised version of the subgroup Proportionality Hypothesis to be true, we would have to establish that, for each set of characteristics θ , relative frailty is approximately constant across ages and socio-economic subgroups. Unfortunately, data from the ONS are not available at this level of granularity (e.g., infection rates and death counts by socio-economic group and vaccination status are not available). This lack of detail is particularly relevant during Phase 2, when the proportions in the different vaccination groups were changing significantly from month to month (as shown in Figure 7). This means that we are not able to assess the subgroup Proportionality Hypothesis in Phase 2 (or indeed Phase 3).

We do, however, have data at the national level throughout all three phases and so this does allow us to calculate the aggregate version of relative frailty:

$$RF(t, x) \equiv IFR(t, x)/m_A(x) \quad (6)$$

where there is no subdivision by socio-economic subgroup, by vaccination status, or other characteristics. RF is shown in Figure 13. It declined substantially over time (since each line

mimics the shape of the corresponding IFR for each age group), but it stabilised from the middle of 2022.

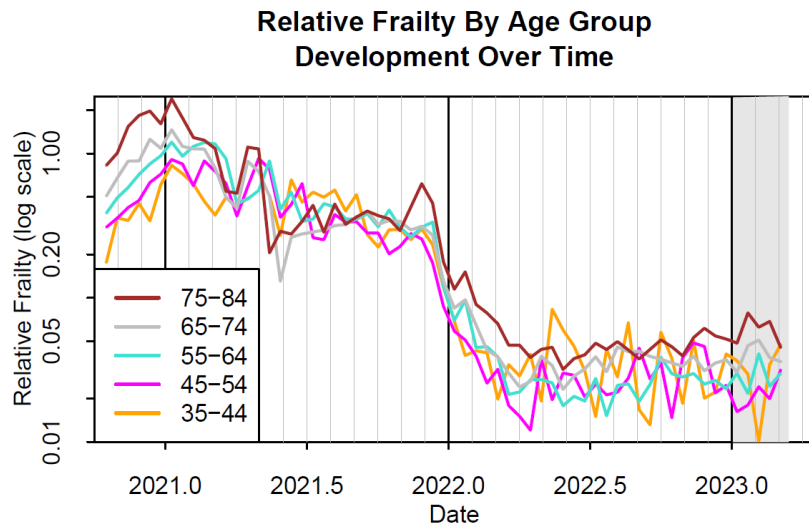


Figure 13: Relative frailty by age group over time, October 2021 – February 2023. The grey region on the right indicates the period where death occurrences might be underreported due to late registration and, hence, might be subject to later revision.

We therefore have the data to assess the aggregate Proportionality Hypothesis. Specifically, we can see from Figure 13 that relative frailty had some dependence on age in the last quarter of 2020 before vaccination began, with a steady age gradient (e.g., in December 2020, the brown line for ages 75-84 is higher than the orange line for ages 35-44).

The strict age gradient disappeared in the middle in 2021 during the vaccine rollout (e.g., the brown line is no longer consistently above the orange line), but then re-emerged during the early part of 2022, following the full rollout of the booster vaccines. It is clear that vaccination had a significant impact on relative frailty, mimicking the underlying IFR discussion above. But, without more granular data, it is difficult to disentangle precisely the impact of vaccinations from the impact of new variants and the development of new treatments.²⁶ So, although there are some interesting comments to be made on relative frailty, the lack of data granularity means that we cannot formally test for the continuing validity of the aggregate Proportionality Hypothesis during Phase 2 in the way that we were able to do during Phase 1.

Nevertheless, an alternative and powerful graphical illustration of the aggregate Proportionality Hypothesis can be seen in Figure 14. In the left-hand panel, we plot both the IFR in December 2020 (solid red line) and the all-cause death rate (solid black line) in a pre-Covid year (2018) for ages above 40. The plot shows an approximately parallel relationship between the IFR by age at the end of Phase 1 and the all-cause death rate by age (on a log scale). This mimics what we saw in Figure 5, but now with the effect of the infection rate filtered out. It is this *approximate* parallel relationship that justifies one of the ways we express the aggregate Proportionality Hypothesis, namely that infection fatality rates are *approximately* proportional to all-cause death rates by age.

²⁶ For example, anti-viral drugs, such as Nirmatrelvir-ritonavir (Paxlovid), and therapeutic interventions.

In the right-hand panel, we plot RF in December 2020 (solid blue line) and the all-cause death rate (solid black line) in a pre-Covid year (2018) for ages above 40. The plot shows the relative flatness of the RF curve across ages at the end of Phase 1 compared with the all-cause death rate curve.²⁷ It is this *approximate* constancy of the RF curve that justifies another way in which we express the aggregate Proportionality Hypothesis, namely that relative frailty is *approximately* constant across ages (relative to the dependence of all-cause death rates on age).

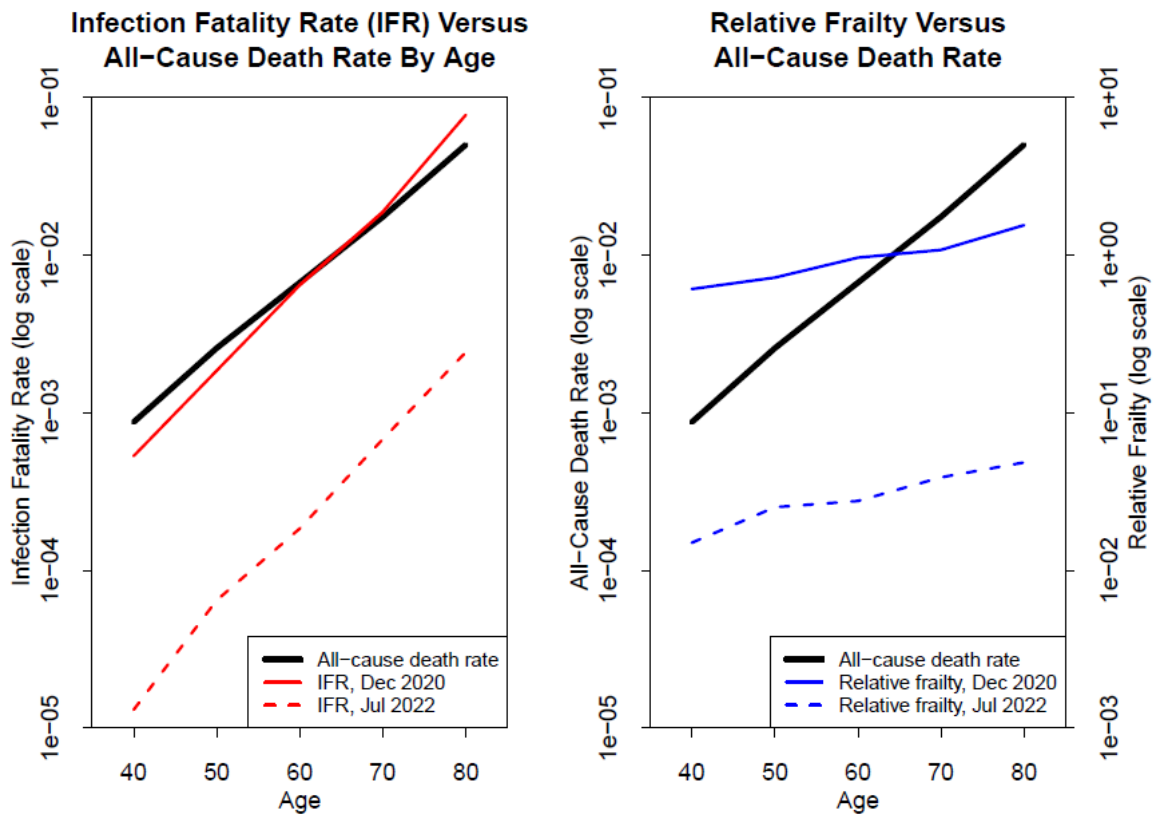


Figure 14: Left: the infection fatality rate (IFR) above age 40 in December 2020 (solid red line) and July 2022 (dashed red line) in comparison with the all-cause death rate for 2018 (solid black line). Right: relative frailty (RF) above age 40 in December 2020 (solid blue line) and July 2022 (dashed blue line) in comparison with the all-cause death rate for 2018 (solid black line).

5. Phase 3: mid-2022 onwards – transition from pandemic to endemic state

5.1 What do the data reveal in detail?

Data from the middle of 2022 onwards can be seen towards the right-hand sides of Figures 2 and 7 to 13. A key feature of these plots is the cyclical nature of infection rates, hospitalisations and death rates. The wave amplitudes are much lower than the initial waves

²⁷ The slight upward tilt of the solid blue RF line in the right-hand panel of Figure 14 follows from the slight upward tilt of the solid red IFR line relative to the solid black all-cause death line in the left-hand panel (see Equation (5)).

of the pandemic and the troughs are higher. Also, trends flatten out for some variables (e.g., relative frailty in Figure 13) as the virus transitioned from a pandemic to an endemic state.

Figure 7 also suggests that most of the adult population at all ages had either been vaccinated (with at least three doses) or infected with Covid resulting in a very high percentage carrying antibodies. During Phase 2, there was a rebalancing between age groups, discussed in Section 4, in terms of antibody prevalence with a resulting significant impact on the age distribution of infections, hospitalisations and deaths. In Phase 3, the high level of antibody prevalence makes it reasonable to assume that all age groups were approximately the same in terms of their vulnerability to infection without having to account for (as, for example, in early 2021) differing proportions of vaccinated and unvaccinated (or antibody positive and negative).

In younger age groups, there was a higher proportion of people who had not been vaccinated, but who had been infected compared to older age groups. By contrast, in the older age groups, there was a higher proportion of people who had had at least three vaccines – and it seems that it was the third booster vaccine that kept antibody levels high for much longer. Either way, from mid-2022, antibodies in the population stayed high for the remainder of the period covered by the data. A consequence is that relative frailty was much lower in Phase 3 than in Phase 1, reflecting the high levels of prior infections and/or vaccinations as well as improved treatments for serious cases of Covid-19. This therefore provides a good basis for assessing the Proportionality Hypothesis in Phase 3.

5.2 Assessing the Proportionality Hypothesis

Strong support for the aggregate Proportionality Hypothesis in Phase 3 comes from both Figure 13 and Figure 14. Figure 13 shows that, after July 2022, relative frailty for each age group remains reasonably constant (i.e., has no discernible trend), although there is some volatility across different months. In Figure 14, the dashed red and blue lines (for July 2022) are approximately parallel to the solid red and blue lines (for December 2020) – although at a lower level – so the same conclusions as in Section 4.2 hold, namely that infection fatality rates are *approximately* proportional to all-cause death rates by age (left-hand panel) and relative frailty is *approximately* constant across ages (right-hand panel).

As with Phase 2, the data in Phase 3 lack the granularity to allow us to check the validity of the subgroup Proportionality Hypothesis across different socio-economic groups or by vaccination and infection status.

6. The impact of Covid-19 on other causes of death

The ONS monthly mortality analysis also reports monthly death rates (ASMRs) for the top 10 causes that month. These can be compared to pre-pandemic levels by month, allowing us to gain some insight into the indirect impact of the pandemic on other causes of death.

Figure 15 shows monthly ASMRs for 9 of the leading causes of death over the period March (the top 5 only were reported initially by the ONS) or April (the top 10) 2020 to March 2023, with the named cause as the *underlying* cause, rather than rates with the named cause being *mentioned* on the death certificate. The blue lines (except for Covid-19) show the average for the previous 5 years. The orange line shows the actual monthly rates during the pandemic.

A number of different patterns can be observed. The blue lines reveal the seasonal nature of death rates to varying degrees: very strong seasonality for flu and pneumonia; negligible for

long-term diseases such as cancers. These pre-pandemic patterns match the findings of Marti-Soler et al. (2014): strong seasonality for cardiovascular diseases; no seasonality for cancers. Marti-Soler et al. (2014) also find that the seasonality is stronger for countries in colder climates compared with those nearer the Equator. Seasonality in cardiovascular diseases is reviewed in more detail by Fares (2013) who finds that the reasons for seasonality are quite varied, depending on the specific cause of death, so the causes of seasonality are complex and not always clear.

The most striking plot in Figure 15 is that for flu and pneumonia. The main observation is that, during the pandemic, standardised mortality rates were much lower than normal, but we also see much less seasonality than in normal years. The strong implication is that the combination of lockdowns, social distancing and home working substantially reduced people's exposure to pneumonia pathogens (bacterial and viral pneumonia), resulting in fewer cases and fewer deaths. This is backed up by data from the Royal College of General Practitioners (RCGP). The RCGP publish weekly reports on the numbers of new cases on communicable diseases²⁸ and the pattern of reported cases of pneumonia during the pandemic relative to the 5-year average for the same week in the year is very similar to the flu and pneumonia death rate.

At the other end of the spectrum, in the case of lung and bowel cancers, the pandemic had effectively no impact on the level or weekly pattern of standardised mortality rates. This is because these are long-term illnesses, so Covid-19 had no impact on the average level of deaths and little impact on accelerating death (i.e., no impact related to the waves of Covid-19). We can also observe that the lung cancer standardised mortality rate during the pandemic was a bit lower than the 5-year average, but this mostly likely reflected the gradual, pre-pandemic downward trend in lung cancer death rates associated with falling numbers of smokers.

²⁸ <https://www.rcgp.org.uk/representing-you/research-at-rcgp/research-surveillance-centre/public-health-data>

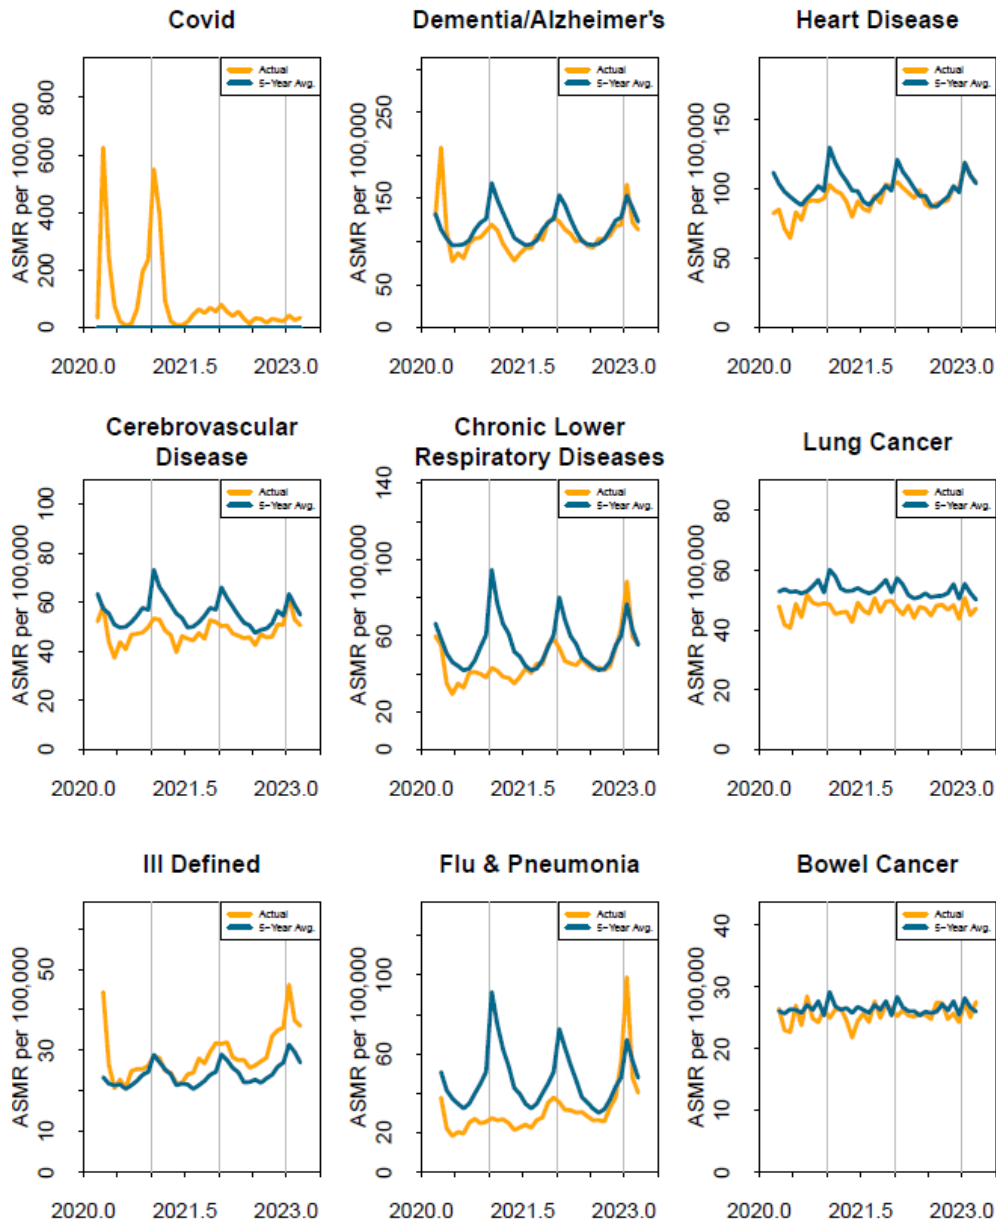


Figure 15: Monthly age-standardised mortality rates (ASMR) for 9 of the leading causes of death (deaths with the named cause as the underlying cause of death), April 2020 to March 2023. Orange lines: monthly ASMR's during the pandemic. Blue lines: 5-year average ASMR for each specific month.

For other causes of death in Figure 15, we can see varying degrees of seasonality in pre-pandemic standardised mortality rates, but much less seasonality during the pandemic, as with influenza and pneumonia. To take one example, chronic lower respiratory diseases typically have a chronic obstructive pulmonary disease (COPD), such as emphysema and bronchitis, as the main cause. Unlike flu and pneumonia, COPD is a long-term illness. But, unlike cancers, the strong seasonality suggests that the end of life is often triggered by the onset of a seasonal illness. During the pandemic, the change in behaviour again dampened down the seasonality in COPD death rates considerably, smoothing out the deaths over 2020 and 2021. Overall deaths were lower and the relative size of the reduction was greater than that for lung cancer. But does this mean that we will see a wave of delayed COPD deaths in the years to come (so instead of deaths *accelerated* by Covid-19, we have *decelerated* deaths due to COPD which are linked to social distancing, etc.)? Individuals diagnosed with COPD

are at increased risk of death from other causes, such as cardiovascular disease (Sin et al., 2007). So, the decelerated COPD deaths might appear with something other than COPD as the underlying cause of death.

In contrast with the general lower rates observed in Figure 15, some recent medical studies point to the greater onset of certain neurological diseases (in comparison with other respiratory diseases; Taquet et al., 2022) and heart disease (Xie et al., 2022; Wang et al., 2022). Taquet et al. (2022) noted that the onset rates for certain disorders could be dependent on the variant of Covid-19 with which individuals had been infected. Using the same large-scale health-insurance database, Xie et al. (2022) and Wang et al. (2022) concluded that cardiovascular events were up to twice as likely in the period after infection with Covid-19 compared with the control population (no Covid-19 infection). Wang et al. (2022) found that individuals who were hospitalised within 30 days of a positive test for Covid-19 had a much higher risk of cardiovascular problems, including death, than those who were not hospitalised. Furthermore, those who were not hospitalised had a lower risk of many cardiovascular events than the control population. This tentatively suggests a link, again, between underlying biological frailty and hospitalisation: those who are at higher risk of suffering a cardiac event prior to infection are more likely to be hospitalised following infection. However, Wang et al. (2022) did not directly address this question. So, although a Covid-19 infection is likely to increase cardiovascular risk in the short term, the increased rates observed might be explained, in part, by residual heterogeneity in individual risk, even after controlling for known risk factors.

Uusküla et al. (2022) used individual medical records for the general Estonian population to compare outcomes for people who had been infected with Covid-19 (with three levels of severity) versus those who had not. They found that, post-infection, death rates were significantly higher for most causes of death (cardiovascular, cancer, respiratory, and other) than the control group (no infection) for a period of weeks to as much as a whole year (for those aged 60+). Whether or not these cases are labelled as ‘long-Covid’,²⁹ the results collectively point to variation in individual biological frailty. A Covid-19 infection leaves an individual in a temporary state of heightened frailty that gradually reverts back to normal over a period of weeks or months: for example, in fighting the infection, the body uses up resources that take time to renew.³⁰

Bhaskaran et al. (2022) conducted a similar study based on individual health records in England. They drew similar conclusions that people discharged from Covid-19 hospital admission had a significantly higher risk of death following discharge for a period of time. Additionally, they found similarly increased mortality risk for people discharged from influenza-related hospital admissions during the period 2017-2019.

These studies allow us to broaden our observation above: following recovery from a severe illness, an individual will be subject to higher mortality risk for a period of time than would otherwise be the case after taking account of other risk factors.

Dale et al. (2023) provide other reasons for increased death rates in the general population from various potential future causes. Using primary care prescription data in the UK, they found that prescription initiation rates for key cardiovascular preventative medicines fell significantly during the pandemic, with a prediction of higher future cardiovascular mortality as a result. Contrasting with Figure 15 (top right, deaths *due to* ischaemic heart disease), the

²⁹ Long-Covid is defined as symptoms, such as extreme tiredness, feeling short of breath, loss of smell, and muscle aches, lasting longer than 12 weeks after a Covid-19 infection; <https://www.nhs.uk/conditions/covid-19/long-term-effects-of-covid-19-long-covid> (accessed 6/8/23).

³⁰ This observation equally applies to the periods following other adverse health events, not just Covid-19.

numbers of deaths with ischaemic heart disease (IHD) *mentioned* on the death certificate were significantly higher than expected (that is, due to IHD or with IHD as a secondary or contributing factor).³¹ Partial explanations for this difference are emerging (e.g., delayed prescription initiation or missed cardiovascular operations), but more analysis is required to understand this phenomenon better.³² From a statistical perspective, we note that the excess deaths due to ‘ill-defined’ causes (Figure 15, bottom left; ICD 10: R00-99) could explain a large proportion of the increase in deaths with IHD mentioned. Verifying this would require publication or access to more detailed death records than are currently publicly available. At the present time, it seems plausible that the increase in ill-defined deaths could be the result of a combination of (a) delayed diagnosis and treatments, and (b) long-Covid.

There is also significant concern about increased cancer deaths in the short to medium term, resulting from late diagnosis of new cancers or delayed treatments for cancer during the national lockdowns. Morris et al. (2021) reported a significant reduction in suspected bowel cancer referrals and colonoscopies following the start of the first Covid-19 lockdown, resulting in around 3,500 fewer formal cancer diagnoses over a 6-month period. These missing cases will, inevitably, emerge later when the cancer has advanced and so the survival prospects for the patient will be lower. Englum et al. (2021) also report similar falls in the diagnosis of a range of cancers. Luo et al. (2022) model the effects of late diagnosis and delayed treatment in Australia. They estimate, for example, that the pandemic will result in 460 additional deaths and 437 accelerated deaths from bowel cancer in Australia. The timing and magnitude of additional cancer deaths will depend on how much a cancer advances over the 3 to 6 months of a diagnosis delay, what the subsequent impact is on survival and remission probabilities, and the general time lag between stage of cancer (1 to 4) at diagnosis and death.

A detailed analysis by cause of death is beyond the scope of this paper, but we note that, over time, a clearer picture will emerge in terms of additional cancer and other deaths that resulted indirectly from the Covid-19 pandemic. It is likely that the consequences of the Covid-19 pandemic will have a long tail.

7. Extrapolative models and the Proportionality Hypothesis

We now discuss how extrapolative mortality models might need to be modified (stopping short of a full exercise in model fitting) to incorporate the Proportionality Hypothesis by considering the Lee-Carter (1992) (LC) and the CBD-X model (Dowd et al., 2020). We will consider the national death rate only, $m(t, x)$. In related work, van Berkum et al. (2023) consider how to generalise the multi-population Li and Lee (2005) model to capture weekly effects (including both seasonality and Covid-19 waves). What we outline here is similar in spirit to van Berkum et al (2023) with the additional consideration of the Proportionality Hypothesis.

³¹ See, for example, the Office for Health Improvement and Disparities (OHID) visualisation tool at <https://app.powerbi.com/view?r=eyJrIjoiYmUwNmFhMjYtNGZhYS00NDk2LWFIMTAAtOTg0OGNhNmFiNGM0IiwidCI6ImVINGUxNDk5LTRhMzUtNGIyZS1hZDQ3LTVmM2NmOWRIODY2NiIsImMiOiJh9> (accessed 28/6/2024).

³² As an aside, the numbers of excess ‘with IHD’ deaths have recently been declining, so perhaps what has been observed is a temporary effect. See <https://app.powerbi.com/view?r=eyJrIjoiYmUwNmFhMjYtNGZhYS00NDk2LWFIMTAAtOTg0OGNhNmFiNGM0IiwidCI6ImVINGUxNDk5LTRhMzUtNGIyZS1hZDQ3LTVmM2NmOWRIODY2NiIsImMiOiJh9> (accessed 28/6/2024).

7.1 The Lee-Carter model

The standard LC model for all-cause death rates is $\log m(t, x) = \alpha_x + \beta_x \kappa_t + \epsilon(t, x)$, the sum of age and age-period terms plus $\epsilon(t, x)$, a set of independent, zero-mean error terms. If we first assume a constant Covid-19 infection rate and constant relative frailty across all ages, then the Proportionality Hypothesis would result in an all-cause death rate (including Covid-19) of $\log m(t, x) = \alpha_x + \beta_x \kappa_t + \delta$. The additional δ term is consistent with what we see in Figure 5, but, even with constant infection rates, this does not capture the time dependency. More generally, in the pandemic setting and being mindful of the Proportionality Hypothesis, the LC model needs to be modified by adding at least one further age-period term:³³

$$\log m(t, x) = \alpha_x + \beta_x \kappa_t + I_x \delta_t + \epsilon(t, x).$$

The interpretation of I_x and δ_t and justification for this form is as follows. We start with the aggregate Proportionality Hypothesis (Equation (4)) modified to incorporate a time dimension

$$m_c(t, x) \equiv \tilde{m}_A(t, x) \widetilde{IR}(t, x) RF(t, x)$$

where $\tilde{m}_A(t, x)$ represents the underlying all-cause death rate in the absence of Covid. The all-cause death rate including Covid-19 is then $m(t, x) = \tilde{m}_A(t, x) \{1 + \widetilde{IR}(t, x) RF(t, x)\}$. It follows that

$$\log m(t, x) \approx \log \tilde{m}_A(t, x) + \widetilde{IR}(t, x) RF(t, x).$$

Now suppose that we can write $\widetilde{IR}(t, x) = I_x^1 \delta_t^1$. This assumes that the relative age distribution of infections remains stable over time (as a first approximation), while δ_t^1 models the waves of the pandemic over time (e.g., Figures 2 and 8). Similarly, suppose that we can write $RF(t, x) = I_x^2 \delta_t^2$, where I_x^2 models the relative variation by age in the relative frailty (Figure 14), and δ_t^2 captures the decline in the relative frailty over time (Figure 13). The net result of this is that $\log m(t, x) \approx \alpha_x + \beta_x \kappa_t + I_x \delta_t + \epsilon(t, x)$ where $I_x = I_x^1 I_x^2$, and $\delta_t = \delta_t^1 \delta_t^2$.

It might be that more than one additional age-period component is required to capture properly the impact of Covid-19 on mortality rates. The assumptions in the preceding paragraph require stability in the relative age-profile of infections and relative frailty. But we saw (Figures 8 and 13) that these both tilted one way and then the other during 2021, implying that more than one additional age-period component might be necessary. This would be particularly relevant if (as in van Berkum et al, 2023) mortality rates are being modelled on a weekly basis rather than annual and over the full course of the pandemic.

7.2 A CBD-X model

We will take a simple version of the CBD-X models from Dowd et al. (2020) without a cohort effect:

$$\log m(t, x) = \alpha_x + \kappa_{1t} + \kappa_{2t}(x - \bar{x}) + \epsilon(t, x).$$

³³ Zhou and Li (2022) and Schnürch et al. (2022) take a similar approach.

This model can capture some of the Covid-19 dynamics in a way that is consistent with the Proportionality Hypothesis. With constant infection rates and relative frailty across ages, the resulting parallel shift (as in the LC model) would be captured by the κ_{1t} period effect, while the κ_{2t} period effect in combination with the age effect $(x - \bar{x})$ picks up some of the variation by age in infection rates and relative frailty (as the latter varies through Phase 2). But an additional age-period component similar to the modified LC model ($I_x \delta_t$) might be required if the linear age effect turns out to be insufficient.

8. Implications for insurance: future extreme scenarios and catastrophe bonds

In this section, we raise a number of issues relevant for insurance companies in the light of the pandemic. They are discussed only briefly here, and a more detailed analysis is left for further work.

Drivers of extreme mortality events include terrorist attacks (e.g., New York, September 2001), tsunamis (e.g., Sumatra, 2004), earthquakes (e.g., Haiti, 2010) and pandemics (e.g., Spanish Flu, 1918; Covid-19, 2020). Past experience³⁴ suggests that (deadly) pandemics are likely to be the most extreme mortality events even though they are likely to be much less frequent. It follows, therefore, that when modelling extremes at, for example, the 99.5% confidence level, we should focus on pandemics.

8.1 Future extreme scenarios

Our discussion in the preceding sections suggests the following: when modelling an extreme pandemic-related mortality event, it would be inappropriate to simply model this by adding a fixed percentage to regular death rates.³⁵ Instead, stochastic models need to be developed which allow for significant variation within a country, for example: at regional and sub-regional level; between urban and rural areas; between different socio-economic groups; and between different age groups. This variation will introduce a form of risk diversification, and insurers whose liabilities are well diversified across the groups above should benefit from that diversification through a lower provision for extremes. Conversely, an insurer that holds a more concentrated group of lives (e.g., a high percentage located in London and from a particular socio-economic group) might be expected to post higher reserves against extreme mortality events.

A model that assumes all groups experience the same magnitude of mortality shock effectively assumes that all groups are perfectly correlated, but the Covid-19 pandemic has proved otherwise. The challenge here concerns how much credit should be given to a well-diversified portfolio of lives, and, of course, we really have very little data with which to calibrate a model that allows for variation between groups. This remains an open question, but it should not be ignored.

³⁴ Past experience, of course, might not be the best guide to the future. Future larger extremes could also include terrorism involving biological agents (perhaps leading to a pandemic) as well as nuclear missile attacks and other acts of war (although these are typically not covered in standard insurance contracts).

³⁵ Other approaches to handling a Covid-19-type shock have been proposed by Schnürch et al. (2022) and Hanika (2024), but these do not consider the level of granularity proposed here.

8.2 Parametric mortality catastrophe bonds

The last two decades have seen a small but steady issuance of mortality catastrophe bonds, mainly by insurers and reinsurers; they have been used principally as hedging instruments by the same issuers. Many of these use national mortality indices to determine the payments to the issuer and bondholders, for example, parametric mortality catastrophe bonds (as discussed in Blake et al., 2019, or Blake and Cairns, 2021). The linkage to national mortality indices assumes a high correlation (in an extreme scenario) between the national mortality index and the (re)insurer's own portfolio of lives and associated sums-at-risk. However, the Covid-19 pandemic has revealed significant variation between regions and other groups, and so it is not clear that the assumption of a *high* correlation between portfolio and *national* mortality is valid. For example, based on Figure 3, national mortality might have been 15% higher than normal, but an insurer that is concentrated in highly urban areas might have experienced as much as 25% higher mortality.

This might cause issuers of mortality catastrophe bonds to prefer linkage to a customised mortality index (i.e., based on their own mortality experience). However, an alternative would be to subdivide the national mortality index into regional and socio-economic mortality indices,³⁶ in a way that mimics the current linkage in a parametric bond to several national mortality indices.³⁷

9. Conclusions

Although Covid-19 was evidently present in China towards the end of 2019, the pandemic finally took hold around the world early in 2020 resulting in a range of different public health measures, including temporary lockdowns and the development and rollout of vaccines, designed to slow down the spread of the virus and reduce cases and fatalities. These measures and their timing, as well as differing behavioural responses by individuals, resulted in quite different patterns of spread in different countries even during the first significant wave of the pandemic in the Spring of 2020.

Nevertheless, despite the differences in the timing and magnitude of the pandemic waves in different parts of the world, we believe that many features of the disease's impact in England are likely to apply to other countries as well. To identify the impact, researchers would need to filter out the effects of varying infection rates over time and between different age and socio-economic groups. We then conjecture that, *given an individual has been infected*, risk factors such as age, socio-economic status, ethnicity, and region have a similar impact on health outcomes in different countries. In particular, we would expect the Proportionality Hypothesis to hold, with Covid-19 infection fatality rates being proportional to all-cause death rates across ages and subgroups.

Supporting evidence for this comes from Atanasov et al. (2023), one of the few studies that has data available with sufficient granularity to be able to conduct a similar analysis to ours. They identify a pattern similar to our analysis of the subgroup Proportionality Hypothesis (Figure 6), using a measure which they call the Covid Excess Mortality Percentage (CEMP). Their data consists of death counts rather than ASMRs for the US state of Indiana, subdivided by age group, sex, race and socio-economic quintile. They found a very strong correlation

³⁶ In England, these might be based on national data subdivided by region and deprivation level, or using a customised index such as the Longevity Index for England (Cairns et al., 2023) or a proprietary index (e.g. Club Vita).

³⁷ Chen et al. (2024) have considered the design of a pandemic catastrophe bond based on infection and mortality rates.

between Covid deaths during the first wave of the pandemic and all-cause deaths in a non-Covid year by subgroup: a result that is consistent with the subgroup Proportionality Hypothesis. As with our analysis underpinning Figure 6, they do not make any adjustments for potentially different infection rates in different subgroups. Their further subdivision by age group also supports the subgroup and aggregate proportionality hypotheses by age.

Global insurers are recommended to investigate new extreme pandemic mortality scenarios that take into account the significant random variation between different geographical areas, age groups and socio-economic subgroups. They are also recommended to design mortality catastrophe bonds that incorporate a linkage to sub-population mortality indices.

The Proportionality Hypothesis also has implications for the pricing and valuation of annuities post-pandemic. This is investigated further in Cairns et al. (2024c).

Finally, we note that the next pandemic might be closer than we fear. In August 2024, the World Health Organization added the Black Death plague, bird flu and the mpox virus to its pandemic watchlist (Newey, 2024).

Conflicts of interest: the authors declare no conflicts of interest.

Competing interests: Andrew Cairns and David Blake are guest editors for the special issue of the *European Actuarial Journal* that published the best papers presented at *Longevity 18: The Eighteenth International Longevity Risk and Capital Markets Solutions Conference* that was held at Bayes Business School, London, on 7-8 September 2023.

Appendix A: Data sources

- Human Mortality Database (HMD): Short Term Mortality Fluctuations database
<https://www.mortality.org/Data/STMF>
- Office for National Statistics (ONS): deaths registered weekly in England and Wales
<https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/deaths/datasets/weeklyprovisionalfiguresondeathsregisteredinenglandandwales>
- ONS: Coronavirus (Covid-19) antibody and vaccination data for the UK
<https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/conditionsanddiseases/datasets/coronaviruscovid19antibodydatafortheuk>
- ONS: Deaths involving Covid-19 by vaccination status, England: deaths occurring between 1 January 2021 and 31 May 2022
<https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/deaths/bulletins/deathsinvolvingcovid19byvaccinationstatusengland/deathsoccurringbetween1january2021and31may2022>
- ONS: Deaths by vaccination status, England
<https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/deaths/datasets/deathsbyvaccinationstatusengland>
- ONS: Number of positive Covid-19 tests by genetic lineages seen in samples by week
<https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/conditionsanddiseases/datasets/covid19infectionsurveytechnicaldata>

- NHS: Daily Covid-related hospital admissions
<https://www.england.nhs.uk/statistics/statistical-work-areas/covid-19-hospital-activity/>
- ONS: Monthly mortality analysis, England and Wales
<https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/deaths/datasets/monthlymortalityanalysisenglandandwales>
- ONS: Monthly mortality analysis, England and Wales monthly bulletins
<https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/deaths/bulletins/monthlymortalityanalysisenglandandwales/latest>
- ONS: Antibody prevalence
<https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/conditionsanddiseases/articles/coronaviruscovid19latestinsights/antibodies>
- Coronavirus Data website: Vaccination uptake by age group
<https://coronavirus.data.gov.uk/details/vaccinations?areaType=nation&areaName=England>
- ONS: Infection survey, England
<https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/conditionsanddiseases/datasets/coronaviruscovid19infectionsurveydata>
- Royal College of General Practitioners (RCGP): Communicable and respiratory disease reports
<https://www.rcgp.org.uk/representing-you/research-at-rcgp/research-surveillance-centre/public-health-data>

Appendix B: Summary of findings

In this section, we summarise the findings of the paper.

What do the data reveal?

- Prior to the introduction of vaccination:
 - There were no differences in infection and death rates between males and females, although average ages at death for females were about 2 years higher than for males.
 - For those who became infected, there was a clear dependency between severity, hospitalisation, lethality and age. Given that someone had become infected, age had a significant positive impact on subsequent hospitalisation and death rates. Death rates were higher for older people.
 - There was considerable variation in death rates by region, sub-region and other subgroups (e.g., socio-economic groups).
 - However, after adjustment for regional effects, Covid-19 did not disproportionately affect more deprived subgroups. Instead, it reflected pre-existing health inequalities.
 - Accordingly, differences in Covid-19 death rates between groups were most likely because of:

- existing differences in the average biological frailty of different socio-economic groups (which we identify with the average all-cause death rate in the coming year for a particular group, excluding external causes) as well as by age;
 - differences in Covid-19 infection rates between groups, as well as by age.
- During the course of vaccine rollout in 2021 and 2022:
 - Vaccination provided significant but not complete protection against infection. It reduced infection rates and, for those who became infected, it reduced the average severity of the infection, as well as hospitalisation rates and lethality. Further, vaccination appears to have had a much more important effect than behavioural changes in all age groups
 - Death rates involving Covid-19 were very significantly lower for people in the fully vaccinated group compared with the unvaccinated group. Those who drop out of the vaccination programme after 1 or 2 vaccines ended up with higher all-cause mortality than the unvaccinated group, but this could be because the drop-outs were already in very poor health and some would have died before completing the vaccination programme. In short, there were clear benefits from a programme of full vaccination.
 - In the older age groups, antibody levels gradually declined after the first and second vaccinations, but were much more persistent after the third (booster) vaccine. It seems that it was the third vaccine, or the use of an improved vaccine, that kept antibody levels high for much longer.
 - By mid-2022, antibody levels had reached very high levels in all age groups. This indicates that almost everyone by that time had either received one or more vaccinations or they had been infected with Covid-19. Specifically, it meant that most people who had chosen not to get vaccinated had now had at least one dose of Covid-19.
 - The infection fatality rate was, by the middle of 2022, 1/20th to 1/30th of what it was at the end of 2020: being a combination of the impact of vaccination, new less lethal variants, and improved hospital treatments for Covid-19. Over the same period, the infection hospitalisation rate fell to 1/10th of what it had been, indicating the reduction in the severity of the infections.
 - The combination of vaccination and prior infection conferred stronger protection than either one on its own.

The Proportionality Hypothesis

- We propose two versions of the Proportionality Hypothesis: the subgroup hypothesis and the aggregate hypothesis.
- The subgroup Proportionality Hypothesis can be expressed in three mathematically equivalent ways:
 - The Covid-19 infection fatality rate is approximately proportional to the all-cause death rate across ages and subgroups
 - The Covid-19 death rate is approximately proportional to the product of the all-cause death rate and the infection rate across ages and subgroups

- Relative frailty (the infection fatality rate of a Covid-19-infected individual *relative* to that individual's all-cause mortality rate, excluding external causes) is approximately constant across ages and subgroups.
- The aggregate Proportionality Hypothesis can be expressed at the national level in the same three ways.
- On the basis of evidence from Phase 1, when Covid-19 death rates were not affected by vaccination and other factors, we conclude that the data are consistent with the subgroup Proportionality Hypothesis. In particular, Covid-19 infection fatality rates by age and subgroup are approximately proportional to all-cause death rates by age and subgroup. Further, relative frailty:
 - while having some dependence on age, has much lower age dependence than all-cause death rates;
 - does not depend on socio-economic subgroup (specifically, data support the hypothesis that relative frailty does not depend on deprivation decile).
- This implies that the differences we observe in Covid-19 death rates by age and subgroup are likely to reflect: existing differences in the average biological frailty (i.e., the average all-cause death rate) by age and subgroup; and differences in infection rates by age and subgroup.
- On the basis of evidence from Phase 2, we conclude that there is some supporting evidence for the aggregate Proportionality Hypothesis. Infection fatality rates are again approximately proportional to all-cause death rates by age. Similarly, the variation by age in relative frailty is much less than the variation by age in all-cause death rates (although the situation is complicated by the fact that relative frailty also depends on vaccination status, which vaccines have been received and when, prior infection status, and which variant is dominant at the time of infection); and the synchronisation of male and female death rates by age.
- In Phase 3, when the great majority of the population had either been fully vaccinated or infected with Covid-19, we find that the Proportionality Hypothesis still holds at the aggregate level: Covid-19 infection fatality rates by age group at the national level continue to be approximately proportional to all-cause death rates by age. Relative frailty continues to have only modest age dependence as in Phase 1, but its level is much lower than in Phase 1, reflecting the now high levels of protection resulting from vaccinations and/or prior infections, as well as improved treatments for serious cases of Covid-19; it also has much lower variation by age than all-cause death rates, as required by the hypothesis.

Other causes of death

- Lockdown and behavioural changes during the pandemic have had an impact on death rates from some diseases: most notably flu and pneumonia. Death rates from some other causes have exhibited less seasonality than in normal years, perhaps linked to lower seasonality in flu and pneumonia. However, some death rates, such as those relating to cancer, have not been impacted at all so far, although it is anticipated that these death rates might rise temporarily in the next few years due to late diagnosis and treatments of cancers during the pandemic. It is likely that the consequences of the Covid-19 pandemic will have a long tail.
- Some studies point to temporarily elevated death rates from other causes such as heart disease following recovery from Covid-19. In other words, the biological frailty of an individual who has recently recovered from Covid-19 remains at an elevated level for

a few weeks or months, while the individual makes a more general recovery from the infection, especially those who were hospitalised. This pattern of elevated biological frailty is similar to people recovering from a severe case of influenza.

Implications for the future

- Insurers should revisit their extreme pandemic mortality scenarios. Models should allow for significant random variation between different geographical areas, age groups and socio-economic subgroups.
- Mortality catastrophe bonds might not be as effective as hedging instruments as had previously been anticipated, unless they incorporate a linkage to sub-population mortality indices.

References

- Altarawneh, H.N., Chemaitelly, H., Ayoub, H. H., Tang, P., Hasan, M. R., Yassine, H. M., Al-Khatib, H.A., Smatti, M. K., Coyle, P., 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., Bertollini, R., and Abu-Raddad, L. J. (2022) Effects of previous infection and vaccination on symptomatic Omicron infections. *New England Journal of Medicine*, 387(1): 21-34.
- Atanasov, V., Barreto, N., Whittle, J., Meurer, J., Weston, B. W., Luo, Q. E., Franchi, L., Yuan, A. Y., Zhang, R., and Black, B. (2023) Understanding Covid-19 vaccine effectiveness against death using a novel measure: Covid excess mortality percentage. *Vaccines*, 11(2):379. doi: 10.3390/vaccines11020379. PMID: 36851256; PMCID: PMC9959409.
- 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. <https://doi.org/10.1371/journal.pmed.1003871>
- Blake, D., and Cairns, A.J.G. (2021) Longevity risk and capital markets: The 2019-20 update. *Insurance: Mathematics and Economics*, 99: 395-439.
- Blake, D., Cairns, A.J.G., Dowd, K., and Kessler, A.R. (2019) Still living with mortality: The longevity risk transfer market after one decade. *British Actuarial Journal*, 24: e1, 1-80.
- Bobrovitz, N., Ware. H., Ma, X., Li, Z., Hosseini, R., Cao, C., Selemon, A., Whelan, M., Premji, Z., Issa, H., Cheng, B., Abu Raddad, L.J., Buckeridge, D., Van Kerkhove, M., Piechotta, V., Higdon, M., Wilder-Smith, A., Bergeri, I., Feikin, D., Arora, R.K., Patel, M., and Subissi, L. (2023) Protective effectiveness of prior SARS-CoV-2 infection and hybrid immunity against Omicron infection and severe disease: a systematic review and meta-regression. *Lancet Infectious Diseases*, 23: 556-567.

Cairns, A.J.G., Kleinow, T., and Wen, J. (2024a) Drivers of mortality: Risk factors and inequality. *Journal of the Royal Statistical Society, Series A*.
<https://academic.oup.com/jrsssa/advance-article/doi/10.1093/jrsssa/qnae017/7615081>

Cairns, A.J.G., Kleinow, T., and Wen, J. (2024b) Longevity Index for England – LIFE. Version 3.2. [online app] https://andrewcairns.shinyapps.io/LIFEapp_Version3B/ (Accessed 13/8/2024)

Cairns, A.J.G., Blake, D. and Mathur, R. (2024c) ADM’s APPLE: Covid-19 and the Accelerated Deaths Model. Working paper, forthcoming.

Carannante, M., D’Amato, V., Haberman, S., and Menziatti, M. (2023a) Frailty-based Lee Carter family of stochastic mortality models. *Quality & Quantity*.
<https://doi.org/10.1007/s11135-023-01786-6>

Carannante, M., D’Amato, V., and Haberman, S. (2023b) Effect of Covid-19 frailty heterogeneity on the future evolution of mortality by stratified weighting. *Journal of Demographic Economics*, 89: 513-532.

Chen, A., Li, H., Mao, Y., and Zhou, Z.Q. (2024) Learning from Covid-19: A Catastrophe Mortality Bond Solution in the Post-Pandemic Era. Working paper.
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4461995

Dale, C., Takhar, R., Carragher, R., Katsoulis, M., Torabi, F., Duffield, S., Kent, S., Mueller, T., Kurdi, A., Nguyen Le Anh, T., McTaggart, S., Abbasizanjani, H., Hollings, S., Scourfield, A., Lyons, R. A., Griffiths, R., Lyons, J., Davies, G., Harris, D., Handy, A., Mizani, M. M., Tomlinson, C., Thygesen, J. H., Ashworth, M., Denaxas, S., Banerjee, A., Sterne, J. A. C., Brown, P., Bullard, I., Priedon, R., Mamas, M. A., Slee, A., Lorgelly, P., Pirmohamed, M., Khunti, K., Morris, A. D., Sudlow, C., Akbari, A., Bennie, M., Sattar, N., and Sofat, R. (2023) The impact of the Covid-19 pandemic on cardiovascular disease prevention and management. *Nature Medicine*, 2: 219–225.

Dowd, K., Cairns, A.J.G., and Blake, D. (2020) CBDX: A workhorse mortality model from the Cairns-Blake-Dowd family. *Annals of Actuarial Science*, 14: 445-460.

Englum, B.R., Prasad, N.K., Lake, R.E., Mayorga-Carlin, M., Turner, D.J., Siddiqui, T., Sorkin, J.D. and Lal, B.K. (2022) Impact of the Covid-19 pandemic on diagnosis of new cancers: A national multicenter study of the Veterans Affairs Healthcare System. *Cancer*, 128: 1048-1056.

Fares, A. (2013) Winter cardiovascular diseases phenomenon. *North American Journal of Medical Sciences*, 5: 266-279.

Ferdinands, J. M., Rao, S., Dixon, B. E., Mitchell, P. K., DeSilva, M. B., Irving, S. A, Lewis, N., Natarajan, K., Stenehjem, E., Grannis, S. J., Han, J., McEvoy, C. H., Ong, T. C., Naleway, A. L., Reese, S. E., Embi, P. J., Dascomb, K., Klein, N. P., Griggs, E. P., Liao, I.-C., Yang, D.-H., Fadel, W. F., Grisel, N., Goddard, K., Patel, P., Murthy, K., Birch, R., Valvi, N. R., Arndorfer, J., Zerbo, O., Dickerson, M., Raiyani, C., Williams, J., Bozio, C. H., Blanton, L., Link-Gelles, R., Barron, M. A., Gaglani, M., Thompson, M. G., Fireman, B. (2022) Waning of vaccine effectiveness against moderate and severe Covid-19 among adults

in the US from the VISION network: test negative, case-control study. *BMJ*, 379: e072141. doi:10.1136/bmj-2022-072141

Hanika, M. (2024) A Covid-19 stress test for life insurance: insights into the effectiveness of different risk mitigation strategies. *European Actuarial Journal*, 14: 525–550.

Huang, H., Milevsky, M.A., and Salisbury, T.S. (2017) Retirement spending and biological age. *Journal of Economic Dynamics and Control*, 84: 58-76.

Lee, R. D., and Carter, L. R. (1992) Modeling and forecasting US mortality. *Journal of the American Statistical Association*, 87: 659–75.

Li, N., and Lee, R.D. (2005) Coherent mortality forecasts for a group of populations: An extension of the Lee-Carter method. *Demography*, 42: 575-594.

Luo, Q., O’Connell, D.L., Yu, X.Q., Kahn, C., Caruana, M., Pesola, F., Sasieni, P., Grogan, P.B., Aranda, S., Cabasag, C.J., Soerjomataram, I., Steinberg, J., and Canfell, K. (2022) Cancer incidence and mortality in Australia from 2020 to 2044 and an exploratory analysis of the potential effect of treatment delays during the Covid-19 pandemic: A statistical modelling study. *Lancet Public Health*, 7: e537–48.

Marti-Soler, H., Gonseth, S., Gubelmann, C., Stringhini, S., Bovet, P., Chen, P.-C., Wojtyniak, B., Paccaud, F., Tsai, D.-H., Zdrojewski, T., and Marques-Vidal, P. (2014) Seasonal variation of overall and cardiovascular mortality: A study in 19 countries from different geographic locations. *PLoS ONE*, 9(11): e113500. doi:10.1371/journal.pone.0113500

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

Morris, E.J.A., Goldacre, R., Spata, E., Mafham, M., Finan, P.J., Shelton, J., Richards, M., Spencer, K., Emberson, J., Hollings, S., Curnow, P., Gair, D., Sebag-Montefiore, D., Cunningham, C., Rutter, M. D. , Nicholson, B. D., Rashbass, J., Landray, M., Collins, R., Casadei, B., and Baigent, C. (2021) Impact of the Covid-19 pandemic on the detection and management of colorectal cancer in England: A population-based study. *Lancet Gastroenterology & Hepatology*, 6: 199–208.

Newey, S. (2024) WHO adds Black Death plague, bird flu and mpox to new pandemic watchlist, *Daily Telegraph*, 13 August 2024, <https://www.telegraph.co.uk/global-health/science-and-disease/who-pandemic-black-death-bird-flu-pathogens-disease>

Nyberg, T., Ferguson, N. M., Nash, S., Webster, H., Flaxman, S., Andrews, N., Hinsley, W., Lopez Bernal, J., Kall, M., Bhatt, S., Blomquist, P., Zaidi, A. Volz, E., Aziz, N. A., Harman, K., Funk, S., Abbott, S., Hope, R., Charlett, A., Chand, M., Ghani, A. C., Seaman, S. R., Dabrera, G., De Angelis, D., Presanis, A., and Thelwall, S. (2022) Comparative analysis of the risks of hospitalisation and death associated with SARS-CoV-2 Omicron (B.1.1.529) and Delta (B.1.617.2) variants in England: A cohort study. *Lancet*, 399: 1303-1312.

Office for National Statistics (2020) Deaths involving Covid-19 by local area and socioeconomic deprivation: Deaths occurring between 1 March and 30 June 2020. <https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/deaths/bu>

lletins/deathsinvolvingcovid19bylocalareasanddeprivation/deathsoccurringbetween1marchand30june2020

Office for National Statistics (2021a) Deaths involving Covid-19 by vaccination status, England: Deaths occurring between 1 January and 31 October 2021. <https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/deaths/bulletins/deathsinvolvingcovid19byvaccinationstatusengland/deathsoccurringbetween1januaryand31october2021#monthly-age-standardised-mortality-rates-by-vaccination-status-non-covid-19-deaths>

Office for National Statistics (2021b) Coronavirus vaccine hesitancy in younger adults: June 2021. <https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/healthandwellbeing/articles/coronavirusvaccinehesitancyinyoungeradults/june2021>

Office for National Statistics (2023) Deaths involving Covid-19 by vaccination status, England: Deaths occurring between 1 April 2021 and 31 December 2022. <https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/deaths/bulletins/deathsinvolvingcovid19byvaccinationstatusengland/latest>

Savcisen, 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.

Schnürch S, Kleinow T, Korn R, Wagner A (2022) The impact of mortality shocks on modelling and insurance valuation as exemplified by Covid-19. *Annals of Actuarial Science*, 16(3):498–526.

Sherris, M., and Wei, P. (2021) A multi-state model of functional disability and health status in the presence of systematic trend and uncertainty. *North American Actuarial Journal*, 25: 17-39.

Spiegelhalter, D., and Masters, A. (2021) *Covid by Numbers*. Pelican Random House: Dublin.

Taquet, M., Sillett, R., Zhu, L., Mendel, J., Camplisson, I., Dercon, Q., and Harrison, P. J. (2022) Neurological and psychiatric risk trajectories after SARS-CoV-2 infection: An analysis of 2-year retrospective cohort studies including 1284437 patients. *The Lancet Psychiatry*, 9: 815-827.

Uusküla, A., Jürgenson, T., Pisarev, H., Kolde, R., Meister, T., Tisler, A., Suija, K., Kalda, R., Piirsoo, M., and Fischer, K. (2022) Long-term mortality following SARS-CoV-2 infection: A national cohort study from Estonia. *Lancet Regional Health Europe*, 18: 100394.

van Berkum, F., Melenberg, B., and Vellekoop, M. (2023) Estimating the impact of Covid-19 on mortality using granular data. Working paper, Research Centre for Longevity Risk, University of Amsterdam.

Vaupel, J.W., Manton, K.G., and Stallard, E. (1979) The impact of heterogeneity in individual frailty on the dynamics of mortality. *Demography*, 16:439-454.

- Verity, R., Okell, L.C., Dorigatti, I., Winskill, P., Whittaker, C., Imai, N., Cuomo-Dannenburg, G., Thompson, H., Walker, P. G. T., Fu, H., Dighe, A., Griffin, J.T., Baguelin, M., Bhatia, S., Boonyasiri, A., Cori, A., Cucunubá, Z., FitzJohn, R., Gaythorpe, K., Green, W., Hamlet, A., Hinsley, W., Laydon, D., Nedjati-Gilani, G., Riley, S., van Elsland, S., Volz, E., Wang, H., Wang, Y., Xi, X., Donnelly, C. A., Ghani, A. C., and Ferguson, N. M. (2020) Estimates of the severity of coronavirus disease 2019: A model-based analysis. *The Lancet Infectious Diseases*, 20: 669-677.
- Wang, W., Wang, C.-Y., Wang, S.-I., and Wei, C.-C. (2022) Long-term cardiovascular outcomes in Covid-19 survivors among non-vaccinated population: A retrospective cohort study from the TriNetX US collaborative networks. *eClinicalMedicine*, 53: 101619.
- Ward, H., Atchison, C., Whitaker, M., Ainslie, K.E.C., Elliott, J., Okell, L., Redd, R., Ashby, D., Donnelly, C.A., Barclay, W., Darzi, A., Cooke, G., Riley, S., and Elliott, P. (2020) Antibody prevalence for SARS-CoV-2 following the peak of the pandemic in England: REACT2 study in 100,000 adults. Imperial College, Working Paper.
- Wen, J., Cairns, A.J.G., and Kleinow, T., (2021) Fitting multi-population mortality models to socio-economic groups. *Annals of Actuarial Science*, 15: 144-172.
- Wen, J., Cairns, A.J.G., and Kleinow, T. (2023) Modelling socio-economic mortality at neighbourhood level. *ASTIN Bulletin*, 53: 285-310.
- Xie, Y., Xu, E., and Al-Aly, Z. (2022) Long-term cardiovascular outcomes of Covid-19. *Nature Medicine*, 28: 583-590.
- Zhou, R., and Li, J. S.-H. (2022) A multi-parameter-level model for simulating future mortality scenarios with Covid-alike effects. *Annals of Actuarial Science*, 16, 453–477.