



UvA-DARE (Digital Academic Repository)

Estimating the impact of the COVID-19 pandemic using granular mortality data

van Berkum, F.; Melenberg, B.; Vellekoop, M.

DOI

[10.48550/arXiv.2209.06473](https://doi.org/10.48550/arXiv.2209.06473)

Publication date

2022

Document Version

Submitted manuscript

[Link to publication](#)

Citation for published version (APA):

van Berkum, F., Melenberg, B., & Vellekoop, M. (2022). *Estimating the impact of the COVID-19 pandemic using granular mortality data*. (v1 ed.) ArXiv.

<https://doi.org/10.48550/arXiv.2209.06473>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Estimating the impact of the COVID-19 pandemic using granular mortality data

Frank van Berkum,^{*} Bertrand Melenberg[†] & Michel Vellekoop[‡]

September 15, 2022

Abstract

We present an extension of the Li and Lee model to quantify mortality in five European countries during the COVID-19 pandemic. The first two factors are used to model the pre-COVID mortality, with the first layer modelling the common trend and the second layer the country-specific deviation from the common trend. We add a third layer to capture the country-specific impact of COVID-19 in 2020 and 2021 in excess of the pre-COVID trend. We use weekly mortality data from the Short Term Mortality Fluctuations Database to calibrate this third factor, and we use a more granular dataset for deaths in the Netherlands to assess the added value of more detailed data. We use our framework to define mortality forecasts based on different possible scenarios for the future course the pandemic.

1 Introduction and motivation

In this paper we use a three-layer Li and Lee model [14] to quantify mortality in five European countries during the COVID-19 pandemic, namely, Belgium, France, Germany, Great Britain, and the Netherlands. The first two layers model the pre-COVID mortality using annual data, where the first layer quantifies the common trend in the five countries and the second layer the country-specific deviation from the common trend. Our model adds a third layer to capture the country-specific impact of the new cause of death COVID-19 in the years 2020 and 2021 in excess of the pre-COVID trend. To model and quantify this third factor, we make use of weekly data that have become available in the Short Term Mortality Fluctuations (STMF) dataset, which forms a part of the Human Mortality Database (HMD) [19]. We annualize the weekly outcomes to generate annual mortality forecasts. We supplement our model with possible scenarios for mortality rate predictions taking into account the impact of COVID-19.

The Li-Lee model is an extension of the Lee-Carter model [13]. In the Lee-Carter model a factor which determines the impact of mortality changes per age is estimated, and a time-dependent factor that describes the development over time when averaged among all ages, which therefore takes the form of a time series. Next to the Li-Lee model [14], which extends the Lee-Carter model by adding extra layers, many other modifications have been proposed, introducing, for example, additional factors [5, 20], cohort effects [21], or a more appropriate model for measurement noise [4]. What these and alternative models have in common (see [6] for a good overview) is the goal of an improved description of the development of survival probabilities over time for different age groups, since this allows better forecasts for all sorts of statistics for human survival and better pricing and risk management models for financial products that depend on it.

^{*}Corresponding author. University of Amsterdam, Dutch Actuarial Association's Working Group for Mortality Research & PwC the Netherlands, f.vanberkum@uva.nl.

[†]Tilburg University & Dutch Actuarial Association's Committee for Mortality Research, b.melenberg@tilburguniversity.edu.

[‡]University of Amsterdam & Dutch Actuarial Association's Committee for Mortality Research, m.h.vellekoop@uva.nl.

The calibration methods for parameters in statistical models such as the Lee-Carter model rely on historical data and these are usually available in the form of yearly observations of deaths and population numbers or related statistics. We use annual data from 1970 up to and including 2019, retrieved from the HMDB, to calibrate the first two factors of the Li-Lee model, i.e., the pre-COVID mortality. However, since COVID-19 only plays a role in the years 2020 and 2021, we use weekly data to model the third layer of the Li-Lee model. The transition from yearly to weekly data means that some adjustments in the model need to be made.

Looking at more granular data necessitates that seasonal fluctuations in deaths during the year are taken into account since the mere assumption that deaths will be distributed over the different weeks of the year in a uniform manner can already be refuted after a first casual glance at the data. But if the actual distribution of mortality over weeks can be reliably deduced from historical data, and a calibrated mortality model is available for observations before the start of the pandemic, then weekly observations of deaths and exposures during the pandemic should allow us to assess its impact in terms of a time series describing its severity and a factor per age-group which determines how much its members will be affected.

We perform our analysis for the two genders separately, since it is well-known by now that gender is an important risk factor for death due to COVID-19. The risk for men is higher than for women, while at the same time women may be more prone to what has become known as the “long COVID” form of the disease, which can cause debilitating symptoms for a very long time period but does not seem to be fatal. Such a difference between the genders is partially explained by the distinct reactions of the immune systems which have been observed for the body’s response to infection with the SARS-CoV-2 virus [25].

Apart from age and gender, we will not consider other characteristics which may influence an individual’s change in survival probabilities as a result of the SARS-CoV-2 virus.¹ We also will not make use of data that try to measure the number of deaths with known cause of death COVID-19 either. The actual number of deaths that are directly related to getting the disease will far outnumber those that have been confirmed by administrative and other data [1]. But more importantly, our aim is also to include indirect effects on mortality as a result of the pandemic as well.²

The two-layer Li-Lee model that we use to quantify the pre-COVID mortality shows in the first layer a clear common trend. The country-specific deviations of the common trend in the second layer turn out not to be stationary. Instead, we model these deviations, like the first factor, using random walks with drifts. These drift terms are not statistically significantly different from zero. Therefore, we set these drift terms equal to zero in the mortality projections.

To quantify the factors which extend the traditional Li-Lee model, we do not only need death counts, but also exposures at a weekly frequency. These exposures are not available, so we have to determine these ourselves. Given the available data, we can only approximate their values. Therefore, as a comparison, we also briefly introduce the impact of COVID-19 using a compositional data (CoDa) analysis. Such a CoDa analysis only requires death counts, but not the corresponding exposures. We conduct this CoDa analysis using Dutch weekly death counts per individual age in the years 2020 and 2021, which were provided by Statistics Netherlands. Focusing on the years 2020 and 2021, the outcomes of the CoDa analysis turn out to be quite similar to the outcomes based on the three-layer Li-Lee model.

¹Other risk factors, which can be medical, socioeconomic or pertaining to lifestyle choices, have been shown to increase the risk of a fatal outcome after an infection. In particular, cardiovascular disease [24], autoimmune diseases [9], diabetes [16], and obesity and increased blood pressure [10] appear to have a significant impact [8]. In addition, a study of COVID-19 deaths in British hospitals found a sharp increase in the risk for people with high scores on a measurement scale for deprivation, a variable that characterizes negative socioeconomic factors such as poverty, lack of social contacts and a lower level of education. Only part of this effect could be explained by the more common occurrence of existing medical conditions that are known to increase the risk to become severely ill as a result of COVID-19 [26]. Data from the UK also shows that the risk is not the same for people with different ethnicities: in a study that explicitly corrected for age, gender, preexisting conditions and socioeconomic factors, people with a white skin were found to have a lower risk of COVID-19 mortality than those with a different skin colour. A relatively recent large cohort study that took many of these different factors into account found that among the non-medical indicators, old age, male sex and black skin colour are the most severe ones [9].

²Such as the deferred care mentioned in footnote 1 for other conditions and the reduction in the number of deaths due to the flu or traffic accidents.

The weekly data in the Short Term Mortality Fluctuations (STMF) dataset only includes the death counts over five year age ranges. We use the Dutch weekly death counts per individual age in combination with the corresponding approximated exposures to investigate the impact of using five year age groups compared to individual ages. Comparing the resulting outcomes shows that the time trends are quite similar, but the age effects, after aggregating to five year intervals, might be quite sensitive to possible cohort effects. This applies in particular to ages that correspond to the baby boom generation, borne after the second world war.

After annualizing the weekly outcomes, we use our results to generate mortality forecasts. Since the impact of COVID-19 on future mortality is quite uncertain at this stage, we present these mortality forecasts for different scenarios, where each scenario represents a possible future evolution of COVID-19.

Our approach differs from related research. For example, Robben et al. [22] and Schnürch et al. [23] mention that the age profile of the impact may be different from the age profile for other causes of death, but this profile is not estimated. In the first paper, the maximum likelihood estimation problem for the time series of mortality dynamics from [11] is simply modified by giving less weight in the likelihood function to the observation years during the pandemic. By varying the corresponding weighting parameter and by adjusting the starting values of mortality projections using a modified version of the approach of Lee and Miller [12], different projections for the future forces of mortality in Belgium are generated and the effect on life expectancies is determined. Schnürch et al. [23] provide a comparative analysis of the extra deaths in 2020 using Lee-Carter models for different countries, and Cairns-Blake Dowd models for robustness checks. The authors do not introduce a new age-dependent factor but focus on the change in the time series for mortality; their approach can therefore be interpreted as an extension of the approach by Chen & Cox [7], where transitory jumps in mortality time series are assumed.

Two papers in which a distinctive age pattern is addressed explicitly are Liu and Li [15] and Zhou and Li [27]. In the first paper, the consequences of an age effect of a sudden mortality shock are analyzed, but under the assumption that the shock has been observed in the past and that it only affects the year in which it occurred. In [27] the Lee-Carter model is extended by a new age-and-time effect for COVID-19 and a new time series representing the overall impact over all ages. Parameters are estimated using a penalized quasi-likelihood maximization. The estimated age distribution of the effect of COVID-19 follows the shape of the pre-COVID distribution for mortality changes. This is even true for higher ages, possibly because it is hard to separate the effect of the pandemic from ordinary mortality changes in this approach.

The calibration results of the three-layer Li-Lee model for the five countries that we include in our study show a different pattern. Different age groups are affected differently in the different countries, but in all cases we find an increasing trend in age and negligible effects for the youngest ages. This is in contrast to the results reported in [27] but in line with studies in the epidemiological literature, such as the infection fatality ratios reported in the metastudy of Driscoll et al. [17]. We believe that this shows the merit of the adjustments to existing estimation methods which we propose in this paper: using a Li-Lee model for different countries to identify a common pre-pandemic trend, using weekly data during the pandemic to estimate the age distribution of its impact, and using robustness checks based on a more granular dataset for individual ages and a Compositional Data analysis to verify that our extrapolation method for unknown exposures is sound.

The structure of the remainder of this paper is as follows. Section 2 introduces the model that will be used to analyse weekly mortality observations. This section describes how pre-COVID mortality assumptions are derived, the estimation of seasonal effects in mortality, and how COVID-19 age and week effects are calibrated. Section 3 first presents results using Compositional Data analysis for which only death counts are needed. Then, using granular Dutch mortality death and exposure information, results are shown for the COVID-19 age and week effects (including various sensitivities), and a comparison is made between the COVID-19 effects in the Netherlands, Belgium, France, Germany and Great Britain. Section 4 illustrates how estimated COVID-19 age and week effects can be used to construct scenarios for mortality rate predictions that are adjusted for the impact of COVID-19. Finally, Section 5 concludes.

2 Model specification and calibration

In this section, we describe the framework for estimating the impact of COVID-19 on the level of mortality. We start with introducing our COVID-19 mortality model as an extension of the Li-Lee model for weekly observations to which an additional term is added to capture the impact of COVID-19. Then, we describe how the baseline level of mortality is calibrated using the usual Li-Lee model, and we investigate seasonal patterns in recent weekly mortality observations. Finally, we describe how weekly exposures and death counts are obtained and how the full model is calibrated.

2.1 A COVID-extension for the Li-Lee model

To assess the impact of the pandemic, we propose a model in which we distinguish a baseline specification for mortality in a country prior to the pandemic, an adjustment for seasonal effects when we make the transition from yearly to weekly data at the start of 2020, and a new age-dependent factor and time series to capture the effect of the pandemic. We thus assume that the logarithm of the force of mortality in country c for age-group³ χ and gender g in week w of year t has the following structure:

$$\ln \mu_{\chi,t,w}^{c,g} = \ln \mu_{\chi,t}^{c,g} + \ln \phi_{\chi,w}^{c,g} + \mathfrak{B}_{\chi}^{c,g} \mathfrak{R}_{t,w}^{c,g}. \quad (1)$$

The first term, $\mu_{\chi,t}^{c,g}$, is the force of mortality for the year t in a (two-layer) Li-Lee model [14]

$$\ln \mu_{\chi,t}^{c,g} = B_{\chi}^g K_t^g + \alpha_{\chi}^{c,g} + \beta_{\chi}^{c,g} \kappa_t^{c,g}. \quad (2)$$

This force of mortality combines a Lee-Carter specification in the first term, which forces of mortality for all countries in our chosen peer group have in common, with the last two terms that define a country-specific deviation from the common dynamics. The age-dependent parameters B_{χ}^g , $\alpha_{\chi}^{c,g}$ and $\beta_{\chi}^{c,g}$ and the time series K_t^g and $\kappa_t^{c,g}$ are calibrated using yearly historical data for time periods before the pandemic; the precise procedure will be discussed in the next subsection. The forecast values for $t = 2020$ and $t = 2021$ based on these parameters determine the baseline values for pre-pandemic mortality $\mu_{\chi,t}^{c,g}$ in those years since these are then based on historical data before COVID-19 had any impact.

In Equation (1) for the force of mortality per week, we add two more terms. The first, $\ln \phi_{\chi,w}^{c,g}$, is introduced because mortality is not evenly distributed over the different weeks of the year: there is usually more mortality in the cold winter months and less mortality during the milder months. The quantity $\phi_{\chi,w}^{c,g}$ describes this fluctuation of mortality over the different weeks and will be called the *seasonal effect*. The last term in (1), the product $\mathfrak{B}_{\chi}^{c,g} \mathfrak{R}_{t,w}^{c,g}$, is the third layer of our Li-Lee model. It combines a new age effect $\mathfrak{B}_{\chi}^{c,g}$ with a new time effect $\mathfrak{R}_{t,w}^{c,g}$ which is 0 for $t \leq 2019$. Our specification thus preserves the model structure of a Lee-Carter or Li-Lee model, while making it possible to work with a finer dataset of weekly instead of yearly data.

The age effect $\mathfrak{B}_{\chi}^{c,g}$ is expected to be very different from the values that are found for B_{χ}^g and $\beta_{\chi}^{c,g}$ because we know that excess mortality in 2020 and 2021 was largest among the highest age-groups. For this reason, it is less accurate to describe the effect of COVID-19 by only making an adjustment in the time series K_t^g and $\kappa_t^{c,g}$ for $t = 2020$ and $t = 2021$ while retaining the existing pre-pandemic model structure.

Note that we do not include a term that depends only on the age and the week (similar to the terms A_x^g and $\alpha_x^{c,g}$), because this would mean that we make the a priori assumption that there could be a lasting effect of the virus, even if values of the corresponding time series $\mathfrak{R}_{t,w}^{c,g}$ would converge to zero in the future.

When analyzing mortality on an annual basis it is common to assume the force of mortality $\mu_{x,t}$ to remain constant during the year. Since we consider mortality on a weekly basis, it may seem more appropriate to assume the force of mortality $\mu_{x,t,w}$ to gradually move from $\mu_{x,t,1} = \mu_{x,t}$ to $\mu_{x,t,w_t} = \mu_{x,t+1}$ where w_t equals the last week in year t . However, in Section 4 we show how the impact of COVID-19 on the level of mortality can be incorporated in mortality

³Note that χ may refer to an individual age $\chi = \{x\}$ or a group of ages $\chi = \{x_1, x_2, \dots, x_n\}$.

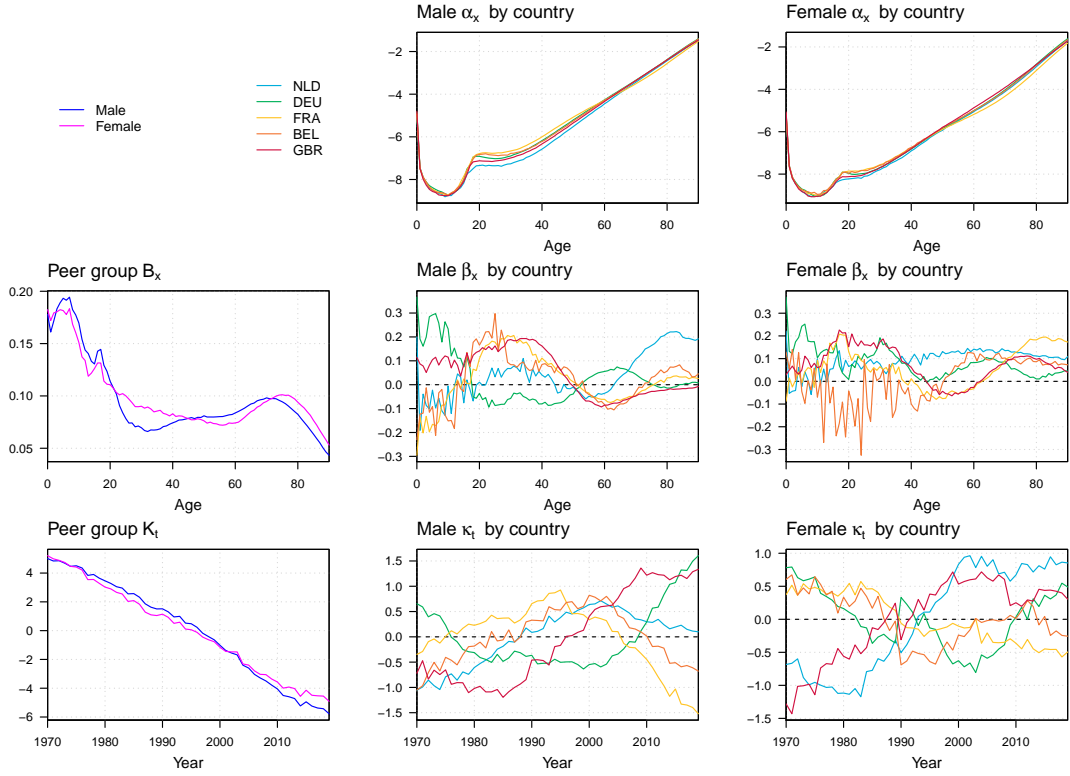


Figure 1: Estimated parameter values for the baseline model in (2). The left column shows B_x^g and K_t^g on the middle and the bottom row respectively. The middle column shows $\alpha_x^{c,g}$, $\beta_x^{c,g}$ and $\kappa_t^{c,g}$ on the top, middle and bottom row respectively when the gender g is equal to male, the third column shows the same information as the second column but for females.

forecasts. For this purpose, it turns out to be more convenient to assume that the baseline level of mortality remains constant during the year.

2.2 Estimation of the baseline

The baseline mortality model in (2) contains common age-dependent parameters B_x^g , country-specific age-dependent parameters $\alpha_x^{c,g}$ and $\beta_x^{c,g}$, and the time series for the common trend K_t^g and country-specific deviations from the trend $\kappa_t^{c,g}$. We use maximum likelihood methods to calibrate these parameters for individual ages, i.e., with $\chi = \{x\}$, for the time period before the pandemic started. The Human Mortality Database contains historical data in terms of deaths ($D_{x,t}^{c,g}$) and exposures ($E_{x,t}^{c,g}$) for individual ages for our peer-group of countries, which consist of the Netherlands, Belgium, Germany, France, and the United Kingdom of Great Britain and Northern Ireland, which we will abbreviate in the country set $\mathcal{C} = \{\text{NLD, BEL, DEU, FRA, GBR}\}$.

Data from the years $t \in \mathcal{T}^{\text{EU}} = \{1970, \dots, 2019\}$ and ages $x \in \mathcal{X} = \{0, \dots, 90\}$ are used for the calibration and we determine the relevant parameters in two separate steps. We first estimate the parameters which are common to all countries, B_x^g and K_t^g , using the aggregated deaths and exposures

$$D_{x,t}^g = \sum_{c \in \mathcal{C}} D_{x,t}^{c,g}, \quad E_{x,t}^g = \sum_{c \in \mathcal{C}} E_{x,t}^{c,g}.$$

Under the distributional assumption that deaths conditioned on exposures follow the Poisson distribution

$$D_{x,t}^g \sim \text{Poisson}(E_{x,t}^g \mu_{x,t}^{\text{comm},g}) \quad (3)$$

$$\ln \mu_{x,t}^{\text{comm},g} = A_x^g + B_x^g K_t^g \quad (4)$$

we can find parameter estimates (A_x^g, B_x^g, K_t^g) by maximizing the log-likelihood

$$\ln L^{\text{comm}} = \sum_{g \in \{m, f\}} \sum_{x \in \mathcal{X}} \sum_{t \in \mathcal{T}} \left(D_{x,t}^g (A_x^g + B_x^g K_t^g) - E_{x,t}^g e^{A_x^g + B_x^g K_t^g} \right) + C \quad (5)$$

with C a constant which does not affect the optimization of the likelihood. Once the estimates B_x^g and K_t^g have been determined (as well as A_x^g , which will no longer be needed in the sequel), we can start the second stage of the calibration based on the specification

$$D_{x,t}^{c,g} \sim \text{Poisson} \left(E_{x,t}^{c,g} \mu_{x,t}^{c,g} \right),$$

with $\mu_{x,t}^{c,g}$ as defined in (2) with $\chi = \{x\}$. This implies that the remaining parameters $\alpha_x^{c,g}$, $\beta_x^{c,g}$ and $\kappa_t^{c,g}$ can be determined by maximization of the log-likelihood

$$\ln L = \sum_{g \in \{m, f\}} \sum_{c \in \mathcal{C}} \sum_{x \in \mathcal{X}} \sum_{t \in \mathcal{T}} \left(D_{x,t}^{c,g} (\alpha_x^{c,g} + \beta_x^{c,g} \kappa_t^{c,g}) - E_{x,t}^{c,g} e^{B_x^g K_t^g + \alpha_x^{c,g} + \beta_x^{c,g} \kappa_t^{c,g}} \right) + \tilde{C}. \quad (6)$$

We impose the parameter restrictions $\|B^g\| = \|\beta^{c,g}\| = 1$ (with $\|\cdot\|$ the Euclidian norm) and $\sum_{t \in \mathcal{T}} K_t^g = \sum_{t \in \mathcal{T}} \kappa_t^{c,g} = 0$ when we maximize (5)-(6) for all countries c and genders g to ensure that no identification issues can arise.

Estimated parameters for the Li-Lee baseline model. The estimated parameters are shown in Figure 1. The B_x^g parameters show that the highest improvements in mortality in the observed period occurred at ages below 20 years. For higher ages the sensitivity to the K_t^g parameter is relatively stable, though ages 70-80 seem to benefit a bit more than average. The general trend in mortality across the countries, as represented by the slope of the K_t^g parameter, is stable, indicating that mortality (averaged over the countries) improved at a stable pace. The $\alpha_x^{c,g}$ parameters exhibit the well-known structure: the lowest level of mortality at around 10 years, the accident hump for young adults, and mortality linearly increasing by age for higher ages. Differences between the countries are small except for ages close to the accident hump, and French female at higher ages seem to have structurally lower levels of mortality than the other countries considered.

The country-specific improvements, captured by the $\beta_x^{c,g}$ and $\kappa_t^{c,g}$ parameters, are less straightforward to interpret. The $\kappa_t^{c,g}$ parameters indicate that countries may experience periods with higher and periods with lower improvement rates than the peer group. For example, the German male κ_t parameter decreases from 1970 to 1985, then increases until 2010, and remains steady afterwards. Such periods of higher and lower improvement rates can alternate from one year to the other, as indicated by the abrupt change in slope in those period effects. The $\beta_x^{c,g}$ parameters, which indicate how sensitive ages are towards changes in the $\kappa_t^{c,g}$ parameter, are not consistently above or below zero. Therefore, we cannot say in general that for a specific dataset mortality increased less or more severely for all ages compared to the peer group.

Forecasting period effects in Li-Lee baseline model. We have used data until 2019 to calibrate the Li-Lee baseline model. To calibrate the new age-dependent factor and time series for the effect of the pandemic, we need to predict the baseline mortality rates for the years 2020 and 2021. In Section 4 we investigate the impact of the pandemic on forecasts of cohort life expectancies, for which we need to forecast mortality far into the future. Hence, we need to specify a time series model for the calibrated period effects K_t^g and $\kappa_t^{c,g}$.

The common approach to forecast the period effects in the peer group of countries is using a random walk with drift, see, for example, [14] and [11]. The assumption of a random walk with drift implies that the shared mortality trend continues in the future. For the country-specific period effects, an autoregressive model of order one, including a constant, is often used. We aim to calibrate one large time series model for all period effects simultaneously. Preliminary analyses highlighted that for various country-specific period effects the mean reversion parameter is larger than one, indicating that the time series is not stationary. We therefore impose a random walk with drift for the country-specific period effects. For projection purposes, we neglect the

drift term for the country-specific period effects to ensure mortality in the different countries does not diverge. This approach is justified by the fact that the estimated drift terms for most country-specific period effects are not statistically different from zero.

We define the assumed time series for the period effect of the peer group of countries and for the period effect of the country-specific deviations as:

$$\begin{aligned} K_t^g &= K_{t-1}^g + \theta^g + \varepsilon_t^g \\ \kappa_t^{c,g} &= \kappa_{t-1}^{c,g} + \delta^{c,g} + \zeta_t^{c,g}. \end{aligned}$$

We assume that the error terms ε_t^g and $\zeta_t^{c,g}$ for $g \in \{m, f\}$ and $c \in \mathcal{C}$ follow a multivariate normal distribution with mean vector $\mathbf{0}_{12}$ and covariance matrix Σ . The parameters θ^g , $\delta^{c,g}$ and all elements of the covariance matrix Σ are estimated using maximum likelihood estimation techniques.

2.3 Incorporation of the seasonal effect

The specification for the weekly force of mortality in (1) contains a parameter for seasonal effect. We consider two approaches to incorporate a seasonal effect. If no seasonal effect is estimated in advance, which means $\ln \phi_{\chi,w}^{c,g}$ is taken to be zero, it will be included in the time series $\mathfrak{R}_{t,w}^{c,g}$. An alternative approach is to include a pre-determined estimator based on historical data which ensures that the product $\mathfrak{B}_{\chi}^{c,g} \mathfrak{R}_{t,w}^{c,g}$ only represents the observed deviation from the baseline predictions that are based on statistical information before the virus struck. We will designate the first approach with the term ‘‘time series for seasonal effect plus COVID-19’’ and the second approach with ‘‘time series for COVID-19’’. By comparing these two methods, we can investigate whether the two choices can lead to different conclusions about the impact of the pandemic.

To illustrate the second approach, Figure 2 shows how mortality spreads over different weeks of the year for the different countries. Values are aggregated over the two sexes and over all ages $x \in \mathcal{X}$. The gray lines for the years 2010 to 2019 show how mortality fluctuates over the weeks and we see a clear variation in time: for example, the severe flu wave around the tenth week of 2018 in the Netherlands is clearly visible in one of the gray lines. A value of 100% in the figure corresponds to the situation in which mortality is uniformly distributed in a year. The green line is the average over all gray lines, and thus equals the observed average effect per week. Mortality during the year is not evenly distributed: as expected, more people die in the winter months than during the summer months. Using cyclic cubic splines, we have estimated a smooth effect based on the annual historical observations shown here, and the result is represented by the orange line. The values at the beginning and at the end of the year ensure a smooth transition from week 52 to week 1.

2.4 Estimation of the impact of COVID-19

Once the Li-Lee model for the baseline and the seasonal effects have been calibrated, we can estimate the remaining parameters $\mathfrak{B}_{\chi}^{c,g}$ and $\mathfrak{R}_{t,w}^{c,g}$ that describe the effect of COVID-19. We make the assumption that during the pandemic the number of deaths for given exposures still follows a Poisson distribution which now includes the additional factors. This implies that for $t = 2020$ and $t = 2021$ we impose

$$D_{\chi,t,w}^{c,g} \sim \text{Poisson} \left(E_{\chi,t,w}^{c,g} \mu_{\chi,t}^{c,g} \phi_{\chi,w}^{c,g} \exp(\mathfrak{B}_{\chi}^{c,g} \mathfrak{R}_{t,w}^{c,g}) \right)$$

(where $\mu_{\chi,t}^{c,g}$ and $\phi_{\chi,w}^{c,g}$ are set to the earlier calibrated values), and the required dataset for calibration is

$$\{(E_{\chi,t,w}^{c,g}, D_{\chi,t,w}^{c,g}), \chi \in \mathcal{X}, g \in \{m, f\}, c \in \mathcal{C}, t \in \mathcal{T}, w \in \mathcal{W}_t\}.$$

We choose a subset \mathcal{X} of ages and the collections $\mathcal{T} = \{2020, 2021\}$ and $\mathcal{W}_{2020} = \{1, \dots, 53\}$ and $\mathcal{W}_{2021} = \{1, \dots, 52\}$ for weeks⁴ after January 1st, 2020. The parameters that describe the impact

⁴The year 2020 counted among the used calendar convention NEN 2772/ISO 8601 an (incomplete) 53rd week and 2021 an (incomplete) 0th week; a correction has been made in the datasets by merging the two.

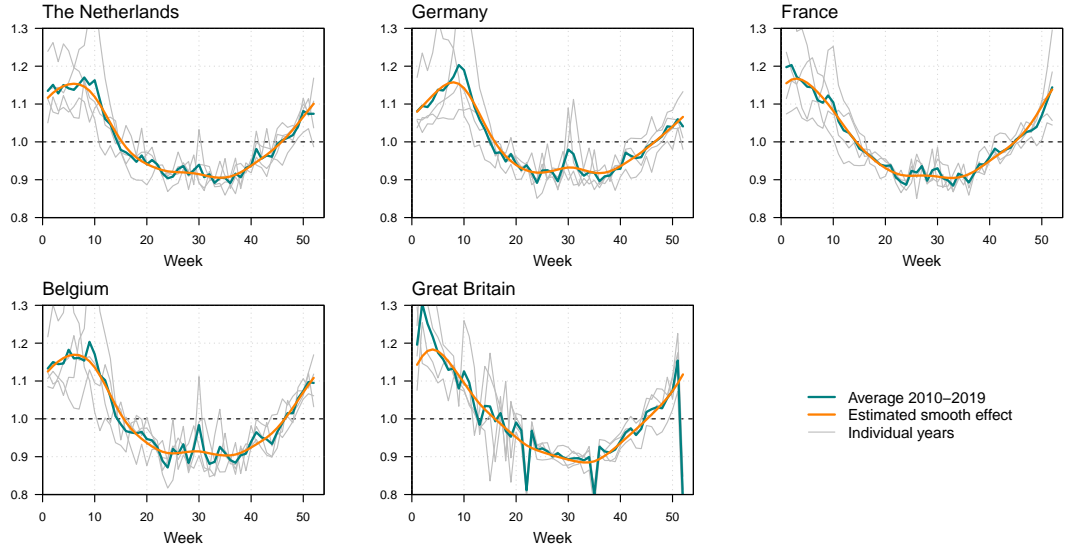


Figure 2: Observed fraction of annual mortality per week in the years 2010-2019 for the five countries considered, and estimated seasonal effect, aggregated across ages and sexes.

of the pandemic thus follow by determining for $g \in \{m, f\}$ and $c \in \mathcal{C}$:

$$\left(\widehat{\mathfrak{B}}_{\chi}^{c,g}, \widehat{\mathfrak{K}}_{t,w}^{c,g}\right) = \arg \min_{(\mathfrak{B}_{\chi}, \mathfrak{K}_{t,w})} \sum_{\chi \in \mathcal{X}} \sum_{t \in \mathcal{T}} \sum_{w \in \mathcal{W}_t} \left(D_{\chi,t,w}^{c,g} \mathfrak{B}_{\chi} \mathfrak{K}_{t,w} - D_{\chi,t,w}^{\text{pred},g,c} \exp(\mathfrak{B}_{\chi} \mathfrak{K}_{t,w}) \right), \quad (7)$$

where we define

$$D_{\chi,t,w}^{\text{pred},c,g} = E_{\chi,t,w}^{c,g} \mu_{\chi,t}^{c,g} \phi_{\chi,w}^{c,g} \quad (8)$$

for the expected number of deaths in a certain week based on given exposures during the week and the baseline mortality calibration, while possibly taking into account a seasonal effect.

Estimation of weekly deaths and exposures by age. The weekly exposures during the pandemic in (8), i.e., the values of $E_{\chi,t,w}^{c,g}$, cannot be based on observed data so we must generate estimates based on the population data that we have at our disposal.⁵ Various population statistics are available in the Eurostat database. Unfortunately, only yearly population estimates are given, and the most recent population estimate at the time that this paper was written refers to the measurements on 1 January 2020 for $c \in \{\text{NLD}, \text{BEL}, \text{DEU}, \text{FRA}\}$ and 1 January 2019 for $c = \text{GBR}$.

We determine the population on January 1 of year $t+1$ as $P_{x,t+1} = P_{x-1,t} - C_{x,t}$, where $C_{x,t}$ denotes the number of people that died in year t that would have had age x at 31 December of that year; these data can be obtained from the Eurostat database⁶. Some adjustments are necessary to estimate weekly exposures. We combine the population sizes from Eurostat with available weekly mortality observations per age-group, $D_{\chi,t,w}^{c,g}$, from the *Short Term Mortality Fluctuations* dataset, which works with age-groups of five years, starting from $\{0, 1, 2, 3, 4\}$ until

⁵To generate estimates of weekly exposures, we will make the assumption that each year consists of 52 weeks, the month February consists of 28 days, and the month December consists of 30 days. This results in a year of 364 days, which equals exactly 52 weeks. This approach has some clear drawbacks. For example, week 1 does not start on 1 January in each year, and December obviously has 31 days. However, this approach does take into account the development in the population estimates through the year. In case a year has 53 weeks, we assume that the exposure in week 53 is the same as it is in week 52 of that year. The exposure for a certain week is estimated as the average population during that week multiplied by $7/365$, so in the sequel we can focus on the determination of the population numbers per day during 2020 and 2021.

⁶We used a backtest over the years 2015 up to, and including, 2019 to investigate whether this is a good method to predict the change in population figures for an individual age over the course of a year, and found satisfactory results. Including migration information improves the accuracy of this method, but weekly migration information is not available and is therefore not used when estimating the weekly population sizes.

$\{85, 86, 87, 88, 89\}$, and an open-ended final age-group $\{90, 91, 92, \dots\}$. We transform the values $D_{\chi,t,w}^{c,g}$ for age-groups χ to values $D_{x,t,w}^{c,g}$ for individual ages x in $t = 2020$ and $t = 2021$ by assuming that the proportional distribution of deaths over ages in an age-group in a certain week w equals the historical average of that distribution for age-group over the years 2015 up to (and including) 2019:

$$D_{x,t,w}^{c,g} = D_{\chi,t,w}^{c,g} \cdot \frac{\sum_{t=2015}^{2019} D_{x,t}^{c,g}}{\sum_{x \in \chi} \sum_{t=2015}^{2019} D_{x,t}^{c,g}}. \quad (9)$$

Define $C_{x,t,w}^{c,g}$ as the number of people that died during week w in year t that would have had age x at 31 December of year t . We construct $C_{x,t,w}^{c,g}$ using the approximation⁷:

$$C_{x,t,w}^{c,g} = \left(1 - \frac{w}{w_t}\right) \cdot D_{x-1,t,w}^{c,g} + \frac{w}{w_t} \cdot D_{x,t,w}^{c,g},$$

where w_t represents the number of weeks in year t , so $w_{2019} = 52$, $w_{2020} = 53$, and $w_{2021} = 52$. We could estimate the population size at the first day of week $w + 1$ using

$$P_{x,t,w+1}^{c,g} = P_{x,t,w}^{c,g} - \left(\frac{w}{w_t} \cdot C_{x+1,t,w}^{c,g} + \left(1 - \frac{w}{w_t}\right) \cdot C_{x,t,w}^{c,g}\right).$$

However, we then ignore the fact that people may have their birthday during a week, and we therefore replace the above formula by

$$P_{x,t,w+1}^{c,g} = \left(1 - \frac{w}{w_t}\right) \left(P_{x,t,1}^{c,g} - \sum_{i \leq w} C_{x+1,t,i}^{c,g}\right) + \frac{w}{w_t} \left(P_{x-1,t,1}^{c,g} - \sum_{i \leq w} C_{x,t,i}^{c,g}\right), \quad (10)$$

which accounts for this effect under the assumption that births are distributed uniformly during the year. We initialize this procedure using the last known value of $P_{x,t}$ and we apply this formula for $w = 0, \dots, w_t$, such that $P_{x,t,1} = P_{x,t}$ and $P_{x,t,w_t+1} = P_{x,t+1,1} \approx P_{x,t+1}$.

Once we have a value for the population at the beginning and the end of the week, we can take the average to find the average population during the week, and we then have the required exposure $E_{x,t,w}^{c,g}$ for that week after we have multiplied the result by $\frac{7}{365}$.

For the Netherlands, we have more granular data at our disposal. Population sizes are available for the first day of each month until 1 January 2022. For the other days in the month, we determine the population sizes through linear interpolation. We do not need to project population sizes as in (10), since linear interpolation over monthly estimates is more accurate than projecting monthly estimates over a two-year period. The weekly exposures are then determined by multiplying the average population during the week by $\frac{7}{365}$. The mortality observations $D_{x,t,w}^{c,g}$ are available for individual ages and can be used directly.

3 Empirical Results

In this section, we present the empirical results when the model introduced in the previous section is applied to actual data. First, we analyze the impact of COVID-19 on mortality observations without using exposure information. This approach can be applied relatively easily, since during pandemics mortality observations may quickly become available, whereas exposure observations are often estimated and published on an annual basis only. We calibrate the COVID-19 model to Dutch mortality data for which we have mortality data available for individual ages. We use this granular dataset to investigate different ways to include a seasonal effect, analyze which ages are most affected by COVID-19, and examine the importance of having granular data. Finally, we compare the impact of COVID-19 on the level of mortality over time for the collection of five countries: NLD, DEU, FRA, BEL, and GBR.

⁷This approximation is based on the assumption that at all times mortality is uniformly spread over all people in a Lexis-parallellogram who have the same age at the end of the year, and uniformly spread during the year over all people with a common (rounded) age.

3.1 Analysis based on Dutch death counts only

In this section we show, as comparison, the impact of COVID-19 using only the number of deaths during the weeks of 2020 and 2021. This avoids the use of exposures which may be hard to estimate. We follow [18] in which the Compositional Data (CoDa) analysis is introduced, referring to [2], as an alternative to the Lee-Carter way of modeling mortality. In [18], the CoDa analysis is presented in a number of steps which make use of so-called CoDa operators. These operators are summarized in an appendix of [3]. We follow this approach and apply it to Dutch mortality data for the years 2010 to 2021 and ages 0 to 98. The CoDa analysis cannot deal with zero observations. Since the weekly death counts by age contain many zero observations, we increase all death counts by 1. This adjustment has negligible impact on the shape of the distribution of the deaths over the ages and over the weeks.

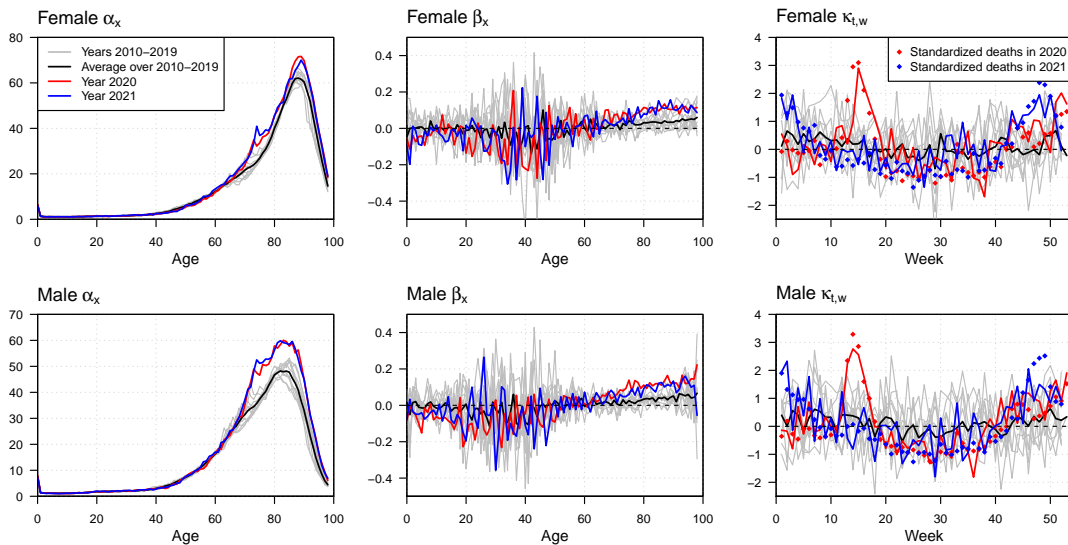


Figure 3: Results from CoDa analysis using Dutch weekly number of deaths. The gray lines show the estimated CoDa parameters for the years 2010 to 2019 and the black line represents the average over these estimates. The estimated CoDa parameters for the years 2020 and 2021 are shown in red and blue respectively. For the years 2020 and 2021, the graphs on the right also show the standardized weekly death counts (normalized using the mean and variance in that year).

Figure 3 shows the CoDa-parameter estimates. In the CoDa analysis applied to mortality data, the α_x parameters show the (average) age distribution of the death counts.⁸ These age distributions in 2020 and 2021 are close to the average over the years 2010 to 2019 for lower ages. For higher ages, however, the age distributions in 2020 and 2021 are higher than the average over the years 2010-2019 and in particular also the age distribution in 2019. This confirms that COVID-19 mainly increased the number of deaths of the elderly. The age distribution of 2021 is close to that of 2020, although there seems to be a slight shift to younger ages. In particular, there seems to be a small increase for the ages between 50 and 70.

The $\kappa_{t,w}$ process in the CoDa analysis shows shifts in the age distribution over the weeks. By construction, both the κ_w process (summed over time) and the β_x parameters (summed over ages) add up to zero. A positive value of κ_w implies a shift in the age distribution of number of deaths from the ages with a negative β_x to the ages with a positive β_x . The κ_w process shows a clear peak between weeks 10 and 20 of the year 2020, the first wave of the COVID-19 pandemic in the Netherlands. It also shows increasing positive values near the end of 2020, followed by positive values at the start of 2021, the second wave of COVID-19. Finally, near the end of 2021 the κ_w process shows the third wave of COVID-19. For comparison we also

⁸This age distribution shows the absolute number of deaths per age, aggregating to the total number of deaths per week.

show the standardized weekly death counts, i.e., the death counts normalized using the mean and variance in that year.

These peaks imply a shift in the age distribution from the young to the old, since the β_x s of the older people (above age 60) are positive, while the β_x values of the younger ages are mostly negative (although there are some exceptions as far as the younger ages are concerned). However, the β_x s for the older males are substantially higher in 2020 than in 2021. The shifts in the age distribution (corresponding to an increase in the κ_w process) of the males in 2021 are far less dramatic than the shifts in 2020. For females, the β_x coefficients of 2021 are close to those in 2020 for ages above 80, but for ages between 60 and 80 the β_x values in 2021 are lower than those in 2020. Thus, the shifts in the age distribution for females in 2021 is particularly in the direction of the very old, above age 80.

3.2 Calibration results based on a Dutch dataset with high granularity

In this section we continue our analysis using the Dutch mortality data as obtained from Statistics Netherlands. First, we investigate the two different approaches for incorporating the seasonal effect, then we compare the results when using different age ranges for calibration, and finally we analyze the importance of using granular data.

Incorporation of seasonal effect. In Section 2.3 we described two approaches for including the seasonal effect:

Method 1 *Time series for seasonal effect plus COVID-19*: the seasonal effect $\phi_{x,t}^{c:g}$ in (1) is set equal to 1, which results in any seasonal effect being captured by the COVID-19 term $\mathfrak{R}_{t,w}^{c:g}$.

Method 2 *Time series for COVID-19*: the seasonal effect $\phi_{x,t}^{c:g}$ is set equal to the estimated smooth seasonal effect as illustrated in Figure 2, which results in the COVID-19 term $\mathfrak{R}_{t,w}^{c:g}$ reflecting the impact of COVID-19 corrected for this seasonal effect.

Figure 4 shows the calibrated parameters. We observe that the differences in the COVID-19 age effects, as shown in the top row, are hardly visible. The COVID-19 age effect is erratic and close to zero for both females and males up to age 60. This indicates that for those ages there was hardly any impact on the level of mortality due to COVID-19 in the years 2020 and 2021. For higher ages, the age effect for females increases from age 60 to 70 and remains constant for ages 70 to 98, and for males the age effect increases steadily from age 60 to 98.

The COVID-19 week effects (bottom row) show substantial differences between the two methods. In the Netherlands, the corona virus was first identified in February 2020, which means that in the first four weeks of 2020 we would expect the COVID-19 week effect to be close to zero. The week effect represented by the dashed lines clearly starts above 0 for males and females. Further, in the summer of 2020 (around week 26) the dashed lines for the week effect are below zero, indicating mortality levels were below expectation if seasonality was not taken into account. The COVID-19 week effect of Method 2 (the solid line) starts close to zero and remains above zero for nearly all weeks included in the dataset; Method 2 therefore seems to capture only the impact of COVID-19 on observed mortality.

The observed weekly mortality death counts as shown in Figure 5 are very volatile for the ages 45 and 55 and do not exhibit the peaks as observed in the COVID-19 week effects in Figure 4. The fitted death counts are close to the seasonally adjusted death counts, and the impact of COVID-19 on mortality at these ages is therefore negligible. For ages 65 and 85, the seasonally adjusted expected deaths clearly exhibit a seasonal pattern, but this pattern is not sufficient to capture the peaks that coincide with the COVID-19 waves. The fitted deaths using Method 1 and Method 2 follow the wave pattern in the observed deaths more closely.

We would like to be able to distinguish between effects induced by a typical seasonal effect and effects due to a pandemic. Method 2, in which we correct for historically observed seasonal effects, allows us to assess the impact of only the pandemic on the level of mortality. Further, Method 2 is more flexible than Method 1, since the seasonal effect is not enforced to affect mortality the same way as the pandemic. Therefore, in the remainder of the analyses we will only use Method 2 when estimating the impact of COVID-19 on the level of mortality.

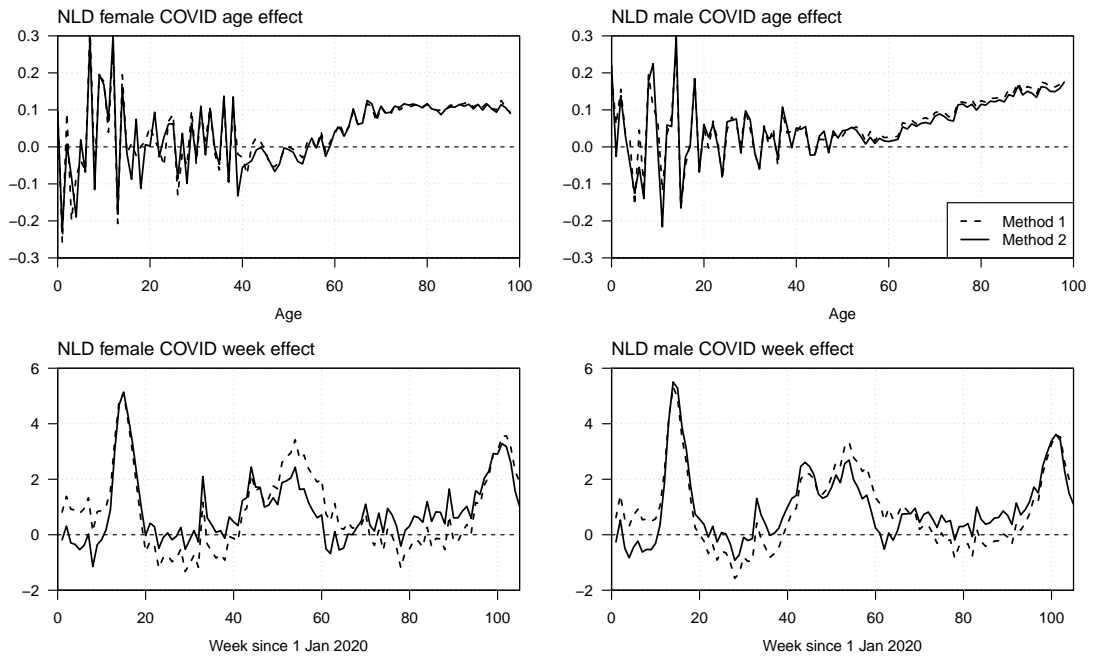


Figure 4: COVID-19 parameters estimated with and without predetermined seasonal effect (solid respectively dashed line) for females and males using Dutch data.

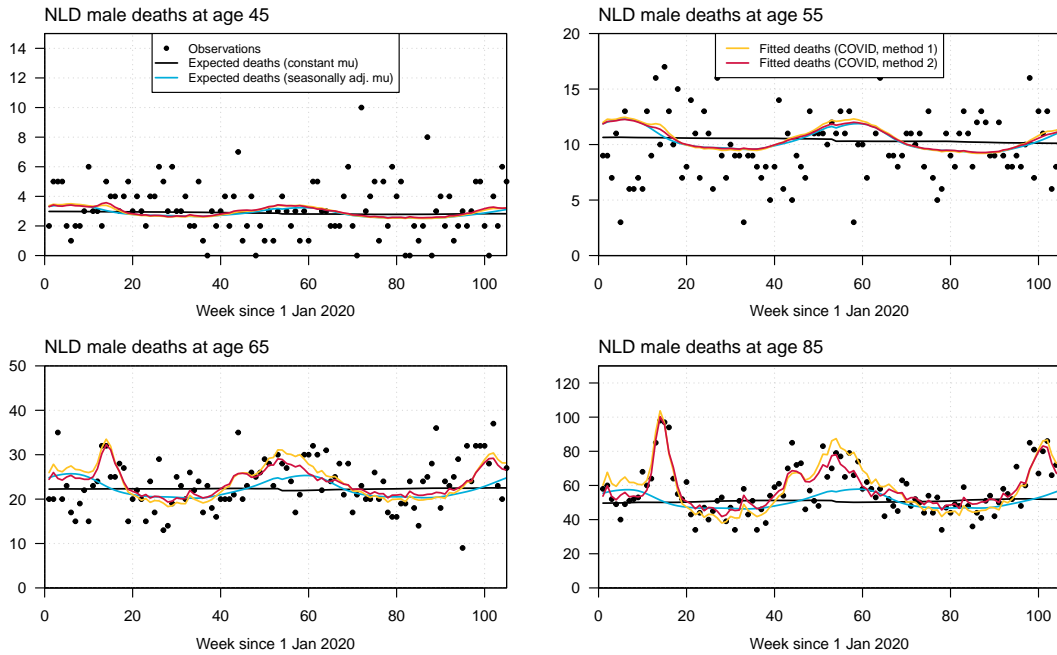


Figure 5: The black dots represent observed weekly deaths for Dutch males at ages 45, 55, 65 and 85 during the years 2020 and 2021. The black line shows the expected deaths assuming a constant force of mortality, and the blue line represents expected death counts taking into account the historical seasonal effect. The yellow and red lines show the fitted deaths using Method 1 respectively Method 2.

Selection of ages. The COVID-19 age effects for ages below 40 in Figure 4 are volatile and seem to be centered around zero. In Appendix A we show estimates using ages 0-98 and ages 40-98. From the parameter estimates in Figure 12, we conclude that the COVID-19 week parameters and age parameters for higher ages are hardly affected when including information from

younger ages. For interpretation of parameters and reliability of estimates and projections, it is desirable that only structural effects are analyzed; noise due to low exposures should preferably be excluded. Since the age effects at the younger ages do not seem to capture a systematic COVID-19 effect, we exclude data for ages below 40 from calibration.

Importance of granular data. For the Netherlands we have weekly mortality data available for both genders and individual ages. The *Short Term Mortality Fluctuations database* (STMF) contains weekly mortality data for many countries, but only for specific age groups. For some countries the age groups are small (age groups of five years), whereas for other countries there are only five age groups.

In this section, we use the dataset from Statistics Netherlands to investigate the importance of having granular data in analyzing the impact of COVID-19 on mortality. We consider three levels of granularity:

- **Level 1**, individual ages with $\chi \in \{0, 1, 2, \dots, 99\}$; this is the most granular type of data;
- **Level 2**, age groups of five years with $\chi \in \{\{0, \dots, 4\}, \{5, \dots, 9\}, \dots, \{90, \dots, 94\}, \{95, \dots\}\}$;
- **Level 3**, five age groups with $\chi \in \{\{0, \dots, 14\}, \{15, \dots, 64\}, \{65, \dots, 74\}, \{75, \dots, 84\}, \{85, \dots\}\}$; this is the least granular type of data.

We use the original dataset with observations by individual ages (Level 1) to artificially construct the datasets based on Level 2 and Level 3 granularity.

Figure 6 shows the parameter estimates when calibrating the COVID-19 model to the three datasets. The top row shows the estimated COVID-19 age effects, and we observe clear differences between the estimates from the datasets with different levels of granularity. The age effects are volatile for all levels of granularity, but the Level 2 and Level 3 estimates show remarkable peaks and dips at the ages 65-75. In the year 2020, these ages correspond with the years of birth 1945-1955 (with a clear peak in 1946), which is the so-called baby boom generation after the second World War. The baby boom generation appears as a cohort effect in the exposures, which in turn results in cohort effects in the observed deaths. For most age groups (either Level 2 or Level 3), the distribution of deaths over the ages is relatively stable over time, except for this baby boom generation. This cohort effect makes the allocation of deaths to individual ages increasingly inaccurate if larger age groups and more historical years are used when applying Equation (9).

In contrast, the COVID-19 week effects estimated using the datasets of different granularity look remarkably similar. This is as expected since the observed deaths are not relocated over different weeks, only over different ages. The impact of COVID-19 on total mortality will therefore be similar, which results in stable estimates of the COVID-19 week effect, regardless of the level of granularity of the dataset.

From this analysis we conclude that datasets with deaths by age group can be used to obtain a first impression of the impact of a pandemic over time. However, if the impact on individual ages must be assessed, one needs to use death counts for individual ages.

3.3 Calibration results based on the STMF dataset

In this section, we calibrate the COVID-19 model to the countries c in the set given by $\mathcal{C} = \{\text{NLD}, \text{DEU}, \text{FRA}, \text{BEL}, \text{GBR}\}$ using mortality data as obtained from STMF. The deaths by age group are allocated to individual ages using Equation (9), and we include data for the ages 40 to 95. We investigate to what extent similarities and differences between countries can be observed.

Figure 7 shows the estimated COVID-19 parameters.⁹ The COVID-19 age effect is close to zero at lower ages for most countries. The exception to this is Great Britain, where the COVID-19 age effect is substantial for all ages included. The volatile behavior between ages 65 and 80 is probably the result of cohort effects in the death counts, as described in the previous section.

⁹Results obtained using ages 0-98 are available in Appendix B.

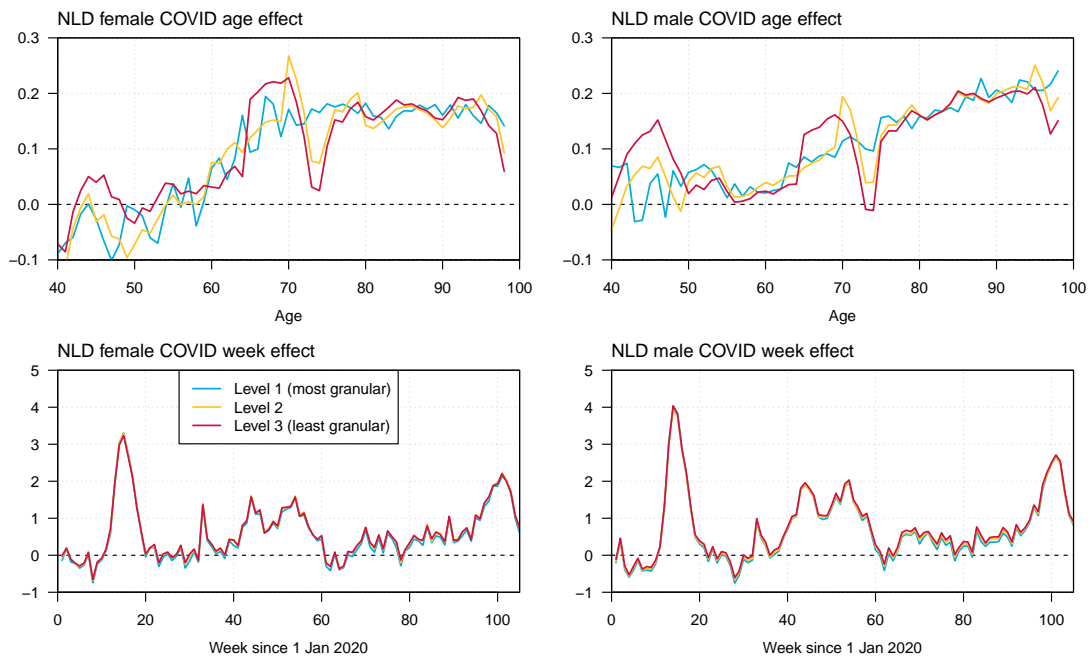


Figure 6: COVID-19 parameter estimates using datasets with different levels of granularity.

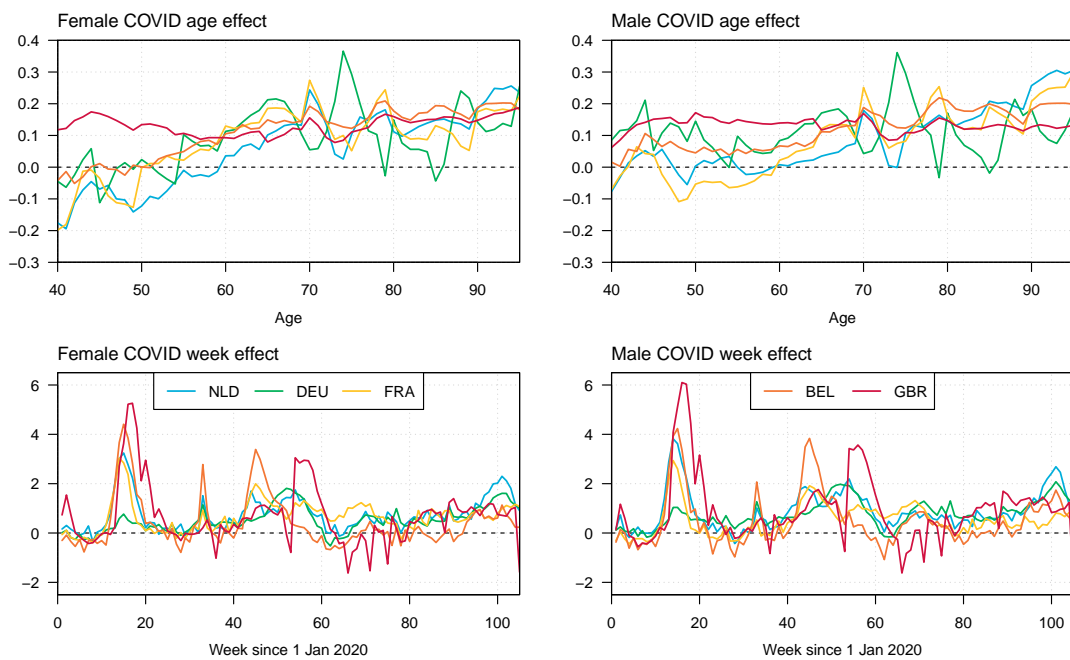


Figure 7: Estimated COVID-19 parameters for various countries using ages 40 to 95.

We observe that the different countries exhibit similar COVID-19 week effects, though there are a few notable differences. Germany experienced a minor first COVID-19 wave compared to the other four countries, whereas later COVID-19 waves were similar to those in other countries. Around week 33, Belgium, the Netherlands and Germany experienced a temporary peak, which may be the result of temporarily relaxing COVID-related restrictions.

Between week 40 and 60 (winter season 2020-2021) the second COVID-19 wave hit Western Europe, but we observe substantial differences between countries. From week 50 onward,

COVID-19 vaccines became available in various countries, though the availability and timing of the vaccine shots for people of different ages differed between countries. Belgium had a high peak at around week 45 after which excess mortality soon disappeared. The Netherlands, Germany and France all experienced lower peaks, but in these countries it also took a few more weeks before excess mortality had vanished. This decrease in excess mortality may have been the result of vaccines being applied to parts of the population but also due to new COVID-19 restrictions. Finally, Great Britain experienced a high peak which was similar to the one in Belgium, but a few weeks later, and measures were taken resulting in excess mortality decreasing rapidly. Great Britain was one of the countries in which vaccines were provided to the population earlier and faster.

It will always remain challenging to assess the effectiveness of policy decisions during a pandemic. But using the best possible data and using improved estimation methods can help to improve such assessments.

4 Forecasting mortality adjusted for COVID-19

In the previous section we have analyzed the impact of COVID-19 on the level of mortality in the years 2020 and 2021. For insurance companies and pension funds it is particularly important to assess the impact of COVID-19 on mortality rates. At the time this paper was written, it was too early to predict whether long term mortality improvement rates should be adjusted upward, downward, or not at all. In this section, we therefore introduce a general framework in which the impact of COVID-19 on the level of mortality in the years 2020 and 2021 can be used to generate scenarios for future mortality.

Transforming weekly effects to annual effects. The impact of COVID-19 was calibrated using weekly mortality observations. From these calibrations, we obtained weekly parameter estimates \mathfrak{B}_x and $\mathfrak{K}_{t,w}$. To generate scenarios for future mortality, we first need to transform these week effects into annual effects \mathfrak{V}_x and \mathfrak{X}_t . We temporarily impose that $\sum_{x \in \mathcal{X}} \mathfrak{V}_x = 1$ such that analytical results can be used for the transformation. Afterwards, we again apply the restriction $\|\mathfrak{V}\| = 1$. To transform week effects into annual effects, we make the two one-year survival probabilities for the years 2020 and 2021 equal to the product of the weekly survival probabilities in those years:

$$\exp(-\mu_{x,t} \cdot \exp[\mathfrak{V}_x \mathfrak{X}_t]) = \prod_{w=1}^{w_t} \exp\left(-\frac{1}{w_t} \cdot \mu_{x,t} \cdot \phi_{x,w} \cdot \exp[\mathfrak{B}_x \mathfrak{K}_{t,w}]\right), \quad (11)$$

for $t = 2020$ and $t = 2021$, and for all $x \in \mathcal{X}$. After taking the natural logarithm on both sides, dividing by $\mu_{x,t}$, taking the natural logarithm again, and summing over all ages $x \in \mathcal{X}$, we find

$$\mathfrak{X}_t = \sum_{x \in \mathcal{X}} \ln \left\{ \frac{1}{w_t} \sum_{w=1}^{w_t} \phi_{x,w} \cdot \exp[\mathfrak{B}_x \mathfrak{K}_{t,w}] \right\}.$$

Next, we determine \mathfrak{V}_x by making survival over both the years 2020 and 2021 equal to surviving over all weeks in those years (which follows directly from (11)):

$$\prod_{t=2020}^{2021} \exp(-\mu_{x,t} \cdot \exp[\mathfrak{V}_x \mathfrak{X}_t]) = \prod_{t=2020}^{2021} \prod_{w=1}^{w_t} \exp\left(-\frac{1}{w_t} \cdot \mu_{x,t} \cdot \phi_{x,w} \cdot \exp[\mathfrak{B}_x \mathfrak{K}_{t,w}]\right).$$

Rewriting this equation results in

$$\sum_{t=2020}^{2021} \mu_{x,t} \sum_{w=1}^{w_t} \frac{1}{w_t} (\exp[\mathfrak{V}_x \mathfrak{X}_t] - \phi_{x,w} \cdot \exp[\mathfrak{B}_x \mathfrak{K}_{t,w}]) = 0,$$

and this non-linear equation in \mathfrak{V}_x can be solved numerically for each age $x \in \mathcal{X}$ separately. Finally, we again renormalize the parameters \mathfrak{V}_x and \mathfrak{X}_t such that $\|\mathfrak{V}\| = 1$.

Figure 8 shows the weekly and annual effects calibrated using the Dutch mortality data from Statistics Netherlands. The weekly and annual COVID-19 age effects are nearly identical, and the COVID-19 year effects are close to the averages of the week effects in that year.

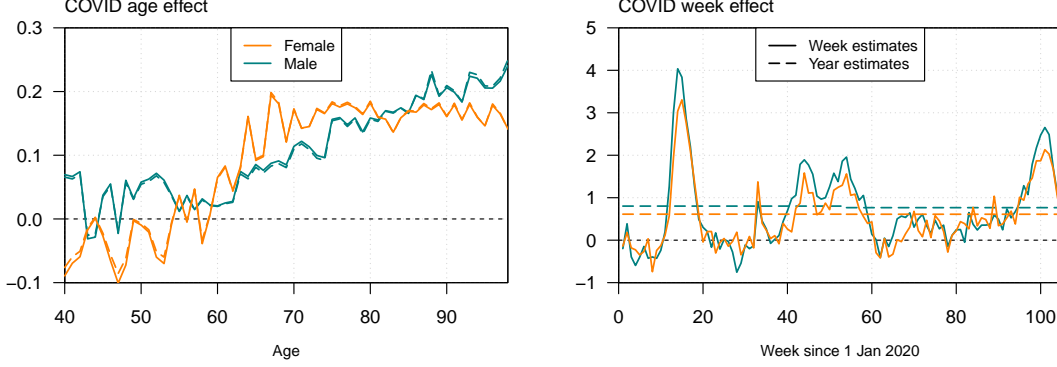


Figure 8: COVID-19 parameters from calibration on weekly data (solid line) and annualized COVID-19 parameters (dashed line).

Future scenarios for the COVID-19 period effect. At the end of the observed period, the week effect as shown in Figure 8 is not far from zero. For prediction purposes, one might be tempted to choose \mathfrak{X}_t equal to 0 for $t > 2021$. However, the week effect also shows that after periods of (almost) no excess mortality, new periods of excess mortality may occur. We therefore investigate how mortality rates could develop under varying assumptions for the future impact of COVID-19 on mortality.

We choose scenarios that can be written as:

$$\begin{aligned} \ln \mu_{x,2021+h} &= \ln \mu_{x,2021+h}^{\text{pre-covid}} + \mathfrak{V}_x \mathfrak{X}_{2021+h} \\ \mathfrak{X}_{2021+h} &= \mathfrak{X}_{\text{start}} \eta^h + (1 - \eta^h) \mathfrak{X}_{\infty}, \end{aligned}$$

for $0 \leq \eta \leq 1$. The value of the COVID-19 period effect in the limit, \mathfrak{X}_{∞} , is chosen to be a multiple of \mathfrak{X}_{2021} where the multiplicative factor is scenario-dependent. The parameter η defines how fast convergence to this limit value takes place. This generic approach for generating COVID-19 scenarios is similar to that of [27] who estimate similar COVID-19 effects using a penalized quasi-likelihood approach. We define the following scenarios:

1. **Completely incidental:** $\mathfrak{X}_{\text{start}} = \mathfrak{X}_{\infty} = 0$.
The mortality forecast is completely based on the pre-COVID forecast.
2. **Completely structural:** $\mathfrak{X}_{\text{start}} = \mathfrak{X}_{\infty} = \mathfrak{X}_{2021}$.
It is assumed that COVID-19 is a new cause of death that will remain permanently with an impact that does not change over time.
3. **Decreasing impact:** $0 < \eta < 1$, $\mathfrak{X}_{\text{start}} = \mathfrak{X}_{2021}$, $\mathfrak{X}_{\infty} = 0$.
It is assumed that the effect of COVID-19 converges to zero after a certain period, for example because herd immunity is (almost) reached in a population.
4. **Growing impact:** $0 < \eta < 1$, $\mathfrak{X}_{\text{start}} = \mathfrak{X}_{2021}$, $\mathfrak{X}_{\infty} = 1.25\mathfrak{X}_{2021}$.
The impact grows, for example when new variants appear which are not impacted by current and new vaccines, or when other mitigating policies imposed by the government become less effective.
5. **New normal:** $0 < \eta < 1$, $\mathfrak{X}_{\text{start}} = \mathfrak{X}_{2021}$, $\mathfrak{X}_{\infty} = 0.25\mathfrak{X}_{2021}$.
The impact decreases to a constant but does not completely vanish, which results in a permanent effect that can be compared to other causes of death such as the flu.

6. **Increased resilience:** $0 < \eta < 1$, $\mathfrak{X}_{\text{start}} = \mathfrak{X}_{2021}$, $\mathfrak{X}_{\infty} = -0.25\mathfrak{X}_{2021}$.

The impact gradually converges to a value below zero, for example because the population that survived the pandemic is stronger than the population before the pandemic.

At the time of writing this paper, no information is available to make an informed decision for the value of η . We choose $\eta = 0.5$ for illustration purposes. Given the other choices made for specifying the COVID-19 scenarios, the predicted period effects for Dutch males are shown in Figure 9; the figure for females looks similar in case the same choices are made.

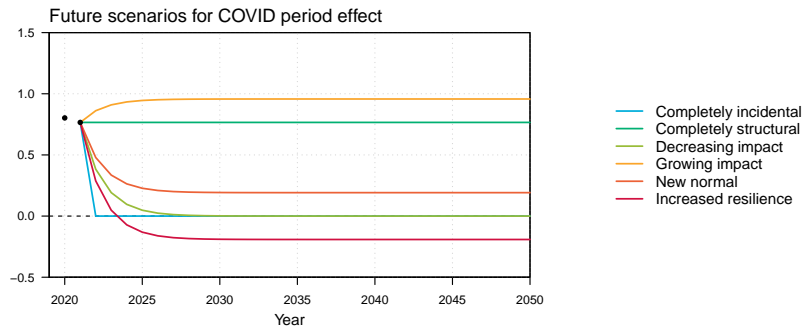


Figure 9: Projected COVID-19 period effects for Dutch males under various assumptions for future development of impact of COVID-19 on mortality.

We also need to specify how the COVID-19 age effect is defined for ages $x \notin \mathcal{X}$. Define x_{\min} and x_{\max} as the lowest respectively highest age included in the set of ages used for calibration, \mathcal{X} (i.e. $x_{\min} = 40$ and $x_{\max} = 95$ in the previous section). Figure 12 suggests that for lower ages the impact of COVID-19 on the level of mortality was negligible, and therefore these ages are excluded from calibration. In line with that approach, we define $\mathfrak{V}_x = 0$ for $x < x_{\min}$. For higher ages, we observe that the estimated parameter \mathfrak{V}_x seems relatively stable at higher ages. Based on that observation, we define $\mathfrak{V}_x = \mathfrak{V}_{x_{\max}}$ for $x > x_{\max}$.

Mortality predictions under various COVID-19 scenarios. For the various scenarios illustrated in Figure 9, we have constructed the corresponding forecasts of mortality rates. Figure 10 shows the mortality forecasts for Dutch males at ages 55, 65 and 85. Though at age 55 mortality rates in 2020 and 2021 were not markedly different from the pre-COVID expectation, at age 65 the level of mortality was elevated and at age 85 it was substantially higher than expected. The pattern of the projected COVID-19 period effects from Figure 9 are clearly visible in the mortality forecasts. When constructing mortality rates, the COVID-19 period effects are multiplied with the COVID-19 age effects from Figure 8 which explains why the forecast of mortality rates at age 55 are hardly affected by the COVID-19 period effect while at age 85 mortality in all scenarios (except for **Completely incidental**) deviates greatly from the pre-COVID forecast. For the scenarios **Decreasing impact**, **New normal** and **Increased resilience**, the impact on mortality diminishes quickly after a few years (given our choice for η), but for **Completely structural** and **Growing impact** the impact remains.

The impact on predicted period life expectancies and cohort life expectancies is shown in the top row respectively bottom row of Figure 11. The period life expectancy at all ages is far below the pre-COVID expectation. Though the COVID-19 age effects are calibrated for the ages 40-95, the period life expectancy at age 0 is also affected, since mortality rates at all ages are used to compute this life expectancy. The gap between the most positive scenario (**Increased resilience**) and the most negative scenario (**Growing impact**) increases to approximately one year around 2040 for the ages shown. This impact remains constant further in the future once the COVID-19 period effect has converged to its limit value, and all remaining future developments are driven by the original Li-Lee model.

The pre-COVID forecast and the scenarios **Completely incidental** and **Decreasing impact** result in comparable projected cohort life expectancies. While in the predicted period life ex-




Figure 10: Projected mortality rates for Dutch males aged 55, 65 and 85 under various assumptions for future development of impact of COVID-19 on mortality.

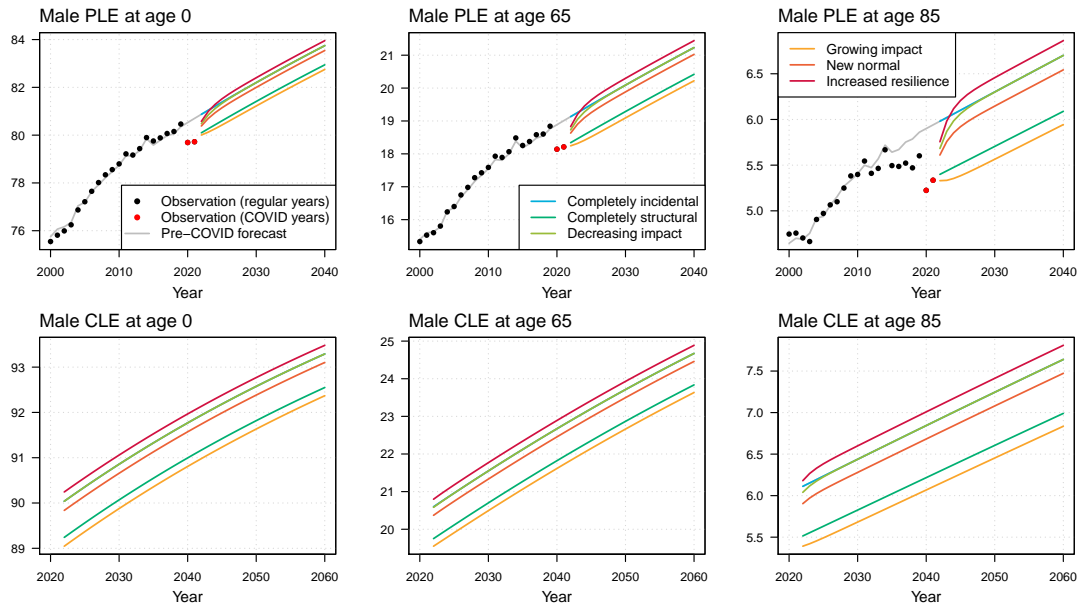


Figure 11: Projected period and cohort life expectancies for Dutch males at birth and at ages 65 and 85 under various assumptions for future development of impact of COVID-19 on mortality.

pectancies differences between scenarios grow over time, in predicted cohort life expectancies the difference between scenarios at the start of the projection is close to the difference in the limit.

5 Conclusions

In this paper, we introduce a model to quantify weekly deviations from expected levels of mortality during a pandemic. The model adds an extra layer to the two-layer Li-Lee model by including an additional seasonal effect to capture regular seasonal patterns and an additional age and week effect to measure deviations from pre-pandemic weekly mortality expectations.

We apply our model to data from Belgium, France, Germany and Great Britain using mortality data from the *Short Term Mortality Fluctuations* database. There are differences between countries in the extent to which mortality at different ages is affected by COVID-19 and how COVID-19 affected mortality through time. Yet, there are also clear similarities since in all countries periods of high excess mortality are followed by periods of lower or no excess mortality.

The application of the model requires the availability of exposures which often are not available on a weekly basis and these must therefore be approximated. Useful insights on the de-

velopment of mortality through a year can be obtained through a Composition Data (CoDa) analysis that can be performed using weekly death counts only.

Most sources of weekly mortality observations provide data by gender and by age groups. Our sensitivity analysis shows that such data can be used to accurately monitor the development of the pandemic through time, which is represented by the COVID-19 week effect. However, the COVID-19 age effects are inaccurate in case of cohort effects exist in population sizes, and use of weekly observations by individual ages is therefore recommended.

The future COVID-19 scenarios analyzed in this paper all assume convergence to pre-COVID long term improvement rates. At this stage, there is insufficient data and information available to determine whether, and if so how long term mortality improvement rates should be adjusted. This paper provides no solution for this problem, and this is a problem that is likely to challenge the imagination of demographers and actuaries for years to come.

Acknowledgment. Certain parts of the approach proposed in this paper have been used in a model to generate a prognosis for future survival probabilities that was recently developed for the Royal Dutch Actuarial Association. The authors gratefully acknowledge all comments from members of the Association's Committee and Working Group on Mortality Research: Wies de Boer, Friso Cuijpers, Corné van Iersel, Marieke Klein, Hans de Mik, Erica Slagter, Janinke Tol, Erik Tornij, Raymond Waucomont, Wouter van Wel, Menno van Wijk, Marco van der Winden and Kim Wittekoek. We also thank Statistics Netherlands, and in particular Lenny Stoeldraijer, for helping us to obtain the required datasets.

References

- [1] D. Adam. The pandemic's true death toll: millions more than official counts. *Nature*, 601:312 – 315, 2022.
- [2] J. Aitchison. *The Statistical Analysis of Compositional Data*. London: Chapman and Hall, 1986.
- [3] M.-P. Bergeron-Boucher, V. Canudas-Romo, J. Oeppen, and J.W. Vaupel. Coherent forecasts of mortality with compositional data analysis. *Demographic Research*, 37:527–566, 2017.
- [4] N. Brouhns, M. Denuit, and J.K. Vermunt. A Poisson log-bilinear regression approach to the construction of projected lifetables. *Insurance: Mathematics and Economics*, 31(3):373 – 393, 2002.
- [5] A.J.G. Cairns, D. Blake, and K. Dowd. A two-factor model for stochastic mortality with parameter uncertainty: Theory and calibration. *Journal of Risk and Insurance*, 73(4):687 – 718, 2006.
- [6] A.J.G. Cairns, D. Blake, K. Dowd, G.D. Coughlan, D. Epstein, A. Ong, and I. Balevich. A quantitative comparison of stochastic mortality models using data from England and Wales and the United States. *North American Actuarial Journal*, 13(1):1 – 35, 2009.
- [7] H. Chen and S.H. Cox. Modeling mortality with jumps: Applications to mortality securitization. *Journal of Risk and Insurance*, 76(3):727–751, 2009.
- [8] A. Clark, J. Mark, and C. WarrenGash et al. Global, regional, and national estimates of the population at increased risk of severe COVID-19 due to underlying health conditions in 2020: a modelling study. *The Lancet Global Health*, 8:1003–1017, 2020.
- [9] J. Elliott, B. Bodinier, and M. Whitaker et al. COVID-19 mortality in the UK biobank cohort: revisiting and evaluating risk factors. *European Journal of Epidemiology*, 36:299–309, 2021.
- [10] M. Spencer Gold, D. Sehayek, S. Gabrielli, X. Zhang, C. McCusker, and M. Ben-Shoshan. COVID-19 and comorbidities: a systematic review and meta-analysis. *Postgraduate Medicine*, 132(8):749–755, 2020.
- [11] Koninklijk Actuarieel Genootschap. Projection table AG 2020, 2020. Available online at: http://www.ag-ai.nl/view.php?action=view&Pagina_Id=1007.
- [12] R. Lee and T. Miller. Evaluating the performance of the Lee-Carter method for forecasting mortality. *Demography*, 38(4):537–549, 2001.
- [13] R.D. Lee and L.R. Carter. Modeling and forecasting U.S. mortality. *Journal of the American Statistical Association*, 87(419):659–671, 1992.
- [14] N. Li and R.D. Lee. Coherent mortality forecasts for a group of populations: an extension of the Lee-Carter method. *Demography*, 42(3):575 – 594, 2005.
- [15] Y. Liu and J. S.-H. Li. The age pattern of transitory mortality jumps and its impact on the pricing of catastrophic mortality bonds. *Insurance: Mathematics and Economics*, 64:135 – 150, 2015.
- [16] S.J. McGurnaghan, A. Weir, and J. Bishop et al. Risks of and risk factors for COVID-19 disease in people with diabetes: a cohort study of the total population of Scotland. *The Lancet Diabetes & Endocrinology*, 9(2):82–93, 2021.
- [17] M. O'Driscoll, G. Ribeiro Dos Santos, and L. Wang et al. Age-specific mortality and immunity patterns of SARS-CoV-2. *Nature*, 590:140 – 145, 2021.

-
- [18] J. Oeppen. Coherent forecasting of multiple-decrement life tables: A test using Japanese cause of death data. Paper presented at the European Population Conference 2008, Barcelona, Spain, July 9–12, 2008.
- [19] University of California at Berkeley and Max Planck Institute for Demographic Research. Human mortality database. www.mortality.org or www.humanmortality.de.
- [20] R. Plat. On stochastic mortality modeling. *Insurance: Mathematics and Economics*, 45(3):393 – 404, 2009.
- [21] A.E. Renshaw and S. Haberman. Lee-Carter mortality forecasting with age-specific enhancement. *Insurance: Mathematics and Economics*, 33(2):255 – 272, 2003.
- [22] J. Robben, K. Antonio, and S. Devriendt. Assessing the impact of the COVID-19 shock on a stochastic multi-population mortality model. *Risks*, 10:26 –, 2022.
- [23] S. Schnürch, T. Kleinow, R. Korn, and A. Wagner. The impact of mortality shocks on modelling and insurance valuation as exemplified by COVID-19. *Annals of Actuarial Science*, page 1–29, 2022.
- [24] A.K. Singh, C.L. Gillies, R. Singh, A. Singh, Y. Chudasama, B. Coles, S. Seidu, F. Zaccardi, M.J. Davies, and K. Khunti. Prevalence of co-morbidities and their association with mortality in patients with COVID-19: A systematic review and meta-analysis. *Diabetes Obes Metab.*, 22:1915–1924, 2020.
- [25] T. Takahashi, M.K. Ellingson, and P. Wong et al. Sex differences in immune responses that underlie COVID-19 disease outcomes. *Nature*, 588:315–320, 2020.
- [26] E. Williamson, A.J. Walker, and K. Bhaskaran et al. OpenSAFELY: factors associated with COVID-19-related hospital death in the linked electronic health records of 17 million adult NHS patients. *medRxiv*, 2020.
- [27] R. Zhou and J.S.-H. Li. A multi-parameter-level model for simulating future mortality scenarios with COVID-alike effects. *Annals of Actuarial Science*, page 1–25, 2022.

A Selection of ages

In Figure 4 we observed that for ages below 40 the COVID-19 age effects are erratic. Since the age effects do not seem to capture a systematic COVID-19 effect at the younger ages, it might be better to exclude those ages from calibration of the model. The COVID-19 age effects are more stable from age 40 onward, though between ages 40 and 60 the estimated effect is close to zero.

For interpretation of parameters and reliability of outcomes it is desirable that only structural effects are analyzed; noise due to low exposures should preferably be excluded. Therefore, we compare parameter estimates if we use different age ranges for calibration of the COVID-19 model. Figure 12 shows the estimated parameters for females and males using ages 0-98 (dashed line) and ages 40-98 (solid line) for calibration, where the former are renormalized such that the norms over the age range 40-98 are equal in both cases. This allows a comparison of the parameter estimates.

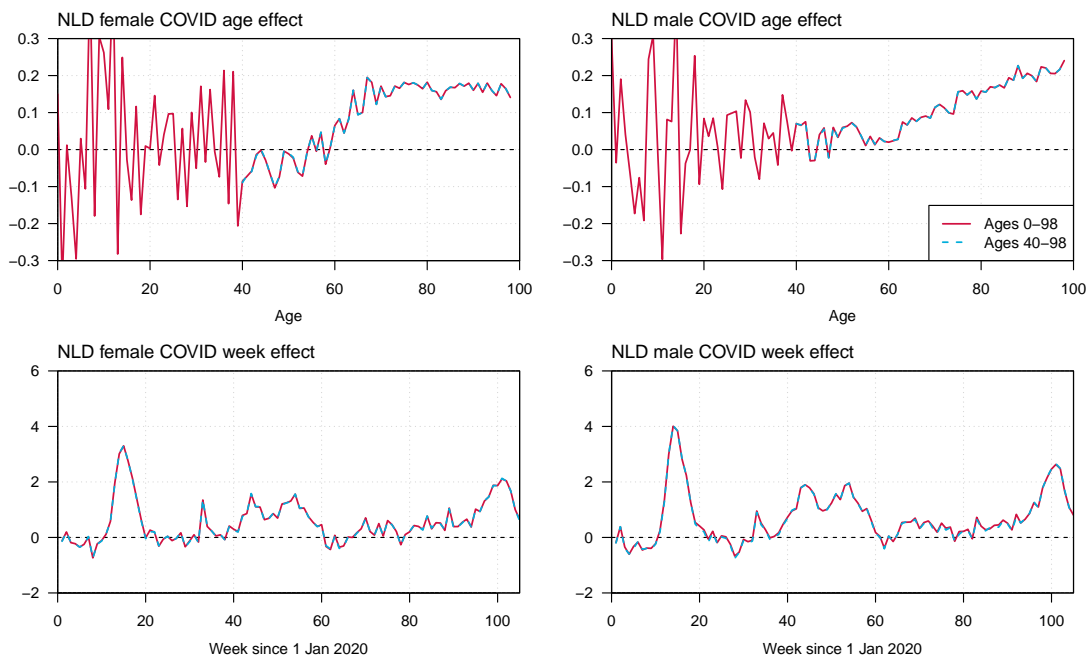


Figure 12: COVID-19 parameters estimated using ages 0-98 (dashed line) and ages 40-98 (solid line).

From the parameter estimates in Figure 12 we conclude that the COVID-19 week parameters and age parameters for higher ages are hardly affected when including information from younger ages. However, the COVID-19 age effects for ages below 40 years are erratic, and therefore we exclude these ages from further analyses.

B Additional calibration results based on STMF data

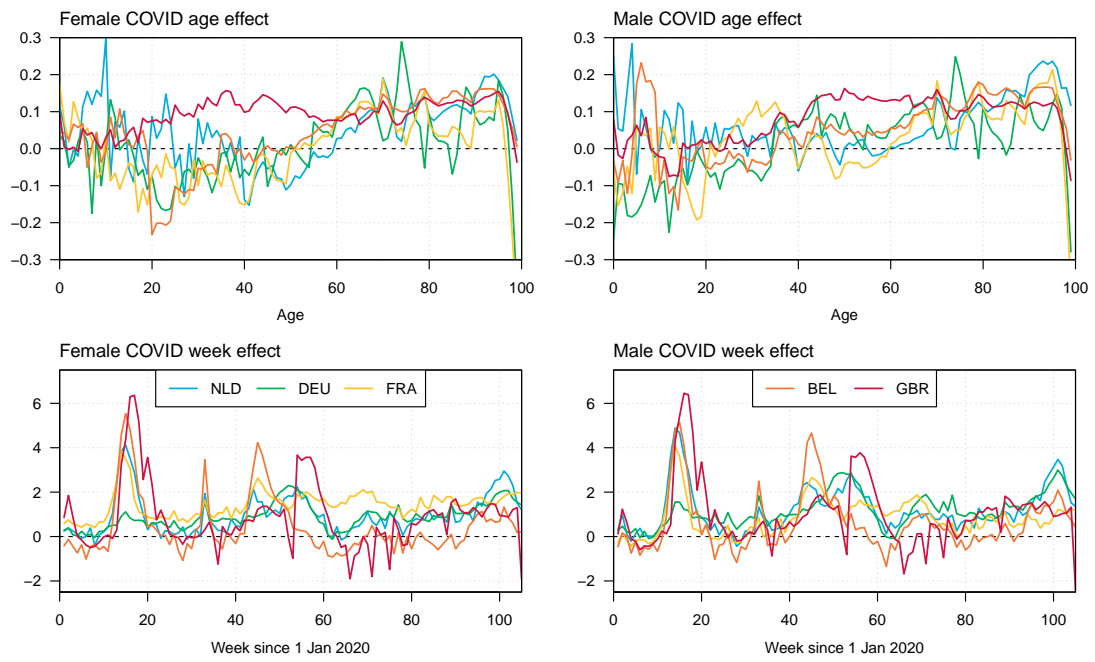


Figure 13: Estimated COVID-19 parameters for various countries using ages 0 up to and including 99.