Mortality Prediction by Time, Temperature and Frailty applied to Pandemic Excesses by Covid and Vaccines

André Redert, PhD
Independent researcher at Rodotti
andre@rodotti.nl
Netherlands, 20 July 2024

Abstract

Objective: Present a mortality model for more objective quantification and localization in time of excess mortality, and analyse the role of covid and vaccines in recent pandemic excesses.

Method: The model is just one equation, *M = W*atT*. Normal and excess mortality are explicitly defined in a generic, non-mechanistic way. *S*ource data is weekly aggregated absolute mortality and average outside temperature, both of which are objective, reliable and available worldwide.

Twelve parameters represent trends over decades, seasonal/weekly variability, health pressure accumulation (HPA) relating to delay between cause and death, frail pool dynamics (FPD) relating to lifetime lost/saved (LTL/LTS), and the natural variability (NV) of remaining random mortality fluctuations. Bayesian probabilistics and brute-force numerical optimization are used to fit the model. Prediction accuracy is explicitly defined and parametrized.

Pandemic excesses are examined by integrating determinants of positive covid tests and vaccinations. Measured parameters are covid Case Fatality Rate (CFR), Vaccine-dose Fatality Rate (VFR), Vaccine Effectivity (VE) against covid mortality, and LTL/LTS by covid and vaccination.

Experiments: Results with 10 EU countries (344M people) and age-stratified datasets in The Netherlands (NL) yield a model fit over 2000-2019 with accuracy ca. 1% of yearly mortality. Predictions for 2020-2023 have 79% variance reduction compared to the Dutch national baseline. Highest all-ages pandemic excess is found in Bulgaria (10σ) and Poland (12σ), while Germany (DE) and NL show no significant excess in any pandemic winteryear (July to June). For ages 0-65 in NL, excess starts in 2020/21, peaks in 2021 (5σ) and remains persistent. FPD predicts mortality deficits to match excesses within a year, but none are observed in any country or age group.

In NL, covid HPA time is found at 1.8 weeks (w) at 95% confidence interval (1.5-2.2w), matching known mean times between positive test and death. Covid FPD time is 23w (16-33w), ca. half of temperature's 44w (33-65w) which represents natural mortality. All-country average VFR is 0.13±0.1% death/dose, while VE starts at ca. 100% and wanes in months to a year, matching existing studies. VFR is 0.09% (0.06-0.15%) for NL and 0.22% (0.12-0.41%) for DE. Dose-to-death HPA time ranges from months to years, while FPD time ranges up to decades, an order above temperature's. Considerable loss of lifeyears is measured due to vaccination in all countries examined. In DE, LTL is 450-4600ky (kiloyear) and in NL 50-500ky. Age-stratified NL results are 13-94ky (ages 0-65), 16-65ky (ages 65-80) and 5-34ky (ages 80+).

Conclusions: Mortality predictions and excess measurements by the proposed model are considerably more accurate than those of the national baseline. Despite low reliability of pandemic determinants, many results are significant and in line with other recent findings. The pandemic health pressure is found to be fully persisting throughout 2020-2023, while covid targets much frailer people than natural causes do, and vaccine fatal events target younger and healthier people. Measured lifetime lost to vaccinations is of the same order as lifetime saved.

The new model is highly promising for generic mortality prediction. It is very different from the state of the art and many options for further research are outlined. A disturbing relation is found between vaccinations and the persisting excess mortality among the young and healthy population, warranting more research in vaccine safety.

Keywords: Mortality, temperature, frailty, pandemic, excess, deficit, covid, vaccine, QALY, lifetime

Statement of Interest

I declare that this work was done with an interest in science and personal safety for myself, loved ones, and humanity. Pro bono, independent, without payroll, not funded. For complete transparency: I applied for a Dutch government grant, within a program created on request of Dutch parliament for independent research into pandemic excess mortality [Zon]. My application was deemed inadmissible and rejected on the basis of my independence. I did continue my work and asked the reader community for support in [Red2,Red3,Red4]. See screenshot; it won't cover my electricity bills, but did make my day.

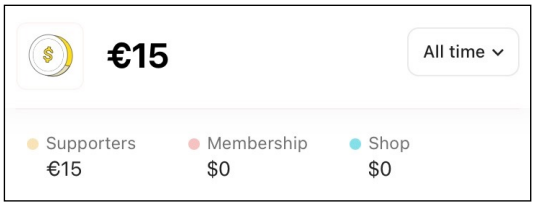

This report took a bigger toll than all my earlier reports combined, on time as well as on my financial situation.

If you like to support my independent research, feel free to visit buymeacoffee.com/andreredert or rodotti.nl/support. It is *highly* appreciated, every bit helps, thank you! Enjoy the report.

1. Introduction

In recent years, many countries around the world reported substantial excess mortality of unknown cause. An essential step in finding the cause is to quantify excess mortality accurately in the first place. Excess is defined as observed mortality minus expected mortality, and the latter depends on models that predict mortality based on various kinds of knowledge and data. National institutes typically provide a yearly mortality prediction, the "baseline", based on statistics of observed mortality in a few past years [Cbs1,Cbs3,Riv1,Riv3], aided by long-term models [Sto] that involve state and expected evolution of e.g. demographics, economy, healthcare, etc. A major issue with baselines is that they are highly subjective [Lev]. Searching for more objectivity, practical guidelines are model simplicity, explicitness and data objectivity.

The most powerful models of Nature are remarkably simple ($F = ma$, $E = mc^2$, $G = 8\pi T$, etc). Mortality models appear generally quite complex, e.g. [Sto] and [Kuh]. The latter uses detailed models based on population tables, longevity trends, etc, and finds substantial excess deaths in Germany during 2020-2022. Complexity is no guarantee for accuracy, however, as variance among nine candidate long-term-mortality models was of the same magnitude as the significance level of excess predicted by the chosen model. On the other hand, the used linear trends may have been too simple, resulting in superficial trend breaks and excess mortality. Both disappear with the use of a quadratic trend, see Figure 1.

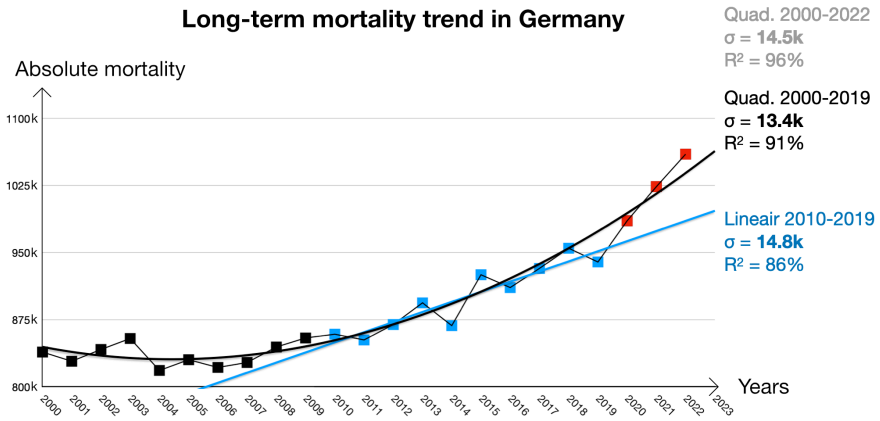

Figure 1: Long-term mortality dynamics in Germany. The blue 10-year linear trend suggests a trend-break near 2007 and substantial excess starting 2020. The black 20-year quadratic trend fits better and predicts no breaks or excesses. A full 23-year quadratic trend fits even better.

Source data varies tremendously in objectivity. Mortality is one of the most objective; death can be determined by anyone, is integer-countable, and such counts are publicly available on various aggregation levels (sub/super national) at weekly time resolution nearly everywhere in the world. On the other hand, death *causes* are far more subjective; e.g. accidental-falls, sudden-deaths and respiratory failure typically involve a long and partially hidden causal chain. The cause named "covid" exists in a context of influenza-like symptoms, mass testing, and novel treatment protocols that are non-default in both use and withholding of early treatment, respiratory supressants, antibiotics, mechanical ventilation, etc.

Publicly reported covid case numbers are among the least objective data, involving testing policies, public test-preparedness, complex chemical/biological processes, and decision models that transform analog curves into binary outcomes. Vaccine doses given over time may seem as objective as mortality, but doses were not always counted well, e.g. due to haste and privacy issues. Also vaccine contents varied not only by design per campaign, but apparently also by batch [Sch]. Demography also seems as objective as mortality, but recent population counts in the UK were so inaccurate as to strongly bias vaccine effectivity (VE) measurements [Nei1]. Further, demographic data typically has lower time resolution than mortality. Their combined use for e.g. population-relative mortality may require temporal interpolation, leading to accuracy loss.

My recent study to improve mortality models [Red4] found that in all of 30 EU countries investigated, mortality had much stronger weekly variability than suggested by the default Poisson model [Poi] which predicts variability equal to the square root of average mortality. This effect accumulates at other timescales, e.g. Germany's yearly-average mortality of ~1M people has an associated Poisson variability of ±1k, while ca ±15k is observed, see Figure 1. A strong population-wide, correlated influence on mortality appears to act on weekly basis, manifesting as increased variance at all time scales. The factor 15 requires that mortality is correlated within groups of *hundreds* of people *for all deaths*. There is no evidence that big families or all residents in care-homes always die together. Broken-heart-syndrome causes couples to die quickly after another [Moo], but even if this occurs for *all* couples, this only explains a factor $\sqrt{2} \approx 1.4$.

Temperature was suggested in [Red4] as cause of mortality correlation, since it is the main element in seasonal and heat-wave mortality, varies strongly on daily basis, and affects mortality within 0-3 days in the case of extreme heat [Xia]. Figure 2 shows a preliminary result: absolute all-cause-mortality (ACM) correlates very strongly with average outside temperature. Temperature plays a role in various causal chains, e.g. sunlight to temperature, sun UV light to vitamin D to immunity, UV light to ozone via ventilation to indoor-desinfection and air purifying, cold bad weather to slippery roads to accidental deadly falls, etc. Relative humidity (RH) influences viral transmission of influenza [Low], and maybe also of common colds and endemic covid. The role of absolute humidity seems to be that it determines indoor RH [Mar1] together with temperature.

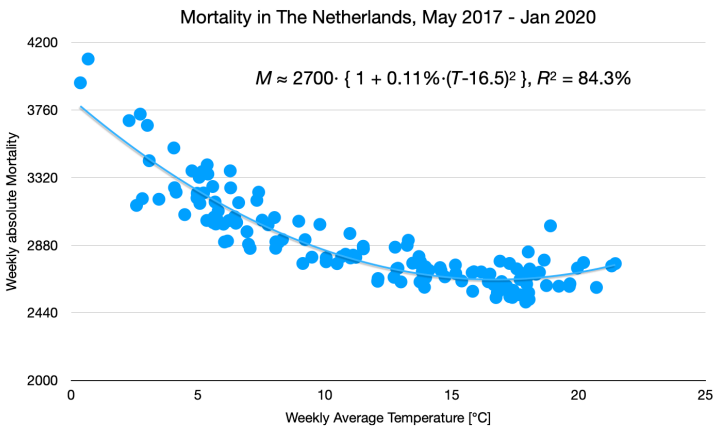

Figure 2: A preliminary result motivating this report. Weekly absolute mortality in The Netherlands correlates strongly with average outside temperature according to a quadratic model.

Temperature data has the strong advantage that it is highly objective and publicly available nearly everywhere in the world at high temporal and regional resolution. The correlation with mortality appears simply quadratic as in Figure 2, with minimum mortality occurring at ca 16.5 °C. This value appeared also in the context of the Dutch national baseline, but there a linear increase of mortality was used and only above that temperature [Cbs3], thereby missing the major mortality component on the cold side.

Age data, used in population tables, may be very objective but its relevance is limited to being a proxy for health. While health is more relevant, its data is generally subjective, unavailable, or inaccessable due to privacy issues. Searching for a way to combine the objectivity of age with the relevance of health, the concept of frailty may be useful. Based on Dutch population tables [Cbs2], age at death is ca. 80±11 years in The Netherlands. According to [Ger], after a healthy adult life, the time between onset of final health decline and death is only 4±1 years, a remarkably stable period hinting at a hardwired biological origin. I coin the population whose frailty is related to the final phase of life as the “frail pool”. This pool is the direct source of natural mortality, and is much more confined by health (4 frail years out of 80 total years, 5% of population) than by age (2x11 out of 80 years, ca 28%).

Figure 3 illustrates the frail pool by what I coin as “frail age”, which is zero when entering the frail pool. Modeling health via frail age may narrow down uncertainties in mortality. With normal age, birth variability resides in health but not age, while death varies strongly in age but not health. With frail age, these uncertainties are redistributed to concentrate at births, while deaths only include the minor uncertainties in frail age. Note that “reverse age”, normal age minus age at death, removes *all* mortal uncertainty which may be of similar use.

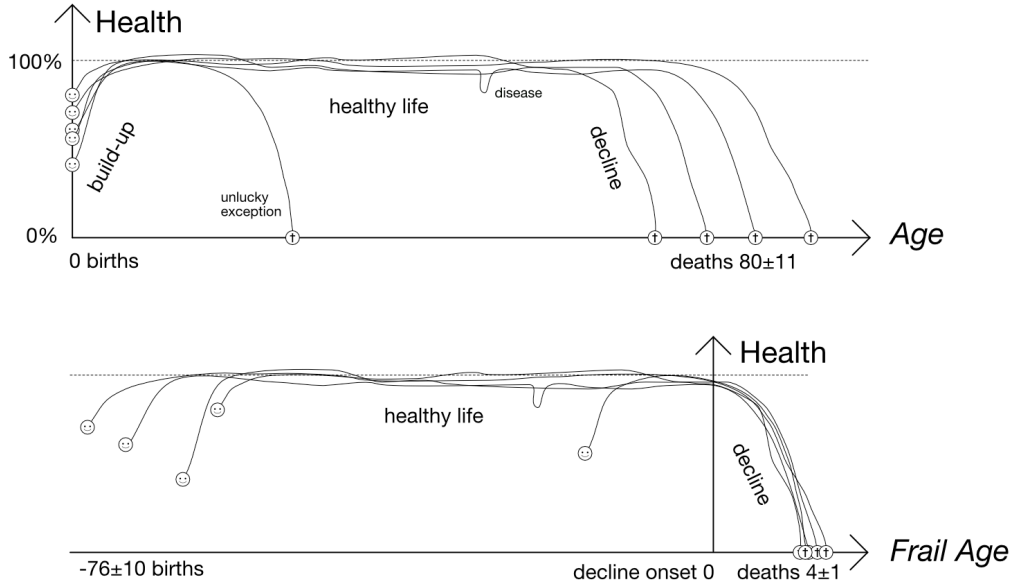

Figure 3: Modeling health via “frail age”. After spending a healthy life of greatly varying duration, people transition to the “frail pool”: ca 5% of the population with various normal ages. The duration of frail life until death is highly constrained to 4±1 years.

A generic advantage of frail age is that it is independent of normal age; it applies to entire populations, subpopulations by age, and stratifications by any modality *other than health*. Further, in between birth and death, migration influences population tables. This happens mainly during healthy life, and thus hardly affects the frail pool. Remigration in the frail does occur, but it effectively only cancels prior emigration.

Even without access to individuals’ health data, some population-wide aspects of the frail pool may be modeled, estimated from observed mortality, and advantageously used. Relevant for this report are mortality deficits after excesses or vice versa and their time lag. The latter enables to

derive life*time* lost/saved, which is much more relevant than death counts. Disease and treatment may accelerate or postpone death, but always have zero net effect on counts over the long term.

Ideally, lifetime lost/saved is adjusted for quality, represented by wellknown quality-adjusted life-years (QALYs). Quality of life and health are strongly intertwined, and health depends much stronger on remaining lifetime (reverse or frail age) than on enjoyed lifetime (normal age). So by nature, the frail age method relates to quality automatically according to the curve in Figure 3 for positive frail ages.

A special challenge for mortality models used for excess measurements is that they must predict not all mortality, but only part of it: *normal, expected* mortality. Without special care, one cannot fit a model and evaluate its predictions by a simple numerical difference measure with respect to observed mortality in fit and prediction periods. Observations reflect *all* mortality, possibly including abnormal, unexpected events. For this reason, one must explicitly define what is normal, such that model fit periods can be selected representative of that normal, or such that models can be defined that are insensitive to certain classes of abnormal events. An example of the latter is a linear model fitted to observed mortality that inevitably does include some abnormal events, where most of those events are zero-sum: an excess is followed some time later by an equal deficit still within the analysis period. The longer the analysis period, the more unexpected events are accounted for automatically; an argument for long-term analyses as in Figure 1.

A relaxed property of mortality models used for excess measurements is that forecasting is not needed. National baselines are used a.o. to prepare for future demographic developments, where forecasting is essential. Excess mortality on the other hand exists by definition in the present and past only. The required prediction is not across time, from known past to unknown future, but only across causes, predicting the expected to quantify the unexpected and localize it in time. Recently, several such excess mortality analyses have been made that are independent of existing national baselines and do not aim at forecasting, attempting to reduce subjectivity. In [Red1], a geographical-differential analysis was used, while [Red2,Red3] used a temporal-differential analysis effectively defining a baseline via a period around a moment in the past. In [Mee], a fixed baseline was estimated and used within the same analysis time frame.

This report's main contribution is a new mortality model, with the following features:

- Mathematically simple, one single equation
- Fits decades of data without trend breaks
- Predicts several years of data at week, month, seasonal and year resolution
- Just 12 parameters, 9-parameter version fittable by default linear regression
- Uses only temperature and absolute mortality data: objective, accurate and widely available
- No population or age data required, but population-relative mortality can be used as well
- Includes frail pool dynamics (FPD)
- Includes delayed mortality by so-called health pressure accumulation (HPA)
- Includes explicit modeling of mortality's observed natural random variability
- Embedded in the Bayesian probability framework
- Explicitly modeled and measured prediction accuracy
- Ideal for time-localized quantification of excess mortality
- More accurate than national baselines *for this purpose*
- Easily integrates determinants for possible causes of abnormal/excess mortality
- Specially suited to determine lost/saved life*time* per cause

The model can be used with population-relative mortality, but this requires additional population data and may not improve results in general due to migration, data (interpolation) errors, etc. Improved results are only expected in special cases, e.g. when both mortality and population are affected by the same type of disruptive dynamics. This may apply to age-limited analyses during recent years of accelerated migration in the EU.

The model's typical use is to predict "normal" mortality by temperature during a period in the past, and if abnormal excesses are encountered when compared to observed mortality, explicit determinants of possible causes are integrated to investigate their contribution. The model's forecasting power for future mortality depends on temperature forecasts. While beyond the scope of this report, Figure 2 suggests global warming will reduce mortality in The Netherlands.

As this report's secondary contribution, the model will be applied to the recent pandemic, to measure excess mortality and covid Case Fatality Rate (CFR), Vaccine-dose Fatality Rate (VFR), Vaccine Effectivity (VE), and lifetime lost/saved (LTL/LTS) by covid and vaccines. Positive covid tests and vaccination doses administered are used as determinants, both of which are unfortunately considerably less objective than temperature. Stringency indices (SI) of non-pharmaceutical interventions (NPIs) are not included.

Next follow the proposed mortality model, experiments, and a conclusion. The experiments involve years 2000-2023 and populations of 10 EU countries Belgium (BE), Bulgaria (BG), Finland (FI), France (FR), Germany (DE), Italy (IT), The Netherlands (NL), Poland (PL), Spain (ES) and Sweden (SE), totalling ca. 344M people. The model is assessed by fit/prediction errors and visual inspection. A comparison with the Dutch national baseline [Cbs1,Cb2] is made and experiments are performed with population-relative mortality and age stratification. Finally, the model is applied to the recent pandemic.

At many occasions, I will use results from preliminary experiments to shortcut possible routes of investigation. Figure 4 provides a preview of the model's results, which may provide some incentive to you as reader to plough through the many pages of this report.

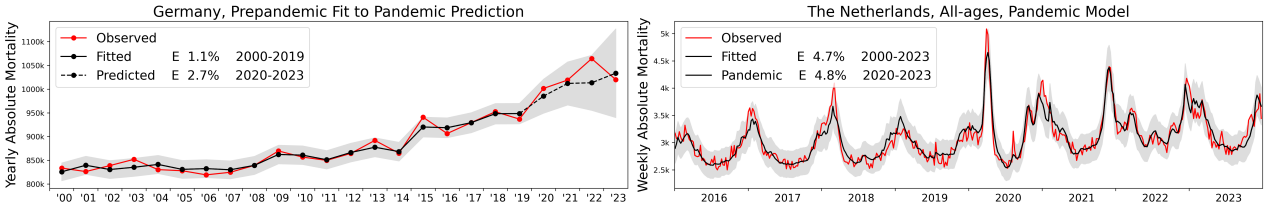

Figure 4: Result previews. Germany) Prepandemic to pandemic prediction shows near-zero excess during the entire pandemic, except for 2022. The Netherlands) Pandemic model fit with covid-test and vaccine-dose determinants.

2. Mortality model

In the next subsections I propose the normal mortality model, starting with a fast lane to the full model, followed by an overview, all details, various considerations, model optimization and assessment, and the use of other determinants that can be integrated in the model. Finally, the pandemic model is presented.

2.1 Fast lane to the full model

According to the preliminary result in Figure 1, the overall trend of mortality $M(t)$ over decades of time t is well-described by a quadratic relation $M = a + bt + ct^2$. According to the preliminary result in Figure 2, weekly mortality is highly correlated with temperature T by $M = a + bT + cT^2$. The strong resemblance between these two equations is remarkable, totally meaningless, and useful for their integration. Each of the three constants in the temperature model will be replaced by a 3-parameter time-trend via $M = a_{ij}t^iT^j$. The a_{ij} is a 3x3 matrix that contains 9 model parameters, 5 for long-term time trend via a_{i0} and weekly variability by temperature via a_{0j}, plus 4 additional cross-terms. The matrix can be estimated by default linear regression. To declutter equations, I will leave out time dependence (*t)* whenever possible, and use Einstein's summation convention: indices repeated in a product are to be summed over.

Health pressure accumulation (HPA) and Frail pool dynamics (FPD) are implemented via a linear filter W parametrized by two time constants τ_{HPA}, τ_{FPD}. These represent mean time between health pressure and death, and between excess mortality and subsequent deficit, respectively. The filter W is applied via temporal convolution denoted by the asterisk operator $*$ (with precedence in between multiplication and addition). Finally, unexpected mortality is denoted by ΔM, consisting of two very different components: normal weekly fluctuations with expected statistical characteristics parameterized by strength parameter α, plus all remaining excess by abnormalities still unaccounted for. Together, the final model equation is adopted as:

$$M = W_j * a_{ij} t^i T^j + \Delta M \qquad \text{(“M is Watt”)}$$

where W_0 is “no filter” (“do not filter long-term trend”) and $W_{1,2} = W$ (apply HPA/FPD filter on T and T^2). I like the aesthetics of the above formula, but will make use of less condensed notation in the following that is easier to read and manipulate, including normal equation numbers.

2.2 Normal and abnormal mortality

Figure 5 shows the model’s main ingredients that determine what is normal mortality, and what not. At the top is the natural process of life in the current civilized world: people are born, live their healthy lives with varying duration, transition to frailty and live just a few years until what is considered *normal* death. The timespans in years are illustrative and may differ per population. At the bottom are three categories of influences on health affecting births, life and deaths. The listed examples per category are illustrative and by no means exhaustive; in the model no mechanisms will be explicitly modeled, only their combined overall effect.

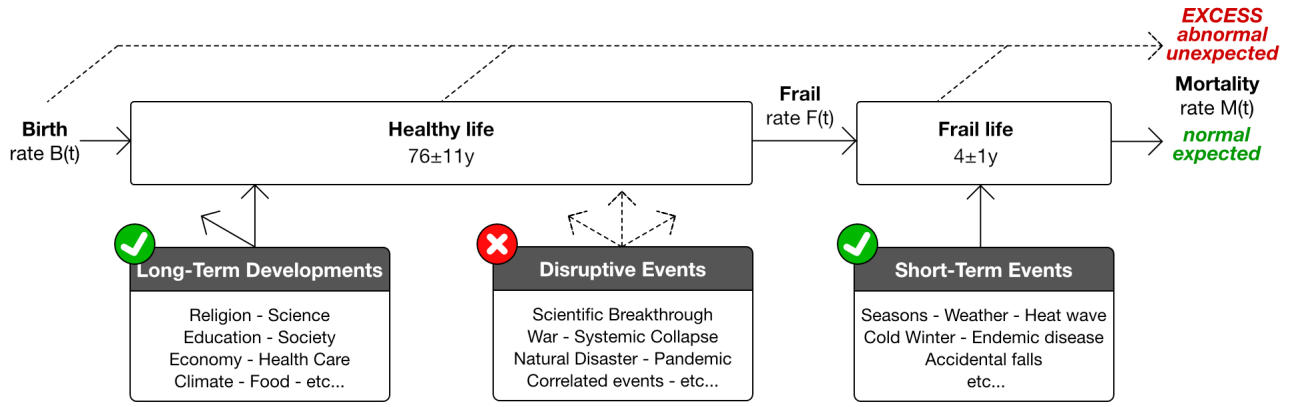

Figure 5: Overview of the mortality model.

Disruptive events may prevent births and cause mortality among all people, including the healthy, skipping the frail phase. On the other hand, they may also *prevent* mortality substantially. These events are *not* modeled, are unexpected/abnormal, and will thus appear as excess/deficit mortality when model predictions are compared with observed mortality. Disruptive events may include scientific breakthoughs in health care, war, population-wide exposure to new pathogens (pandemic) or toxic substances (bioweapon), and other kinds of sudden large-scale events.

Long-term developments act strongly but slowly over timescales of several decades, affecting birth rates and the lifetimes of the healthy. They *define* the number of people available to turn frail, the “frail rate”. But, long-term developments do not *affect* the frail, as frail people’s health decline is very fast and dominated by their frail internal biology.

Short-term events act strongly and instantly, up to time scales of weeks, months, a few years at most. These events have a strong impact on the frail, while the healthy have sufficient health

overhead to withstand or recover from such events. In this report, temperature serves as predictor, or proxy, of all short-term "events". Later in this report, explicit pandemic determinants will be added to the model, such that the pandemic is changed from an unexpected disruptive event into a series of expected short-term events.

Phenomena or their characteristics may be part of multiple categories. For example, accidental falls leading to death are a generic, expected and statistically significant part of frail life. These are short-term events in a personal sense, while their overall rate in the population has long-term dynamics. Further, deadly falls do occur among the healthy, albeit via more disruptive events (high altitude or speed falls) and at a much lower rate.

The frail-age approach does not exclude age-stratified analysis. On the contrary, it is fully compatible as mortality expectations by frailty apply equally well to young populations. The use of age stratification is the same as with other approaches. For example, if the deadly-falls rate among the healthy suddenly doubles, this is a highly significant event. It may go unnoticed in an all-age analysis but would stand out in an age-limited analysis (e.g. up to 65). To accomodate an age-limit, Figure 5's healthy and frail life would have extra output arrows "still healthy at age-limit" and "still alive at age-limit" respectively; these people exit the model alive. Such arrows and people are not shown in Figure 5; they are not needed in any of the following.

The transition from healthy to frail is not unlike crossing a black hole's event horizon: nothing particular happens at that moment other than losing the option of return. A clear definition of the transition to frailty would be subjective and hard to give, e.g. "so many stem cells per liter of blood", but luckily that is not required. If frail life endures for $\mu \pm \sigma$ years, the definition of frailty onset defines μ, and this μ disappears from the model as shown next.

2.3 Births, frailty and zero-sum variability

Mortality rates may vary over time, but over great time scales they equal birth rates on average. Ignoring migration for the moment, the whole life process is always *zero-sum*: no development, event, or medical intervention can cost or save lives, it can only accelerate or postpone death. On time scales of centuries, births are therefore an objective and accurate predictor of mortality:

$$M(t) = B(t - 80 \text{ years}) + \Delta M(t) \tag{1}$$

t	Time (at century resolution)
$M(t)$	Absolute mortality rate (e.g. deaths per century, averaged over a few centuries)
$B(t)$	Similar for birth rate, whose delayed values serve as mortality prediction
$\Delta M(t)$	Random zero-sum fluctuations, all unexpected as *only* births are used for prediction
80 years	Average human life time (may depend on time epoch, region, culture, etc)

A key point in my model is to apply the same zero-sum argument to the frail and the dead. I start the model at the frail transition, skipping births and healthy life altogether, while increasing time resolution to week level and introducing *expected* short-term mortality fluctuations:

$$\begin{aligned} M(t) &= F(t-\mu) &&+ \Delta M_{short-term}(t) + \Delta M(t) \\ &= M_{long-term}(t) &&+ \Delta M_{short-term}(t) + \Delta M(t) \\ &= M_{long-term}(t) &&+ \{W * H\}(t) \quad + \Delta M(t) \\ &= M_{expected}(t) &&+ \Delta M(t) \end{aligned} \tag{2}$$

t	Time (at week resolution)
$F(t)$	Frail rate, the rate of people transitioning from healthy to frail
$\mu \pm \sigma$	Expected lifetime of the frail, in frail age, ca 4±1 years
$\Delta M_{short-term}(t)$	Expected mortality fluctuations on short-term (*zero-sum* like ΔM)
$M_{long-term}(t)$	Timeshifted version of frail rate F, with shift by average frail lifetime μ
$M_{expected}(t)$	All normal mortality expected by the model
$H(t)$	Instantaneous health pressure on the frail by short-term events
$W(\Delta t)$	Health Pressure Accumulation (HPA) and Frail Pool Dynamics (FPD) filter
$\Delta M(t)$	All remaining fluctuations, zero-sum on various time scales

The timeshift between definitions of F and $M_{long-term}$ cancels μ in (2); one does *not* have to define where healthy life ends and frailty starts.

While short-term temporal fluctuations in births $B(t)$ do get imprinted in and travel through population tables over time, they do *not* end up in frail rate F or $M_{long-term}$ significantly. The variations in people's healthy lifetime (ca ±11 years) effectively smooth out any birth fluctuations, as a sort of information shredder. Therefore, $M_{long-term}(t)$ is a very *smooth* function of time, with variations only over multiple decades due to long-term developments.

Immigration and emigration take place for the major part during healthy life. Although they may not be zero-sum (unless world population is considered), their short-term dynamics also do not reach $M_{long-term}$ significantly. Remigration may take place more commonly among the frail, but this only cancels some of the prior emigration.

2.4 Long-term developments

As $M_{long-term}(t)$ is a highly-smooth function over time, it can be easily parametrized over several decades. I bluntly choose to model this smoothness as having at most one extremum per $L = 25$ years, more or less matching the time scale of healthy life's variability (±11y). A full-cycle up/down/back-to-base wave would thus have a duration of at least $2L = 50$ years. For any analysis up to L years, a parabola will then be sufficient to parametrize $M_{long-term}(t)$. For every additional set of L years, a polynomial of one degree higher would be required.

This long-term model states that mortality rates are effectively independent/decoupled across periods of L years. Therefore, taking into account more than L years is not recommended in this model as it would not add to prediction accuracy. In this report, all analyses cover years 2000-2023, and a parabola suffices:

$$M_{long-term} = a + bt + ct^2 = a_i t^i \tag{3}$$

t	Time in weeks (arbitrary unit)
a, b, c	Three arbitrary constants (a_0, a_1, a_2 are the same with index notation)

In light of all of the above, this model is not overly restrictive or mathematical. It models the combined effect of all long-term developments, rather than a specific set of known mechanisms, by a function whose degrees of freedom match those of all underlying mechanisms combined. At

this point, the a, b, c can be easily estimated using linear regression methods via (2) with observed mortality and the fact that both $\Delta M_{short-term}$ and ΔM are zero on average.

In the experimental section, the polynomial order of time denoted by O_t will be investigated, to see if 2 is in fact optimal.

2.5 Short-term events

In this model, short-term events such as heat or disease waves create an instantaneous health pressure $H(t)$ on the frail, which indirectly (via HPA and FPD) determines expected mortality fluctuations $\Delta M_{short-term}$ in (2). The instantaneous health pressure model is based only on a polynomial of temperature, e.g. quadratic $a + bT + cT^2$, *a* highly limited model but motivated in the introduction.

As FPD will filter all long-term components out of H causing $\Delta M_{short-term}$ to be zero-sum, the constant a is irrelevant and can be omitted. Long-term dynamics will affect remaining "constants" b and c. Modeling them each by a long-term parabolic trend over time, justified by the same argument as for (3), one obtains:

$$H = (b_i T + c_i T^2) t^i \tag{4}$$

$H(t)$	Health pressure (or hazard, etc), unit/scale is implicit via (2) and b_i, c_i
b_i, c_i	Six arbitrary constants
$T(t)$	Average weekly temperature, unit is irrelevant and absorbed by b_i, c_i

In the experimental section, the polynomial order of temperature denoted by O_T will be investigated, together with O_t.

2.6 Health pressure accumulation (HPA)

Health pressure does not always lead to instant mortality, but may acculate over a few weeks, leading to delayed mortality. A simple model for such accumulation is by temporal convolution:

$$H_{ACC}(t) = \sum_{\Delta t \geq 0} W_{HPA}(\Delta t) H(t - \Delta t) = \{W_{HPA} * H\}(t) \tag{5}$$

$H_{ACC}(t)$	Accumulated health pressure
$W_{HPA}(\Delta t)$	Relative weight of health pressure Δt weeks ago in accumulated health pressure

The overall scale of weights W_{HPA} is arbitrary; I choose them normalized to unit-sum so H_{ACC} has the same scale as H.

In preliminary experiments, I used brute-force numerical algorithms to optimize W_{HPA}. Figure 6 shows model fit error (precise definition is not relevant here) as a function of *number* of weights. The weights were freely optimized without any inter-weight restrictions. Results show that the accumulation acts within a range up to ca 1 month in the past. Note that *future* temperatures also correlate with current mortality; the weather system causes this as a confounder, and theoretically cremations are a true cause. Adding weights for future weeks did not reduce model error, however, once current and past weeks were already taken into account.

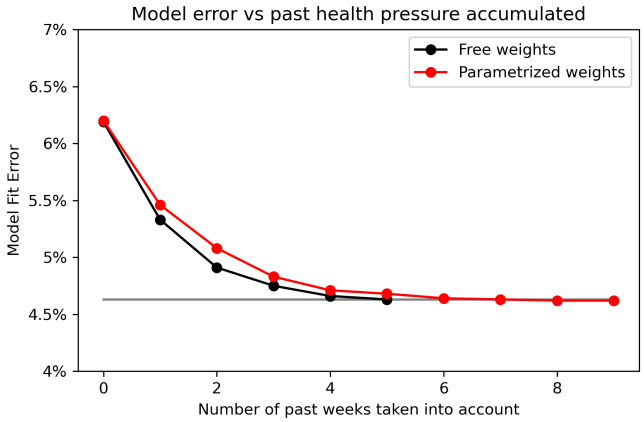

Figure 6: Model fit error as a function of number of past weeks taken into account.

The free weights obtained in multiple experiments were highly variable, as if not fully constrained by the model and observational data. This suggests the weights can be constrained without loss of model accuracy, e.g. by a parametrization. As Figure 6 suggests exponential decline with time passed, I tried:

$$W_{HPA}(\Delta t) = \tau_{HPA}^{-1} e^{-\Delta t/\tau_{HPA}} \tag{6}$$

τ_{HPA} Characteristic time constant of HPA, ca. 1 month

The performance of these weights parametrized by a single constant is nearly identical to that of the free weights. This is a so-called 1st order low-pass filter, that can be implemented very efficiently without limitation to Δt (number of weights) via a default FIR/IIR implementation with coefficients τ_{HPA}^{-1} · [1] / [1 , -1+τ_{HPA}^{-1}] .

The τ_{HPA} relates directly to average delay between health pressure and death. For temperature it is apparently ca. a month. For covid, the time between infection and death (if it occurs) has been measured at ca. 18 days [Mar2]; this is the value expected for covid's HPA time, denoted by $\tau_{HPA;C}$ in my pandemic model in a following subsection.

2.7 Frail pool dynamics (FPD)

Mortality rates are influenced by the frail pool's population size and frail age distribution. The pool is filled and emptied at different rates: $F(t)$ and $M(t)$. When $M > F$ the pool gets depleted over time, the current health pressure will be applied on less people and with a better health status, leading to less deaths until $M = F$ again and a new balance is reached.

The full dynamics of the frail pool are intricate and non-linear. I approximate this behaviour in greatly simplified form using a linear filter W_{FPD}, which can be combined with W_{HPA}:

$$\Delta M_{short-term} = W_{FPD} * H_{ACC} = W_{FPD} * W_{HPA} * H = W * H \tag{7}$$

$W_{FPD}(\Delta t)$ Linear filter representing the frail pool's fill/depletion dynamics

$W(\Delta t)$ Linear filter combining W_{FPD} and W_{HPA}

The W_{FPD} has a so-called *high-pass* characteristic, passing fluctuations but blocking all long-term dynamics. This ensures that $\Delta M_{short-term}$ is zero-sum, or zero-mean.

To determine W_{FPD}, I have done brute-force simulations of the frail pool using the canonical model for mortality by frail age instead of normal age. The canonical model considers population tables $N(a)$ containing the number of people per age a (the population *state*), and the probability rate of mortality by age $r_{\dagger}(a)$ (that together with births and aging determine population *dynamics*). The $r_{\dagger}(a)$ is known to resemble an exponential curve (growing ca. +13% per year of age). I assumed that the frail pool's mortality rate $r_{\dagger}(a_{frail})$ is well-represented by this canonical model, albeit with a much steeper exponential curve (e.g. +30% per frail year) to match death at frail age of 4±1y. I considered various ways in which health pressure adds to $r_{\dagger}(a_{frail})$, e.g. plain multiplicative and/or additive to the age exponent.

The details are completely beyond the scope of this report; the main result is that for health pressures from mild to considerable, but not disruptive, a simple linear so-called 1st order high-pass filter describes the frail pool's simulated behaviour quite well:

$$W_{FPD}(\Delta t) = \delta(\Delta t) - \tau_{FPD}^{-1} e^{-\Delta t/\tau_{FPD}} \tag{8}$$

$\delta(\Delta t)$	Delta function, one for $\Delta t = 0$, and zero otherwise
τ_{FPD}	Characteristic time constant of FPD, ca. 1 year

The first term in W_{FPD} represents excess deaths and the 2nd term the deficits that follow, equal in number but stretched out over a time period that is τ_{FPD} on average. That is, with each death caused by H, the net loss is not a life, as lives are never lost, but life*time* of mean duration τ_{FPD}. The same argument holds the other way round: negative excesses are deficits, which will be automatically followed by (true, positive) excesses, which in total has a net lifetime saving effect.

Similar to W_{HPA}, the W_{FPD} filter can be efficiently implemented without limitation to Δt via a default FIR/IIR filter with coefficients $(1 - \tau_{FPD}^{-1}) \cdot$ [1 , -1] / [1 , -1+τ_{FPD}^{-1}]. The latter requires 2 IIR coefficients, associated with a *scalar* state algorithm. Therefore, in this linear FPD model, the entire frail population's state at some moment is just *one* number, e.g. opposed to many numbers $N(a)$ in the canonical model. In the IIR implementation, the single state value is not the frail pool's population size but its absolute mortality.

Figure 7 illustrates the strong similarity of simulation and the approximation for a fast ramping 3-year-persistent health pressure and a short burst including seasons. For health pressures in a burst/wave, the excess follows health pressure dynamics, while the deficit follows frail pool dynamics. A persistent health pressure leads to excess mortality that starts instantaneously, but wanes smoothly and silently to zero in just a few years, without any deficit. At all times, mortality probability rates remain heightened, and average age at death remains reduced, but absolute mortality rates return to normal. Only once the health pressure is released, a mortality deficit occurs.

Excess figures are often publicly communicated by absolute mortality, or relative to other absolute mortality numbers (e.g. baseline or historic) instead of population-size. This feeds the false perception that excess is linear with instantaneous health threat: no excess, thus all is fine.

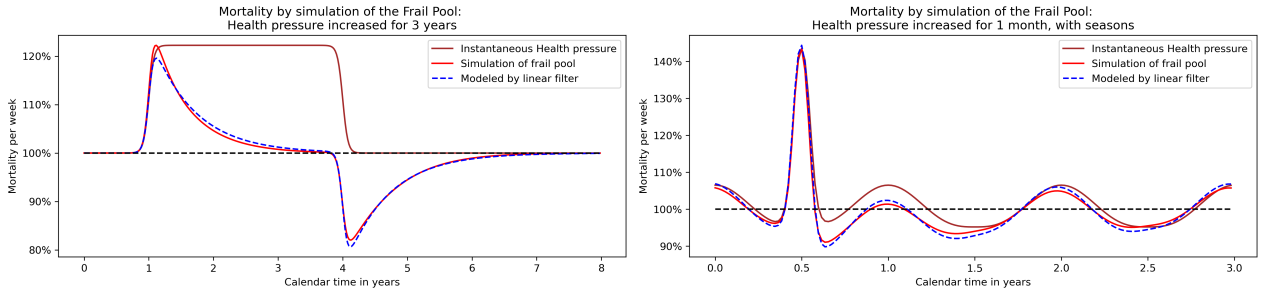

Figure 7: Brute-force simulation of health pressures in a canonical mortality model applied on frail age rather than normal age, and an approximation by a linear filter.

All of the above findings are not specific to the frail pool: the canonical mortality model applied to normal age behaves the same. The main difference is that the frail pool's dynamics are *an order of magnitude* faster, on time scales of years rather than decades.

The only parameter in the frail pool model that relates to the $\mu \pm \sigma$ expected lifetime of the frail, is τ_{FPD}. From here on, no further explicit reference to $\mu \pm \sigma$ is required; their relevant effects are fully represented by τ_{FPD}. In the simulations with $\mu \pm \sigma$ equal to 4±1 year according to [Ger], the τ_{FPD} was found to be ca. 9-12 months. The σ and τ_{FPD} in the order of a year suggest some relationship with seasons, which is not further explored here.

Reality is incredibly more complex than both the canonical and linear FPD model. So it remains to be seen if FPD adds to overall model accuracy. Integration of the canonical model requires more demanding optimization methods to fit the model, whereas the linear model is more forgiving. In this report I only use the linear model.

2.8 Interaction between HPA and FPD

HPA and FPD are distinct mechanisms with time constants that differ by a considerable factor (ca 1 month vs 1 year). Yet, their separate actions appear in my model integrated as a single linear filter W, whose FIR/IIR coefficients are $\tau_{HPA}^{-1}(1-\tau_{FPD}^{-1})$ · [1 , -1] / ([1 , -1+τ_{HPA}^{-1}] * [1 , -1+τ_{FPD}^{-1}]). The multiplicative constant $\tau_{HPA}^{-1}(1-\tau_{FPD}^{-1})$ is irrelevant as it will get absorbed by parameters a_{ij}. Without the constant, the model is symmetrical in HPA and FPD. Exchanging the two time scale parameters provides the same combined filter, so any measurement of their separate values can only be performed when a prior model is applied, e.g. a model *requirement* that $\tau_{HPA} < \tau_{FPD}$.

Figure 8 shows a preliminary brute-force optimization, where overal model fit error (precise definition not relevant here) is given for all $1 \leq \tau_{HPA}, \tau_{FPD} \leq 50$, for three datasets. One can observe the $\tau_{HPA} \leftrightarrow \tau_{FPD}$ symmetry in all cases. The left dataset shows two optimal τ_{HPA}, τ_{FPD} islands at ca. [5,36] and [36,5] weeks, numerically matching findings in the previous sections. This is a preliminary validation for the use of the HPA/FPD model.

The optimal islands are embedded in a broad somewhat hyperbolic-shaped band consisting of near-optimum model-fits. Degenerate cases may arise where HPA and FPD have merged at $\tau_{HPA} = \tau_{FPD}$ and can no longer be distinguished, or have diverged towards infinity at $\tau_{HPA} = \infty$ or $\tau_{FPD} = \infty$ and can no longer be measured. Divergence may relate to (disruptive) events that affect all people including the healthy. As the healthy outnumber the frail substantially, they will dominate τ_{FPD} and increase its value to decades. In the proposed model, expressions with HPA/FPD always involve *inverse* time constants τ_{HPA}^{-1}, τ_{FPD}^{-1} which will just get closer and closer to zero. A τ_{FPD} of 1000w (~20 years) will easily be indistinguishable from $\tau_{FPD} = \infty$.

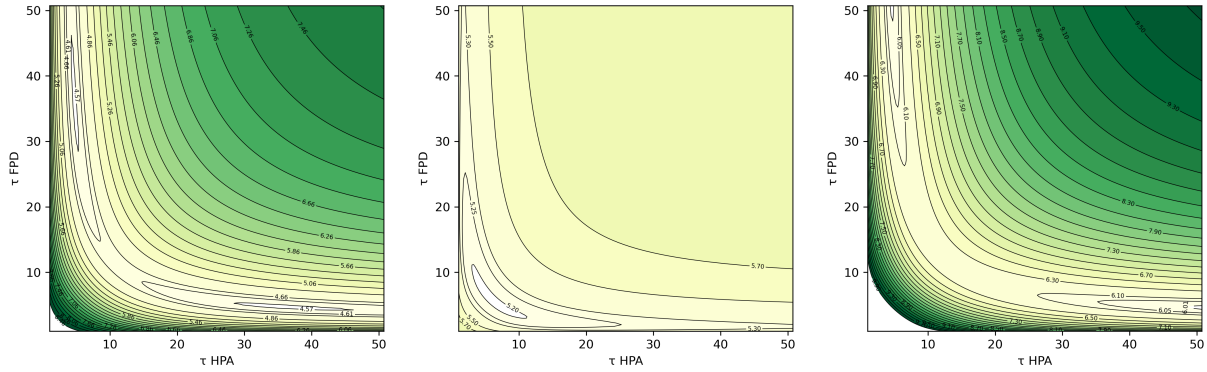

Figure 8: Brute-force model optimization with HPA and FPD time scales $1 \leq \tau_{HPA}, \tau_{FPD} \leq 50$. Isocontours represent overall model fit error. Left) Typical dataset that enables measurement of HPA and FPD. Center/Right) Some datasets lead to degenerate cases where HPA and FPD effects cannot be distinguished (merged) or measured (diverged).

By accepting a possible slight loss of model accuracy, one can *choose* to incorporate merged degeneracy into the model and simplify it to $\tau_{HPA} = \tau_{FPD} = \tau$. This τ represents the merged action of both HPA and FPD, but neither mechanism specifically. The merged filter has FIR/IIR components $\tau^{-1}(1-\tau^{-1}) \cdot [1,-1]/[1,-2+2\tau^{-1},1-2\tau^{-1}+\tau^{-2}]$. The hyperbolic band describing the relation between the three time constants is well approximated by $\tau_{HPA}^{-1/2} + \tau_{FPD}^{-1/2} \approx 2\tau^{-1/2}$. For the [5,36] result in Figure 8, τ is ca 11.

Whether degeneracy is an issue or not, depends on multiple factors. While data quality obviously has an effect on the model's ability to fit and predict mortality, HPA/FPD degeneracy may not. If one is interested in τ_{FPD} values, e.g. for measuring lifetime lost/saved, degeneracy is undesired. But if one is uninterested in τ_{HPA},τ_{FPD} values, the choice for the full/separate or degenerate model should be guided only by fit and prediction accuracies of mortality. If the seperate HPA/FPD model is overfitting, model fits are better but predictions worse, and the explicitly degenerate model is in fact preferred.

In the experimental section, both full and degenerate (merged) HPA/FPD models will be evaluated. Out of interest in τ_{FPD} values, I concentrate on full rather than degenerate models, and handle divergence by estimation of lower bounds.

2.9 Modeling the unexpected: normal random variability and abnormal excesses

The remaining mortality fluctuations ΔM are modeled with two components:

- "Truly" random, ever-present fluctuations caused by a huge inaccessible chain of events
- Abnormal events, excess mortality by unexpected but in principle accessible causes

The law of great numbers justifies modeling the random part probabilistically with a normal distribution, while for the abnormal excesses a something-or-nothing model is used:

$$\Delta M = \mathcal{N}_{0,\sigma_{expected}} + \Delta M_{excess} \quad , \quad \Delta M_{excess} = \begin{cases} 0 & \text{Period of normal mortality} \\ ? & \text{Period with abnormal mortality} \end{cases} \tag{9}$$

$\mathcal{N}_{\mu,\sigma}$	Normal distributed random variable with mean μ and deviation σ
$\sigma_{expected}$	Expected magnitude (deviation) of normal random mortality fluctuations
$\Delta M_{excess}(t)$	Abnormal excess mortality: zero, or not

Both terms are zero-sum, but on vastly different time scales; $\mathcal{N}_{0,\sigma}$ yields a process uncorrelated over time whose average goes to zero over periods of just several weeks, while ΔM_{excess} could have net effects at considerable time scales. In my model, excesses become "normal" and zero-sum automatically at time scales of a few τ_{FPD} (a few years) if only frail people are affected, or up to L (ca 25 years) if also healthy people are affected. This distinction is relevant as it allows to determine whether only frail or also healthy people are affected by abnormal mortality.

My study on the natural variability of mortality [Red4] found $\sigma_{expected}$ is well-modeled by:

$$\sigma^2_{expected} = M_{expected} + \alpha^2 M^2_{expected} \tag{10}$$

α	Scale constant for mortality fluctuations that are linear with mortality
$\sigma_{expected}(t)$	Time-dependent mortality deviation, changes together with (expected) mortality

The first term in (10) represents wellknown Poisson noise, the 2nd represents an omnipresent but sofar not-understood component with $\alpha \approx 3\%$ for weekly mortality [Red4]. Possibly, the proposed temperature model explains part of mortality fluctuations, leading to lower α. Higher α may also be encountered here, as [Red4]'s method to determine α automatically compensated for *any* kind of abnormal mortality lasting a month or longer. In this report, anything-not-predictable-by-temperature-during-normal-years will add to α, such as influenza-winter-variability.

Equation (10) is the only part of the model that is taylored towards absolute mortality. When mortality relative to population size P is used, one can replace M by M/P everywhere in the model, but (10) changes into $\sigma^2_{expected} = (M_{expected}/P)/P + \alpha^2(M_{expected}/P)^2$, that is, an extra P divides the first term.

The something-or-nothing model for ΔM_{excess} determines which periods can be used to fit the model; clearly those with zero excess (no abnormal mortality) are preferred. Any abnormal event outside of the fit period but within a prediction period, can then be detected and quantified by the model. By introducing other determinants in the model, ideally some or all of the mortality in ΔM_{excess} is shifted towards $M_{expected}$.

Altogether, modeling the unexpected leads to one extra parameter α that characterizes the strength of mortality's normal or expected randomness.

2.10 Entire model

The entire model is obtained by combining all previous equations into:

$$\begin{aligned} M &= M_{expected} + \Delta M \\ &= M_{long-term} + \Delta M_{short-term} + \mathcal{N}_{0,\sigma_{expected}} + \Delta M_{excess} \\ &= a_i t^i + W * (b_i T + c_i T^2) t^i + \mathcal{N}_{0,\sigma_{expected}} + \Delta M_{excess} \\ &= W_j * a_{ij} t^i T^j + \mathcal{N}_{0,\sigma_{expected}} + \Delta M_{excess} \end{aligned} \tag{11}$$

The 12 model parameters a_{ij}, τ_{HPA}, τ_{FPD} and α are a non-mechanistic, generic, approximate and incomplete representation of all processes underlying long-term developments, short-term events that correlate with temperature or its square, and the randomness of normal mortality. The units of

both time and temperature are arbitrary (days, weeks, years, K, °C, °F), as scales and zero-points are absorbed by the a_{ij} and/or filtered out by the FPD filter via W.

The parameters are categorized in linear (the nine a_{ij}) and non-linear (the three τ_{HPA}, τ_{FPD} and α). They can be obtained by brute-force numerical methods, and default linear regression under some circumstances discussed in the next subsection .

2.11 Bayesian analysis: parameter fit and accuracy

To fit the model, I follow the Bayesian probability approach that maximizes the posterior probability density of all parameters given mortality and temperature observations:

$$
\begin{aligned}
\phi_{fit} &= \arg\max_{\phi} p_{\phi|M,T} \\
p_{\phi|M,T} &= \frac{p_{\phi,M,T}}{p_{M,T}} \\
p_{\phi,M,T} &= \int_{\Delta M} p_{\Delta M,\phi,M,T} = \int_{\Delta M} p_{M|\Delta M,\phi,T}\, p_{\Delta M,\phi,T} \\
&= \int_{\Delta M} \delta\big(M - M_{expected}(\phi, T) - \Delta M\big) p_{\Delta M|\phi,T}\, p_{\phi} p_{T} \\
&= p_{\Delta M|\phi,T}\big(M - \hat{M}\big) p_{\phi} p_{T} \\
&= K \prod_{t} \frac{1}{\sqrt{\hat{M} + \alpha^2 \hat{M}^2}} e^{-\frac{1}{2}\frac{|M-\hat{M}|^2}{\hat{M} + \alpha^2 \hat{M}^2}}
\end{aligned}
\tag{12}
$$

ϕ	All model parameters
ϕ_{fit}	Parameters fitted to observed mortality and temperature M, T
$p_{\phi\|M,T}$	Posterior probability density of parameters, given observed M, T
$p_{M,T}$	Prior probability density of M, T, constant as M, T are fixed to observed values
p_{ϕ}	Prior parameter probability density, constant/uniform on a restricted domain
$\hat{M}$	Short-hand for $M_{expected}(\phi, T)$
$p_{\Delta M}$	Probability density of normal mortality deviations, normal distribution by (9)-(10)
K	Constant that absorbed all previous other constants

In (12), I make use of the independence between T and ϕ prior to observing M. The argument of the δ function represents the proposed model (11). The probability $p_{\Delta M}$ uses $\Delta M_{excess} = 0$ as the model is fit on a period of normal mortality. The time product runs over all observed weeks selected as fit period, which is in the order of 1000-1250 weeks (20 to 24 years).

The constant (uniform) prior parameter probability density p_{ϕ} reflects completely free parameters, without bias to any sort of prior expectations, albeit with two exceptions. First, physical restrictions such as positivity of τ_{HPA}, τ_{FPD} and α, are reflected by zero p_{ϕ} if restrictions are not met. This is not shown in (12); the restriction is implemented in the maximization by searching only within the physically allowed parameter domain. Secondly, to match the inverse

nature of HPA/FPD parameters (they always appear in equations as τ_{HPA}^{-1} and τ_{FPD}^{-1}) and handle diverged degeneracies properly, these parameters are natively used in inverted way which range uniformly between 0 and 1 without any singular behaviour.

There is, however, no free lunch: by assigning each τ^{-1} a uniform prior distribution over [0,1], each τ automatically acquires a *non-uniform* prior distribution over $[1,\infty >$ with density τ^{-2}. This prior actually reflects sensible expectations that any τ is not extreme or infinite, but weeks, years, decades at most as it certainly remains below my model's long-term-developments time scale L. The bias towards low values by this prior is extremely mild/weak; while τ^{-2} is a valid probality density, both its expected value and variance are infinite similar to Cauchy distributions, providing ample room for high-valued τ. Further, the more than 1000 mortality equations (one per week in a two-decade analysis) overpower even strong priors *very* easily.

Readers comfortable with probability distributions lacking expected value and variance may skip to the next paragraph. For those still here: each linear parameter a_{ij} in my model has a uniform prior distribution over *all* real numbers, which I coin gamma, $\gamma(x)$. It has an integral equal to 1 but is zero everywhere (or 0^+ if you wish). This $\gamma(x)$ lacks far more than expectation and variance, namely even any non-zero value at all. It is similar to the Dirac delta function $\delta(x)$, but infinitely dispersed instead of concentrated. *All* models in use worldwide that apply default linear regression to obtain parameters are equivalent with Bayesian models that use gamma distributions as parameter priors (and normal distributions for all other random variables that must relate additively to some function of the observables). Silently, the super-degenerate $\gamma(x)$ is one of the most used distributions in the world.

A brute-force numerical algorithm is required for the maximization of $p_{\phi|M,T}$ to obtain ϕ_{fit}. I use a proprietary mix of wellknown methods such as full/stochastic/accelerated search and gradient/ quadratic descent. To determine accuracy of parameters and quantities derived from multiple parameters (lifetime lost/saved), I use default Gibbs sampling followed by default statistical analysis of the samples. Depending on tail behaviour of $p_{\phi|M,T}$, parameter mean μ_ϕ, deviation σ_ϕ and 95% Confidence Intervals (CI), Lower/Upper Bounds (CLB/CUB) are well-defined and can be measured. This will be in most cases except when HPA/FPD time constants have diverged; in that case at least their CLBs can still be measured.

For the prediction of mortality I use ϕ_{fit}, with prediction accuracy determined as in the next section. Values of parameters and derived quantities are always presented in terms of μ_ϕ, σ_ϕ and/ or CLB/CUBs.

The maximization of (12) and parameter accuracy analysis can also be performed without brute-force algorithms, by resorting to default linear regression on the *linear parameters only* and a very acceptable approximation:

$$\begin{aligned}\phi_{fit} &= \arg\min_{\phi} -2\ln p_{\phi|M,T} \\ &= \arg\min_{\phi} \sum_t \ln(\hat{M}+\alpha^2\hat{M}^2) + \frac{|M-\hat{M}|^2}{\hat{M}+\alpha^2\hat{M}^2} \\ &\approx \arg\min_{\phi} \sum_t w|M-\hat{M}|^2 \quad , \quad w=\frac{1}{M+\alpha^2M^2}\end{aligned} \tag{13}$$

The bottom row is an accurate approximation where all-but-one $\hat{M}$ are replaced by M; given that the proposed model is of any use, $\hat{M}$ should resemble M up to a few percent. The logarithmic term then becomes constant and disappears from the minimization, resulting in a sum of squares weighted by w. This can be minimized over all a_{ij} by default weighted linear regression, yielding parameter values (ϕ_{fit} and μ_ϕ, these are the same due to the model approximation) and their accuracies (σ_ϕ). Brute force optimization is still required for the non-linear parameters α and the τ's. These, however, are assumed to be more constant than a_{ij} over different periods and regions, so one can reuse the generic values given in the experimental section.

The fit period ideally only contains *normal or expected* mortality, but is in fact allowed to contain some unexpected events on the sufficiently-frail population, as these result in a quick follow-up of matching deficits. A decades-long fit period then suffices to handle the unexpected events via their net zero-sum effect over several years.

I considered and rejected two possible modifications to the fit procedure. First, to improve prediction accuracy one could use temporal weights in (12) or (13) that are skewed towards the prediction period. Preliminary results yielded no benefits at the cost of reduced simplicity.

Secondly, it would be interesting if somehow some unexpected events could be excluded automatically during fitting. For example, exclude the 10% lowest and highest observed mortality, taking into account only an 80% bulk of non-extreme observations. Although interesting, this approach easily leads to bias. Unexpected excesses may be peaked while subsequent deficits are flattened, see Figure 7. This may result in unbalanced exclusion, leading to bias. In The Netherlands, the current baseline actually uses some form of this [Riv1], where 20%-25% of peak mortality is excluded.

2.12 Accuracy of fitted and predicted mortality: error bandwidth

The Bayesian analysis is targeted at optimal parameter fitting, which involves a mix of mortality fitting, mortality's natural variation (via α), and (a very limited set of) priors. To evaluate fit and prediction accuracy of only mortality, I use a simple ratio measure E of remaining unexpected mortality ΔM by observed mortality M, closely related to conventional model fit measure R^2:

$$E = \sqrt{1 - R^2} = \sqrt{\frac{\mu_{\Delta M^2}}{\mu_{M^2}}} = \sqrt{\frac{\mu_{\Delta M}^2 + \sigma_{\Delta M}^2}{\mu_M^2 + \sigma_M^2}} \tag{14}$$

The R^2 would be equally informative as E but less intelligible, as its typical values in this context would be squeezed in a very narrow range (e.g. 99.8%-100%). E resembles approximation (13) in the Bayesian analysis, with main difference that (13) sums over a ratio and (14) is a ratio over sums. The denominators are similar when $\alpha^2 M^2 \gg M$, or $M >$ ca 1k/week at α ca 3%. This typically holds for all-age populations over ca. 5M people.

The E can be computed in both fit and prediction periods as my model is intended for the present and past only; mortality data is always available. The fit period is selected to contain only normal mortality; E_{fit} will be the fit error, or fit "bandwidth" $\pm 2E_{fit}$ around observed mortality in which 95% of all fluctuations reside. The prediction period can be selected either with normal or abnormal mortality; in this report prepandemic data not used for fitting, or pandemic data respectively. $E_{prediction}$ will serve very different roles for these two options.

Although I always use the exact E, these approximations illustrate the different roles of E:

$$E \approx \frac{\sqrt{\mu_{\Delta M^2}}}{\mu_M} = \frac{\sqrt{\mu_{\Delta M}^2 + \sigma_{\Delta M}^2}}{\mu_M} \quad , \quad E_{fit} \approx \frac{\sigma_{\Delta M}}{\mu_M} \quad , \quad E_{prediction} \approx \frac{\sigma_{\Delta M} \text{ or } |\Delta M_{excess}|}{\mu_M} \tag{15}$$

In the denominator approximation, σ_M^2 is negligible with respect to μ_M^2. The σ_M has a weekly random component of magnitude ca. ±3% of μ_M [Red4], a seasonal component of roughly ±10%, and a long-term component with strength in between. The approximation is thereby accurate with relative error < 1% (due to the squared appearance of σ_M^2 vs μ_M^2 and the square root that provides ca. a factor half). E is thereby essentially equal to the intuitive ratio of "root-mean-square" model error $\sqrt{\mu_{\Delta M^2}}$ per average observed mortality μ_M. Despite all media hype, the recent pandemic component is of the same order as seasonal variation and does not invalidate the approximation.

The enumerator approximation differs for fit and prediction period. In the fit period, $\mu_{\Delta M}$ is almost zero due to model parameter a_{00} (not fully as E differs from $p_{\phi|M,T}$). This simplifies E_{fit} to the intuitive ratio of $\sigma_{\Delta M}$ per average observed mortality μ_M. The same holds for the prediction period if it contains only normal mortality; then, $E_{prediction}$ will just be somewhat higher than E_{fit} depending on temporal distance between fit and prediction periods. In prediction periods with abnormal events, however, $\mu_{\Delta M}$ can be substantially non-zero representing a temporal net excess or deficit ΔM_{excess}. Then, $E_{prediction}$ reflects the relative size of that excess or deficit, with $|\Delta M_{excess}|$ as approximate enumerator.

To determine significance of $E_{prediction} \approx |\Delta M_{excess}|/\mu_M$ in prediction periods with abnormal mortality, it must be compared to $E_{prediction;normal-mortality-only} \approx \sigma_{\Delta M}/\mu_M$, that is, the value of $E_{prediction}$ had there been no abnormalities. The latter cannot be measured from observations as these include the abnormalities. Public mortality counts do not come separated in two neat and reliable components "normal" and "abnormal", otherwise this model would have no purpose.

An essential requirement for determining excess significance is to be able to estimate $E_{prediction;normal-mortality-only}$ *without* using observed mortality in the prediction period. For this, in a phase preceeding the fit/prediction on the current dataset, other fit/prediction periods can be chosen in other datasets that are known to contain only normal mortality (prepandemic). Then both E_{fit} and $E_{prediction}$ can be computed, both relate only to normal mortality, and a fit-to-prediction error "amplification" factor can be determined as defined here:

$$A = \frac{E_{prediction}}{E_{fit}} \tag{16}$$

Presumably $A > 1$, rising with temporal distance between fit and prediction periods. With sufficient normal-mortality fit/prediction data sets, average and deviation $\mu_A \pm \sigma_A$ can then be obtained via statistics, as generic properties of the proposed model.

When predicting normal mortality in periods that do include abnormal events, the model's prediction error of (unobservable) normal mortality within the prediction period can then be estimated using the available E_{fit} measure:

$$E_{prediction;estimated-for-normal-mortality-only} = \mu_A E_{fit} \tag{17}$$

Subsequently, the significance of resultant excesses or deficits ΔM_{excess} can then be characterized by a Z-score as in "the excess is Z-sigma significant":

$$|Z| = \frac{E_{prediction}}{E_{prediction;estimated-for-normal-mortality-only}} = \frac{E_{prediction}}{\mu_A E_{fit}} \tag{18}$$

This $|Z|$ is always positive due to the definition of E. During prediction years with abnormal mortality (nonzero ΔM_{excess}), approximation (15) provides Z including sign:

$$Z \approx \frac{\Delta M_{excess}/\mu_M}{\mu_A E_{fit}} \tag{19}$$

A widely used error bandwidth is $|Z| < 2$: the 95% confidence interval of unsignificant excess.

The μ_A strongly depends on the temporal relation between fit and prediction period, e.g. interpolating 2 years centered in 2 decades, denoted as $A_{9\to 2\leftarrow 9}$ may have $\mu_A \approx 1$, while extrapolating 1 year located 3 years in the future of two decades, denoted by $A_{20\overset{3}{\to}1}$ (the number above the arrow indicates years skipped), may have $\mu_A \gg 1$. I assume that there is no systematic difference between forward vs backward extrapolation, e.g. $A_{20\overset{3}{\to}1}$ and $A_{1\overset{3}{\leftarrow}20}$ are equivalent.

In this report, the normal part of mortality during the pandemic will be predicted. As 2023 cannot already be considered fully postpandemic, there are no postpandemic years yet to perform interpolative prediction, and only extrapolation can be used. For years considered in this report, 2000-2023, prepandemic to pandemic extrapolation is denoted by $A_{20\to 4}$ for the whole pandemic, and by $A_{20\to 1}$, $A_{20\overset{1}{\to}1}$, $A_{20\overset{2}{\to}1}$, $A_{20\overset{3}{\to}1}$ for individual years.

With prepandemic (normal mortality) data limited to 2000-2019 in this report, no data is available to determine the required statistics $\mu_{A_{20\overset{s}{\to}p}}$. Therefore, I approximate these by $\mu_{A_{\geq 15\overset{s}{\to}p}}$ that use at least 15 years of available prepandemic years for the model fit. For example, $A_{20\overset{2}{\to}1}$ will be approximated by $A_{\geq 15\overset{2}{\to}1}$ consisting of $A_{15\overset{2}{\to}1}$ (prediction of 2002 and 2017), $A_{16\overset{2}{\to}1}$ (2001 and 2018) and $A_{17\overset{2}{\to}1}$ (2000 and 2019), totalling $6N$ measurements for N datasets.

Once postpandemic years with normal mortality become available, interpolation can be used. If 2024 would be postpandemic, e.g. $A_{15\overset{2}{\to}1\overset{1}{\leftarrow}1}$ and similar can be used. This will not be the case, however, as pandemic excesses not only must have come to an end but also have been followed by matching deficits. At the earliest, this occurs in 2024 as part of the pandemic's final phase.

Depending on arbitrary preference, the required A's (or μ_A's) can be regarded as model evaluation results, or as additional model parameters. I prefer the latter, but stress that the μ_A's are of a different level. They represent the *generic model's* characteristic of prediction accuracy,

estimated using *many datasets*, contrary to the normal 12 parameters that represent a *particular dataset*. The number of required μ_A's depends on which and how many periods one wants to predict, at what model time resolution, and how finegrained over time one wants to estimate prediction accuracy.

Although E and A are "objective", in the end all results must pass the test of human scrutiny. Visual inspection and interpretation of predicted mortality over time remain crucial as final judge.

2.13 Simplicity and number of parameters

The proposed model has only one equation and 12 parameters. Although the model has an integrated model fit error/bandwidth via one of the parameters (α), an explicit parameterless measure E is defined describing model fit error/bandwidth more intuitively. Prediction error bandwidths can be generated at will, represented by additional parameters μ_A estimated from a sufficient number of datasets representative of normal mortality. E.g. when 5 μ_A's are estimated from 10 datasets, they total "0.5 parameters per dataset". This number goes further down the more datasets are used.

Typical national baselines have long-term projections involving a complex multitude of equations and parameters, whose total number is not perse clear. My model's equivalent of these projections are quadratic trends represented by the linear parameters, and the μ_A's. For seasonal variations, the Dutch national baseline at [Cbs1,Cbs3] used no less than 104 (!) parameters, a mean and deviation for each week of the year. It is not clear to what extent implicit smoothness constraints (6-week averages over time are used) lower the *effective* number of parameters. Certainly, many more will remain than the 12 of the proposed model. The Dutch baseline has been updated recently [Riv1] and uses a different approach now.

The number of parameters in the proposed model equals $(O_t + 1)(O_T + 1) + 3$, with O_t, O_T the polynomial orders of time t and temperature T, while 3 is for τ_{HPA},τ_{FPD} and α. If the orders are increased, the number of parameters grows, and the model's fit accuracy improves *by definition*. The only justification for such model expansion is that the model's *predictions* also become more accurate, as will be examined in the experiments.

Alternative equations as $\left(a + bt + ct^2\right)\left(1 + d(T - \breve{T})^2\right)$ are also polynomial in both time and temperature while requiring less parameters than the proposed model at equal polynomial order. Here, a minimum-mortality-temperature $\breve{T}$ is modeled as time-independent (Figure 2 suggests $\breve{T} \approx 16.5$ °C over at least 2.5 years). Preliminary results with several such models showed no clear accuracy benefit, while the non-linear relations among parameters are more computationally demanding when fitting the model.

2.14 Temperature data preparation and reduction

Temperature data is typically available from weather stations all over the world at a fine time resolution of an hour, and at considerable regional resolution, e.g. The Netherlands has 50 weather stations [Knm]. Huge such temperature datasets can be reduced to single weekly, population-wide datapoints for use in the model without loss of accuracy. As the model terms containing temperature are linear in T^j, one can average T^j over all hours in a week, and all regions considered, arriving resp. at μ_T (average temperature) and μ_{T^2}(average of square of temperature), which itself equals $\mu_T^2 + \sigma_T^2$.

One may use appropriately weighted averages for μ_{T^j} to increase model accuracy. Regional averages may use regional population density as weights. Similarly, if daily temperatures have

higher (or less) influence on mortality than during the night, also hourly weights can be used to take into account that difference. In this report, I do not take into account population density for temperature averages; all regions are weighted equally. From preliminary experiments, I found no clear improvement by using hourly weights other than uniform (all 24 hours equally weighted); it appears that night and day temperatures are equally important in terms of health pressure.

Possibly, the human body or housing somehow smooths out temperature over several hours and/or days, thereby remaining sensitive to temperature dynamics only at time scales of more than a few days. At shorter time scales, only average temperature μ_T should then be taken into account, ignoring μ_{T^2} (and σ_T^2). From preliminary experiments, I found no use of taking this effect into account.

A more implementational note, as t changes very slowly compared to T, to high precision the result of $W_j * (a_{ij} t^i T^j)$ in model (11) can be approximated by applying filter W only to temperature as $a_{ij} t^i (W_j * T^j)$, which may simplify data processing. I use this in all results.

2.15 Transient effect at start of analysis

Every analysis starts at some time t_{start}, and ignores the past before it. Simply applying the HPA/FPD filter on data starting at t_{start} is equivalent to applying it to all of time, with health pressure equal to *zero* before t_{start}. While this is fully correct for health pressures as covid that simply did not exist before some date, it is not correct for the everpresent health pressure of natural mortality, represented by temperature. For the latter, a zero start health pressure produces a strong spurious transient effect at the start of the analysis, with duration of several τ_{FPD} of time.

This effect is not an implementation-level detail, but an essential property of the model relating to the state of the frail pool. All health pressure events prior to t_{start}, and mainly only those of a few τ_{FPD} time before, influence the frail pool's state at t_{start}. That state is required at the start of the analysis.

The simplest way to handle the transient effect is to observe health pressure (temperature) starting a few years earlier, at $t_{start} - 2..3\tau_{FPD}$, and apply the filter on that so that transient effects have vanished at t_{start}. The full mortality model is then applied from t_{start} and on, with mortality observations also starting from t_{start}. This solution requires considerable additional health pressure data prior to t_{start}, to determine the frail pool's state at t_{start} which in my simplified model is just *one* number. An alternative is to (gu)estimate the average health pressure (temperature and deviations) before t_{start}.

In this report, all data starts at t_{start} of week 1 of year 2000, and I opted for a guestimate of earlier health pressures, by mirroring available data towards the past:

$$T(t_{start} - \Delta t) \approx T(t_{start} + \Delta t) \tag{20}$$

I use Δt covering all years used in the analysis; for 2000-2023, Δt will run from 0 to 24 years. When this is fed into the HPA/FPD filter W, it automatically filters health pressures including seasonal effects in a weighted way, with most weight assigned close to t_{start} ($\Delta t \approx 0$) and weights declining exponentially with distance from t_{start} via $e^{-\Delta t \,/\, \tau_{FPD}}$.

2.16 Native weekly, derived monthly/yearly and winteryearly time resolution

I fit the model always on weekly data. Afterwards, additional monthly and yearly values of observed, fitted and predicted mortality are derived, whose evaluation is thus based on native weekly optimization. For months, weekly data is accumulated by a uniform filter of ca 4.3 weeks, while retaining weekly time resolution. For years, mortality is accumulated over the 52 or 53 weeks attributed to each year, with yearly time resolution. Data handling would be easier if governments and institutes decided to provide public data on daily basis, to handle the awkwardly incompatible week/month/year overlaps.

Additional winteryearly resolution is introduced, denoted e.g. by 2016/17 meaning July 2016 to/including June 2017. Winteryears capture a single winter's mortality, relevant when measuring mortal events that focus on winters, such as influenza and covid. Normal calendar years mix up two halfs of consecutive winters, and thus always smooth out seasonal events over two years.

The more events are mixed together, the more they smooth out, resulting in lower model errors. Native weekly model fit errors are highest, followed by monthly, winteryearly and finally yearly fit errors. This order is natural, and not the primary factor in deciding which resolution to use.

The choice for a specific resolution depends on intrinsic relevance of the period, as it enables to quantify events within and localize events among these periods. Weekly and monthly resolution enable localization of mortality waves. Yearly resolution enables quantification of net effects of excesses and deficits, and connects to usual reporting of mortality. Winteryears, however, match better with generic seasonal mortality that peaks in winter. They are certainly preferential when they enclose a set of correlated events, e.g. an observed (naturally occurring) deficit in autumn followed by an excess in spring as occurred in The Netherlands in 2019/20.

Model prediction errors differ from fit errors as they are less predictable emergent properties of the world and the proposed model; they may show other behaviour.

2.17 Additional determinants and lifetime lost/saved

When significant excess mortality ΔM_{excess} is detected, one can expand the model with one or more additional determinants X:

$$M = W_j * a_{ij} t^i T^j + W_X * H(X) + \Delta M \tag{21}$$

$X(t)$	Some determinant
$H(t, X)$	Health-pressure function for determinant X
W_X	Determinant-specific HPA/FPD filter, two parameters $\tau_{HPA;X}$, $\tau_{FPD;X}$

Depending on the type of determinant, H may incorporate parametrized polynomials in t and X such as with temperature, or something completely different. Ideally, the additional model term reduces ΔM_{excess} such that $|Z| < 2$, thereby providing evidence that formerly-abnormal mortality becomes "normal" or expected when X is considered.

A major advantage of the proposed model is that lives lost LL and life*time* lost LTL related to X (negative quantities refer to lives and lifetime saved) are very easily determined:

$$\begin{aligned} LL_X &= \sum_t \left(H(X) - \mu_{H;X} \right) \\ LTL_X &= \tau_{FPD;X} LL_X \end{aligned} \tag{22}$$

The $\mu_{H;X}$ is the long-term mean health pressure of determinant X. For temperature, LL_T and LTL_T are typically positive in winter and negative in summer. The precise definition of $\mu_{H;X}$ is not important in this report, as I only use it for pandemic determinants of covid and its vaccines. For these, $\mu_{H;X}$ is zero, as both did not exist prior to the pandemic.

The determination of LTL_X requires that $\tau_{FPD;X}$ can be measured well, without degeneracy. In case of a merged degeneracy ($\tau_X = \tau_{HPA;X} = \tau_{FPD;X}$), the found τ_X may be regarded as a lower-bound for real $\tau_{FPD;X}$ and used to determine a lower-bound on lifetime-lost/saved. In case of a diverged degeneracy ($\tau_{HPA}\tau_{FPD} = \infty$), one can choose to enforce merged degeneracy in the model in order to obtain the lower-bounds on $\tau_{FPD;X}$ and LTL_X, or just resort to lives lost LL_X. In this report, I always determine an explicit 95% CI or CLB of LTL_X via the Bayesian analysis.

2.18 Pandemic determinants: covid tests and vaccinations

The model will be applied on the recent pandemic, to quantify mortality excesses and examine possible causes. During 2020-2023, several covid waves, non-pharmaceutical Interventions (NPIs) and vaccination campaigns took place. Determinants for such pandemic events will be directly integrated into the model, and parameters are optimized for both normal and pandemic-related mortality, thereby creating an exceptional situation for pandemic-related mortal events, namely that they are no longer abnormal. Instead, all that correlate with the determinants have become “normal”.

As there are no periods with other but similar pandemics, the model is fitted on the only pandemic at hand. The goal is not to predict mortality during the pandemic, but to explain observed mortality. A good explanation is equivalent with a model that fits well over the entire period of 2000-2023. Determinant-related parameters obtained in the fit represent pandemic knowledge gained.

My pandemic model includes determinants of positive covid tests C and vaccine dose volumes administered V. Stringency Indices (SI) of NPIs are not considered. Model parameters are covid Case Fatality Rate (CFR), Vaccine-dose Fatality Rate (VFR), and Vaccine Effectivity (VE) against covid mortality. Together, these capture mortal events by covid, the buildup and waning of natural immunity and vaccine-induced protection, and fatal adverse events by vaccination [Hul].

Viral waves and campaigns followed each other quickly during the pandemic, see Figure 9. While mortality waves are expected to follow viral waves, at various times all three of waves, campaigns and mortality coincided to some extent. Clearly, it will not be easy, by far, to disentangle effects among C, V and M, based on these public/aggregated data.

Source [Eur] provides age-stratified vaccination volumes V_l with dose/campaign number l, but some doses overlapped considerably in time (doses 1+2, and 4+5), while others had near-zero volumes (doses 6+7). For increased measurement significance, and to prevent spurious differential overfitting effects, I combined all 7 dose series in just 4 “campaigns” V_l with associated VFR_l: (1+2)/2, 3, 4+5 and 6+7. I averaged the primary series, as the two doses overlapped very strongly in time, and a single dose was in fact regarded as only half of full vaccination. Dose 6+7 data was only available in The Netherlands at the time of data acquisition.

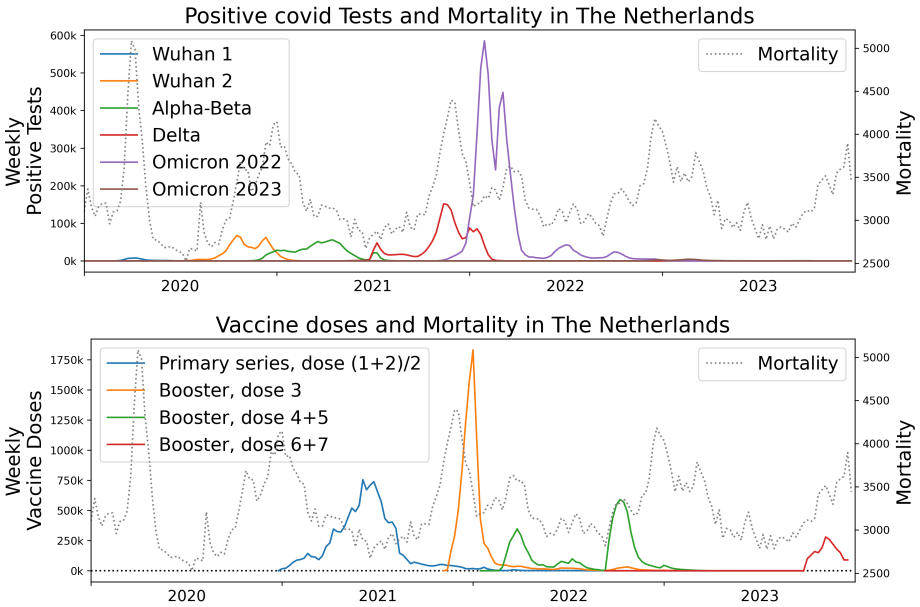

Figure 9: Weekly pandemic determinants for the Netherlands 2020-2023. Positive covid tests per half or whole year, vaccine doses per combined campaign, and mortality for reference.

Test case numbers C from [Eur] are not stratified to age or viral variant. When applied with age-stratified mortality data, I am relying on reasonable proportionality of case numbers over all age categories. Further, I split C into "waves" C_k based on calendar time; four half-year waves for pre-omicron, and two whole-year waves for omicron. This may not be the most accurate representation of reality, but it enables to capture dynamics and pathogenicity of viral variants, and especially the huge differences in testing volumes (e.g. the 1st wave had high mortality but few testing). I gave the periods smooth, slightly overlapping transition boundaries. I implemented this by applying a Gaussian (normal-distributed) smoothing filter with a width σ = 3 weeks on the boundaries of the periods. This leads to error-function-shaped transitions, with a zero net effect on total number of tests.

For some countries, positive tests data is so bad as to be become partly unusable, see Figure 10. During the 1st wave, tests were nearly zero in France. In Sweden, testing was minimal but went up when mortality went down. Both lead to infinite or strongly varying CFR over time during the first half of 2020, incompatible with a single parameter value. In the pandemic experiments, this will be handled by *excluding* the first half of 2020 in the analysis for France and Sweden. Rather than excluding the countries completely, this still enables the analysis of vaccine-related pandemic parameters. In FR and SE, all CFR-derived quantities such as lifetime lost to covid will thus exclude the 1st wave.

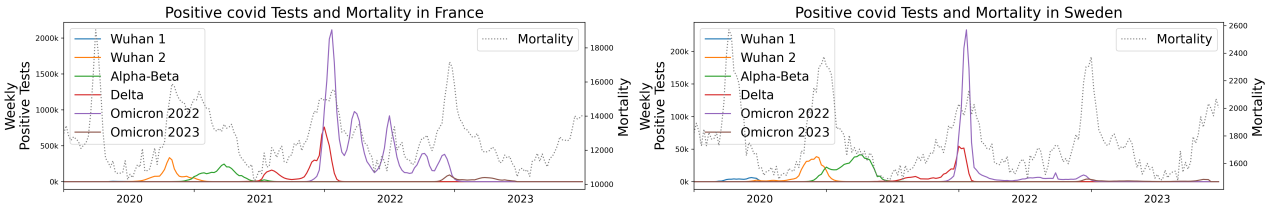

Figure 10: Positive test data in France and Sweden does not represent the 1st mortality wave well (blue line Wuhan 1).

All in all, the race to extract use out of unreliable data resulted in 10 time-series, 6 for covid tests C_k and 4 for vaccination campaigns V_l, and 2 countries requiring special attention.

My pandemic model adds two terms to the normal mortality model:

$$M = W_j * a_{ij} t^i T^j + W_C * CFR_k C_k (1 - VE_{pop}) + W_V * VFR_l V_l + \Delta M \tag{23}$$

$C_k(t)$	Positive covid tests per week, over time for wave/viral-variant k
$V_l(t)$	Absolute number of vaccine doses per week, over time for campaign l
CFR_k	Case Fatality Rate of wave/viral variant k for the *unvaccinated*
VFR_l	Vaccine-dose Fatality Rate for campaign l
W_C	Covid specific HPA/FPD filter, parameterized by $\tau_{HPA;C}, \tau_{FPD;C}$
W_V	Vaccination specific HPA/FPD filter parameterized by $\tau_{HPA;V}, \tau_{FPD;V}$
$VE_{pop}(t)$	Population-average vaccine effectivity over time against covid mortality

Both C and V are aggregated counts over the population which are unlikely to have a non-linear relation with health or mortality; therefore no quadratic effects via C^2 or V^2 are modeled.

Plain values measured for CFR are uninteresting as they are scaled in many ways; the definition of CFR combines relevant covid Infection Fatality Rate (IFR) with the subjectivity and irrelevance associated with test policies and public test-preparedness etc. Covid tests C_k typically undercount true cases by a considerable factor (especially in first and last waves). While IFR is limited to 100%, CFR absorbs the undercount factor, which could even lead to values above 100%. Additionally, covid test data C are not age-stratified, and when used with mortality of age-stratified subpopulations, CFR will absorb the proportion of subpopulation to entire population.

The role of CFR in this model is to represent covid mortality via the productsum $CFR_k C_k$ (in which all above factors cancel), derive lifetime lost, enable VE measurements, and prevent confounding the measurements of VFR.

The CFR captures the time-varying aspects of covid IFR that depends on viral variant, as well as the buildup and waning of natural immunity, all in the unvaccinated population. Without explicit VE modeling, CFR would refer to the entire population and automatically incorporate population-average vaccine-induced protection, albeit at a rough temporal resolution of one fixed value per half/whole year for the entire population. The values measured for CFR would then be lower, assuming $VE > 0$.

My VE model captures the start of protection after vaccination as well as its continuous decline over time *per vaccinated person*, via population-average vaccine effectivity $VE_{pop}(t)$, implemented by default linear filter methods:

$$\begin{aligned} VE_{pop} &= \min\left(\frac{V * VE_{ind}}{P}, 1\right) \\ VE_{ind}(t_V) &= VE_0 e^{-t_V/\tau_{VE}} \end{aligned} \tag{24}$$

$VE_{pop}(t)$	Population-average vaccine effectivity, over calendar time t
$V(t)$	Absolute number of vaccine doses given over time (not cumulative), for all campaigns

$P(t)$	Total number of people in the population over time
$VE_{ind}(t_V)$	Vaccine effectivity for an individual, over *time-since-vaccination* t_V
t_V	Individual time-since-vaccination
VE_0	Initial individual VE, right after vaccination, at $t_V = 0$
τ_{VE}	Time scale of VE waning over time-since-vaccination

Although VE_{ind} may start out high (e.g. VE_0 near 100% was often reported), it wanes quickly over time [Wu], resulting in the rollout of no less than 5 booster campaigns in just 2 years. I have modeled the waning by a simple negative exponential curve with a characteristic time scale τ_{VE}.

This VE model allows for the representaton of so-called “system dead-time”; a VE_0 value *higher than 1* in (24) produces a VE_{pop} that remains saturated at 100% for some time and only then starts dropping, see Figure 11. I make use of this by allowing $VE_0 > 1$ in the model. When $VE_0 > 1$, the initial period during which the vaccine’s effectivity remains at 100% is:

$$\tau_{VE-100\%} = \tau_{VE} \ln VE_0 \tag{25}$$

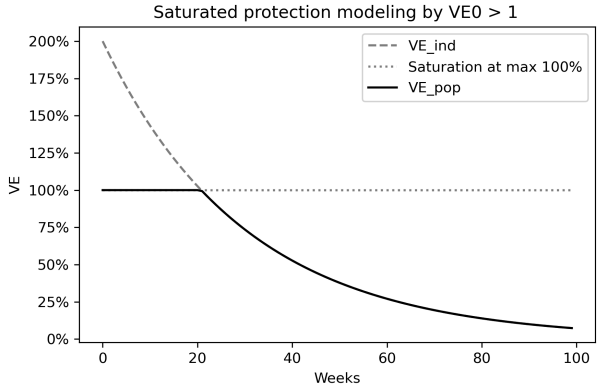

Figure 11: The model can represent a population-wide VE_{pop} that remains at maximum protection of 100% for some time and then starts waning.

My VE model is inevitably approximate and does not capture all kinds of intricacies, such as delayed protection due to immunity build-up after vaccination, and individual VE saturating at 100%, relevant e.g. when a booster dose adds 100% to whatever protection is remaining of the previous dose. My model saturates protection at population-level only.

Similar to the normal mortality model, VE_0 and five τ’s are non-linear parameters that can only be measured by a brute-force numerical method. The CFR_k and VFR_l are linear parameters similar to a_{ij} of the temperature model. The CFR_k, VFR_l and VE_0 are physically restricted to 0-100% in reality, but in the context of this model the upper bound need not hold or be enforced as explained for CFR_k and VE_0. The VFR_l is expected to be very low and to remain below 100% automatically without explicit model restrictions. Essentially, only the zero-lower-bound of 0% remains relevant, and must be enforced by the model’s prior parameter probabilities, similar to the non-linear parameters of the temperature model. Brute force numerical methods can directly enforce the zero-lower-bound. If all non-linear parameters are fixed to default values, and one resorts to linear regression to find all linear parameters, repeated linear regression can be used. If

a spurious negative CFR_k or VFR_l is encountered, the associated C_k or V_l is successively removed, effectively forcing the CFR_k or VFR_l to zero, and linear regression is repeated.

Although the model's zero-lower-bound on CFR or VFR represents real physical constraints, one can rightfully argue that such a bound creates a bias, transforming zero-mean randomness into net-positive values. Therefore, an experiment will be included without the zero-lower-bound.

From the pandemic parameters one can derive lives and lifetime lost to covid and vaccinations, and saved by vaccine-protection:

$$\begin{aligned} LL_C &= \sum_t C_k CFR_k (1 - VE_{pop}) & LTL_C &= \tau_{FPD;C} LL_C \\ LS_{VE} &= \sum_t C_k CFR_k VE_{pop} & LTS_{VE} &= \tau_{FPD;C} LS_{VE} \\ LL_V &= \sum_t V_l VFR_l & LTL_V &= \tau_{FPD;V} LL_V \end{aligned} \tag{26}$$

Despite all my disclaimers on data reliability and model limitations, the pandemic model captures the main features of covid and vaccines. It has 16 parameters, 10 linear (6 CFR_k, 4 VFR_l) and 6 non-linear (VE_0, τ_{VE}, $\tau_{HPA;C}$, $\tau_{FPD;C}$, $\tau_{HPA;V}$, $\tau_{FPD;V}$). For pandemic years, this is 1 additional parameter on average per quarter-year during a 4-year period with highly abnormal mortality dynamics. The full pandemic model including time and temperature has 28 parameters, and will be used to fit mortality of 24 years from 2000-2023 covering 1252 weeks, on average ca 1 parameter per 45 weeks.

Parameters are distributed non-uniformly over pandemic and prepandemic periods, at ca 4.5 and 0.5 parameter per year respectively. A more balanced situation would be if e.g. prepandemic influenza tests and vaccinations were also included, but that is beyond the scope of this report. The non-uniformity may make the model fit *better* in the pandemic period versus the prepandemic period. When the model is evaluated on weekly/monthly/yearly time scales, the number of mortality values is 52/12/1 per year. At 4.5 parameters per pandemic year, the model may show signs of overfitting when evaluated at derived yearly resolution, but this depends on characteristics of the pandemic determinants.

I will experiment with various pandemic models: CFR-only, VFR-only, CFR+VE, etc, and the full pandemic model CFR+VE+VFR.

3. Results

Next follow a description of datasets and five sets of experiments. First, model variants in terms of polynomial order, use of HPA/FPD and population-relative mortality are examined and a selection is made. Secondly, the proposed model is fit to normal mortality during prepandemic years 2000-2019. Thirdly, prediction error bandwidth factors are determined. Fourthly, pandemic predictions are made for years 2020-2023 to measure excess mortality. Finally pandemic determinants are added to the model to assess possible excess causes.

3.1 Datasets

Table 1 provides an overview of the 14 datasets used; the "4-NL" datasets with age-stratified mortality and hourly temperature data, and the "10-EU" datasets with all-age mortality and weekly temperature data. As FR and IT are time-limited, several experiments requiring two full decades of prepandemic data will exclude these countries, referring to "8-EU datasets". The FR and SE datasets have unusable 1st wave covid test data, requiring special handling.

Dataset	Country	Ages	Population	Mortality data years	Effective nr. of weather stations	Weather time resolution	Covid tests 1st wave
NL 0-65	The Netherlands	0-65	14M	2000-2023	33.7	hour	
NL 65-80		65-80	3M				
NL 80+		80+	1M				
NL ALL		All	18M				
BE	Belgium	All	12M	2000-2023	**1.0**	day	**near-zero**
BG	Bulgaria		6.5M	2000-2023	3.0		
DE	Germany		83M	2000-2023	2.6		
ES	Spain		48M	2000-2023	8.6		
FI	Finland		5.5M	2000-2023	2.0		
FR	France		65M	**2011-2023**	7.0		
IT	Italy		59M	**2013-2023**	9.5		
NL	Netherlands		18M	2000-2023	2.3		
PL	Poland		37M	2000-2023	4.0		
SE	Sweden		10.5M	2000-2023	4.6		**inconsistent**
Total	10 countries		344M				

Table 1: All 14 data sets, grouped in 4-NL finegrained datasets, and 10-EU country datasets. Data limitations in bold. The datasets named 8-EU are the 10-EU sets minus FR and IT.

For the Netherlands, data was obtained for years 2000-2023 from [Cbs2,Knm]:

- Weekly absolute mortality for age categories 0-65, 65-80, 80+ and all-ages
- Yearly population size at 1 January for the same age categories
- Dutch national mortality baseline, weekly prediction 2020-2023 all-ages only, version based on prepandemic conditions (the baseline was modified later including pandemic conditions)
- Hourly averages of temperature in 50 subregions

For 10 European countries, the following was obtained from [Eur, Noa, Riv2] and google:

- Observed absolute weekly mortality, all-ages, 2000-2023 (FR and IT from 2011 & 2013 up)
- Daily average and/or min-max temperature, with limited (1 to 10) weatherstations per country
- Weekly number of positive covid tests (all ages only)
- Weekly number of vaccine doses administered (all four age categories)
- Population size at 1 Jan 2022 only

Yearly population sizes were interpolated/extrapolated to weekly values. Weather data was averaged to weekly, whole-population datapoints without regional weighting. When only min/max temperatures were available, average of min/max was used.

3.2 Model variants and selection

The proposed model comes in several variants by polynomial order of time and temperature, the use of HPA/FPD (none, merged, full), and absolute vs population-relative mortality. Based on

native-weekly fit and prediction errors, I will make a selection. Fit accuracy is determined as average over all datasets, with the model fit on prepandemic years 2000-2019 (or as data permits for FR and IT). Prediction accuracies are averaged over 12 datasets (all minus data-limited FR and IT) and over forward/backward 16y-to-4y extrapolation ('00-'15 to '16-'19 and '04-'19 to '00-'03).

Table 2 shows prediction accuracy for all model variants, applied on absolute mortality. Observations are:

- Unanimously, the optimal polynomial order of time is $O_t = 2$, consistent with long-term mortality developments defined in this report, with $L \geq$ 20y (total duration of this experiment)
- The optimal order of temperature is $O_T = 3$, while $O_T = 2$ performs nearly equal
- The use of HPA/FPD lowers prediction error substantially, from ca 7.4% to 6.4% (rel. ca -14%)
- Merged and full HPA/FPD perform practically the same

Average prediction error $\mu_{E_{prediction}}$ [%]

	no HPA/FPD					*Merged HPA/FPD*					*Full HPA/FPD*				
	T^0	T^1	T^2	T^3	T^4	T^0	T^1	T^2	T^3	T^4	T^0	T^1	T^2	T^3	T^4
t^0	12	10	10	9.7	9.7	12	9.5	9.2	9.2	9.2	12	9.5	9.0	8.9	8.9
t^1	12	9.7	9.6	9.2	9.2	12	8.6	8.3	8.2	8.4	12	8.6	8.2	8.2	8.3
t^2	10	7.6	7.5	7.3	7.6	10	6.9	6.49	6.34	7.8	10	6.8	**6.46**	6.36	7.6
t^3	11	9.0	10	9.1	9.5	11	8.5	7.7	9.9	12	11	8.5	10	11	13

Table 2: Average prediction error (native weekly) for model variants, by time and temperature's polynomial order and the use of HPA/FPD. In grey relevant variants, in bold the proposed model.

The four best variants, merged vs full HPA/FPD and $O_T = 2$ vs 3, perform *nearly the same* in terms of mortality prediction. The choice for any of these variants depends fully on whether one has an interest in measurements of parameters τ_{HPA}, τ_{FPD}, and whether such measurements are feasible with sufficient accuracy.

Table 3 shows model fit errors for the four best variants, and estimated non-linear parameters in terms of mean and deviation over all datasets (excluding diverged τ_{FPD}).

	Merged HPA/FPD						**Full HPA/PFD**							
	$O_T = 2$			$O_T = 3$			$O_T = 2$				$O_T = 3$			
Dataset	α [%]	τ [w]	E_{fit} [%]	α [%]	τ [w]	E_{fit} [%]	α [%]	τ_{HPA} [w]	τ_{FPD} [w]	E_{fit} [%]	α [%]	τ_{HPA} [w]	τ_{FPD} [w]	E_{fit} [%]
4-NL	3.9 ±1.4	9.7 ±3.1	5.3 ±0.8	3.7 ±1.4	9.4 ±3.0	5.2 ±0.8	3.9 ±1.4	4.7 ±1.0	28 ±16	5.3 ±0.8	3.7± 1.3	4.0 ±1.5	29 ±13	5.1 ±0.8
10-EU	4.6 ±0.7	12 ±0.9	5.1 ±0.7	4.6 ±0.8	12 ±0.8	5.0 ±0.7	4.5 ±0.6	4.7 ±1.4	43 ±21	4.9 ±0.6	4.4 ±0.6	5.0 ±2.0	43 ±22	4.8 ±0.7
all 14 datasets	4.4 ±1.0	11 ±1.9	**5.1 ±0.7**	4.3 ±1.0	11 ±2.0	**5.0 ±0.7**	4.3 ±0.9	4.7 ±1.3	38 ±20	**5.0 ±0.7**	4.2 ±0.9	4.8 ±1.8	38 ±20	**4.9 ±0.7**

Table 3: Prepandemic 2000-2019 model fit errors (native weekly), non-linear parameters α and HPA/FPD times, by mean and deviation over the given datasets. In bold the model fit error.

Observations are:

- Fit errors E_{fit} and remaining mortality variability α are nearly equal for all model variants, and decrease very slightly with the model's number of parameters: expected and irrelevant
- Across all-age datasets, α and merged HPA/FPD τ are consistent at ca 4.5% and 12w
- Full HPA/FPD τ_{HPA} and τ_{FPD} are more variable over datasets than τ, regardless of O_T

The four variants perform *practically the same* in terms of model fit errors. The accuracy of HPA/FPD parameters (in terms of inter-dataset-variability) does not significantly differ between $O_T = 2$ or $O_T = 3$, so at this point I choose $O_T = 2$ as proposed in the introduction. The minor prediction improvement does not seem worth the 3 extra parameters. And I do like the symmetry $O_t = O_T$.

Given the consistent value $\tau \approx 11$ weeks for merged HPA/FPD, it is interesting to see whether such consistency applies also to for τ_{HPA}, τ_{FPD}, i.e. whether their strong variability in Table 3 is only an issue of accuracy/degeneracy. If so, generic values may be applicable in general. Table 4 show eight model variants, including those with non-linear parameters fixed to results of Table 3.

	Model variants							
	Constant (ref) t^0T^0	Time only t^2T^0	Temp. only t^0T^2	Time+ Temp. t^2T^2	Merged HPA/FPD	**Full HPA/FPD**	Merged HPA/FPD $\alpha = 4.4\,\%$ $\tau = 11\text{w}$	Full HPA/FPD $\alpha = 4.3\,\%$ $\tau_{HPA} = 4.7\text{w}$ $\tau_{FPD} = 38\text{w}$
# linear parameters	1	3	3	9	9	**9**	9	9
# nonlinear parameters	0	0	0	0	2	**3**	0 (fixed)	0 (fixed)
$\mu_{E_{fit}}$ [%]	10.5	8.5	7.9	6.4	5.13	**5.03**	5.23	5.14
wrt ref	-	*-19%*	*-25%*	*-40%*	*-51%*	***-52%***	*-50%*	*-51%*
$\mu_{E_{prediction}}$ [%]	12.2	10.1	10.0	7.5	6.49	**6.46**	6.46	6.40
wrt ref	-	*-17%*	*-18%*	*-39%*	*-47%*	***-47%***	*-47%*	*-47%*

Table 4: Model fit and prediction error for model variants. Note the marginal differences in the four variants at right. The two rightmost variants have non-linear parameters fixed to generic values; all 9 remaining linear parameters can be fit with default linear regression. In bold the proposed model.

All four models that include HPA/FPD in whatever form perform more or less the same. Out of interest in FPD values, I select the model with full HPA/FPD, with 12 parameters, 3 non-linear and 9 linear. The two models with non-linear parameters fixed to generic values can be fitted fully with default linear regression methods and predict mortality *just as good* as the proposed model.

Finally, Table 5 shows results of the model applied on absolute versus population-relative mortality. These results are preliminary, as only 4-NL datasets were available for use. Also, relative mortality results may be sub-optimal as yearly population sizes were interpolated to weekly.

	Absolute/Relative Mortality			
	Native weekly		Derived Yearly	
Dataset	E_{fit} [%]	$E_{prediction}$ [%]	E_{fit} [%]	$E_{prediction}$ [%]
NL 0-65	5.2 / 5.2	5.5 / 5.3	1.2 / 1.0	2.2 / 1.6
NL 65-80	4.8 / 4.9	6.1 / 5.9	1.4 / 1.3	2.4 / 2.7
NL 80+	6.4 / 6.6	9.0 / 8.6	1.6 / 1.4	3.0 / 3.8
NL ALL	4.6 / 4.6	6.2 / 6.2	1.2 / 1.2	2.1 / 2.9
Average	5.3 / 5.3	6.7 / 6.5	1.3 / 1.2	2.4 / 2.8
E Rel vs Abs	+1%	-3%	-7%	+13%

Table 5: Fit and prediction results for absolute vs relative mortality, only for 4-NL datasets.

These results are not simple and unanimous, differing over time resolution, age groups, and fit vs prediction errors. Tentative observations are:

- Native weekly fit and prediction accuracy seem indifferent to absolute or relative mortality, with a slight preference for relative mortality in prediction for age-stratified populations
- Derived yearly predictions using relative mortality seem substantially better for ages 0-65 (-27%), but substantially worse for 65-80, 80+ and all-ages (+13%, +27% and +38%).

This may relate to migration occurring mostly in younger, more healthy people, while mortality occurs in older, more frail people.

Figure 12 shows the population size for ages 0-65 in The Netherlands. From 2003 to ca 2021, the population is quite stable within ca. 1%, but before and after, a “disruptive” event occurs with dynamics stronger than my parabolic long-term developments model can represent.

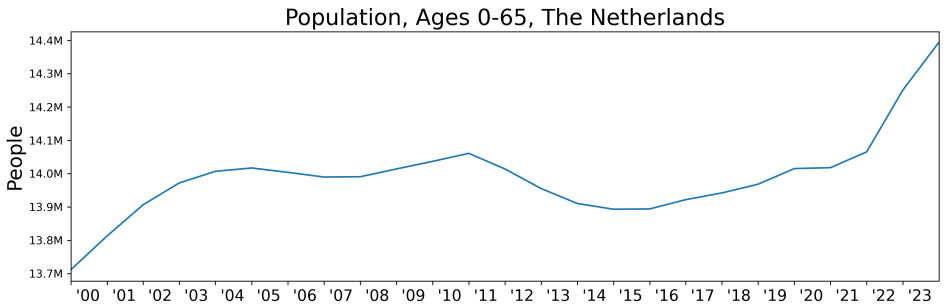

Figure 12: Population growth in The Netherlands, ages 0-65.

The 2000-2003 period is used in Table 5, strongly determining the results. From 2022 to 2023, during the pandemic, a disruptive event occurs of accelerated net immigration of ca. 3%, from 14M in January 2022 to 14.4M in December 2023. Net immigration in The Netherlands was ca. +0.6% during 2016-2021 to ca. +1.6% in 2022 and ca. 1% in 2023 [Cbs4].

These results are in line with expectations discussed in the introduction, and certainly need more examination with additional datasets. If datasets were available that are highly age-stratified (e.g. 5 year bins) and come with weekly-updated population tables, clearly population-relative mortality is always the better choice; such datasets, however, are not (publicly) available.

For the remainder of this report, I will use absolute mortality on all datasets, except for NL 0-65 for which I use population-relative mortality. For compatibility with derived quantities that use absolute numbers (e.g. lifetime lost/saved), I will also use the NL 0-65 dataset with a standardized

population over time, that is, scale its dynamic population to be constant over time, equal to its mean of ca 14M people (and rescale all other relevant data such mortality and pandemic determinants similarly).

3.3 Prepandemic model fit

The proposed model is fitted on all datasets during prepandemic years 2000-2019 (1043 weeks), or winteryears 2000/01-2018/19 (991 weeks) or as data permits (for FR and IT). Figure 13 shows results for NL ALL, with fit error 4.6% at native weekly resolution, and 3.5%, 1.6% and 1.2% at derived monthly, winteryearly and yearly time resolution. This confirms the ordering of fit error magnitudes as a function of resolution.

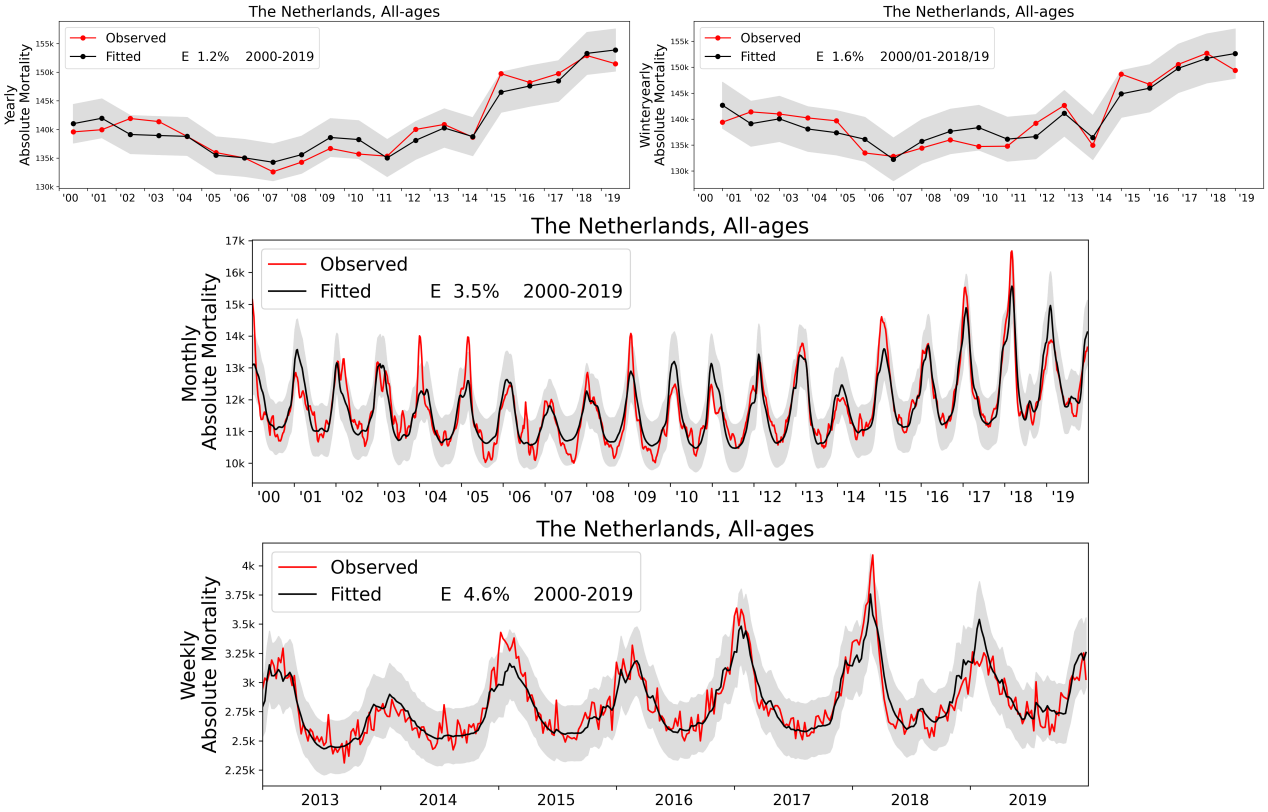

Figure 13: Prepandemic 2-decade model fit for dataset NL ALL at derived yearly/winteryearly/ monthly resolutions, and native weekly resolution (zoomed in on 2013-2019). Bandwidths are 95% confidence intervals (± 2 E_{fit} of fitted mortality).

At both yearly and winteryearly resolutions, the overall parabolic trend and several local details such as the dip/rise in 2014/15 are well-represented by the model. At monthly and weekly resolution, seasonal variations are followed well, such as the mild winter 2013/14 and the strong winter 2017/18. This is due to the temperature part of the model as both observed and modeled mortality for all winters do not follow a simple quadratic trend.

At various occasions, the model does miss some events as winterpeaks and summerdips, see the monthly result. Also, the weekly result shows that my model still does not explain the strong weekly variance in observed mortality, one of the goals discussed in the introduction. This fit result has an α of ca 4%, the parameter for mortality's natural variability. This is higher than the 3% found in [Red4], in accordance with still unexplained weekly variance plus a few missed events that [Red4] accounted for automatically (that method is not applicable for mortality prediction).

Figure 14 shows the same weekly result but with mortality relative to population size. The model fit error is essentially the same, 4.59% for relative mortality versus 4.58% for absolute

mortality. At two decimals precision, relative-mortality fit errors are 4.6%, 3.6%, 1.5% and 1.2%, nearly equal to those of absolute-mortality at all four time resolutions, in line with previous results (Table 5).

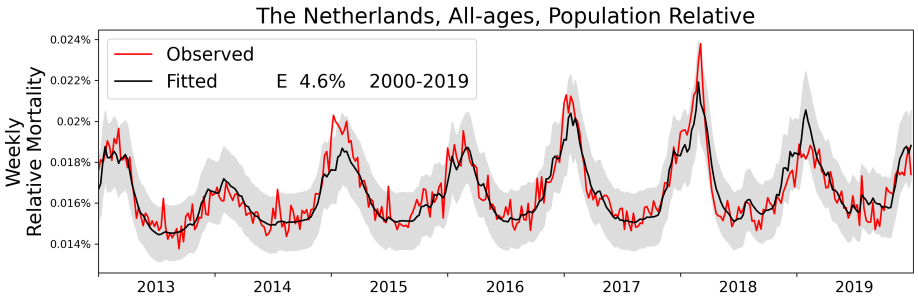

Figure 14: Same result as Fig. 10 with the model applied on mortality relative to population size.

Figure 15 illustrates a few results with model variants time-only, temperature-only, time and temperature but still without HPA/FPD, and the proposed model.

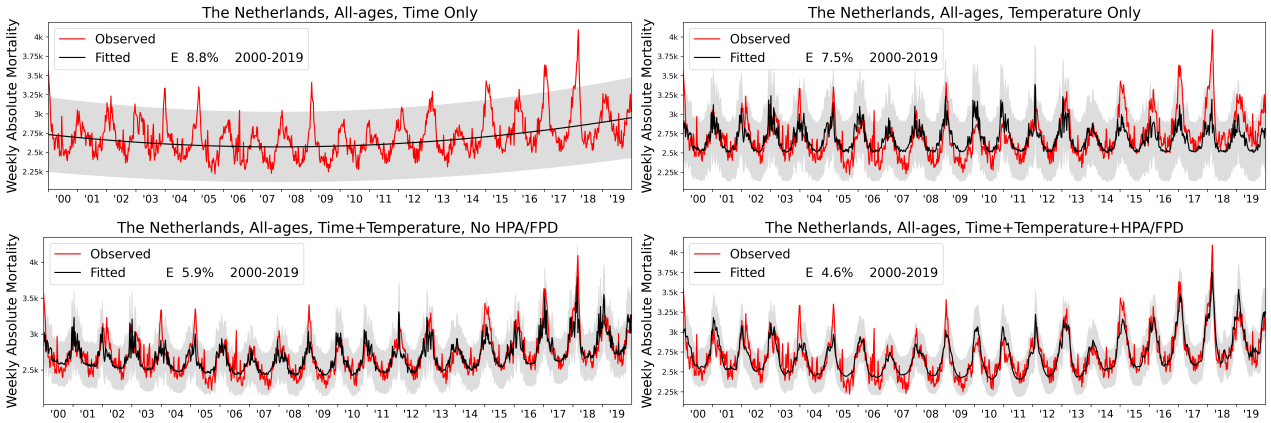

Figure 15: Illustrative fit results for three partial models and the proposed model at bottom-right.

Table 6 and Figures 16-17 show results for all datasets. In general, the model seems to fit reliably similar for different age-groups, countries, and data quality. There is practically no difference in the results for datasets NL ALL vs NL (hourly vs weekly temperature data, 34 vs 2 weather stations), and BE vs ES (1 vs 9 weather stations). Model fit errors are highest at week resolution and lowest at year resolution: winteryears consistently mix mortal events less than calendar years.

The yearly result for DE (Germany) has a fit error of 1.1%, which can be compared to the quadratic trend-only result in Figure 1 with an error of ca. 1.7% (±14.5k per ca. 870k deaths per year). The proposed model outperforms a quadratic trend-only model thus strongly (ca. 34% error reduction, ca. 56% variance), even when not natively fit on years but on weeks.

Table 6 also shows results for nonlinear parameters. On average over all datasets, $\alpha \approx$ 4.3±0.9 %, $\tau_{HPA} \approx$ 4.7±1.3w and $\tau_{FPD} \approx$ 38±20w. The α and τ_{HPA} are stable to some extent over all datasets, but τ_{FPD}'s 95% CI or CLB range per dataset is substantial, and so is its variability over datasets. The 4-NL results shows τ_{FPD} rises with age, suggesting that the few frail among the young are frailer than the many frail among the elderly.

Dataset	E_{fit} [%]				Nonlinear parameters		
	weekly (native)	monthly (derived)	winteryearly (derived)	yearly (derived)	α [%]	τ_{HPA} [w]	τ_{FPD} [w]
NL 0-65	5.2	2.8	1.2	1.0	1.9-2.7	2.8-6.3	5.5-11
NL 65-80	4.8	3.1	1.4	1.4	3.1-3.6	3.0-6.6	12-34
NL 80+	6.4	5.1	2.2	1.6	5.4-6.0	5.4-7.0	27-45
NL ALL	4.6	3.5	1.6	1.2	3.9-4.3	4.0-5.0	33-65
BE	5.4	4.1	1.6	1.6	4.7-5.1	4.4-6.1	29-64
BG	5.9	4.2	1.8	1.6	5.2-5.7	2.9-3.7	>280
DE	4.4	3.5	1.4	1.1	4.0-4.4	5.4-7.0	32-61
ES	5.5	4.5	1.8	1.6	4.9-5.4	3.3-3.9	62-120
FI	5.3	3.9	1.3	1.1	3.9-4.5	3.7-6.2	>43
FR	3.7	2.8	0.8	0.6	3.3-3.8	3.8-5.4	23-47
IT	5.0	3.8	1.1	1.3	4.6-5.4	2.0-2.8	>68
NL	4.7	3.6	1.6	1.2	4.0-4.4	4.0-5.5	29-49
PL	4.8	3.7	1.4	1.4	4.4-4.8	4.0-7.3	19-76
SE	4.5	3.5	1.6	1.0	3.4-3.9	6.0-9.2	13-22
All 14 datasets	5.0±0.7	3.7±0.6	1.5±0.3	1.3±0.3	4.3±0.9	4.7±1.3	38±20

Table 6: Prepandemic model fit errors. Fit period is 20 calendar years (2000-2019), except for winteryearly resolution with 19 winteryears (2000/01-2018/19). Parameters are given by 95% CI, or CLB if diverged. Statistics over datasets are average and deviation of dataset's mean parameter values (not shown, only included if not diverged).

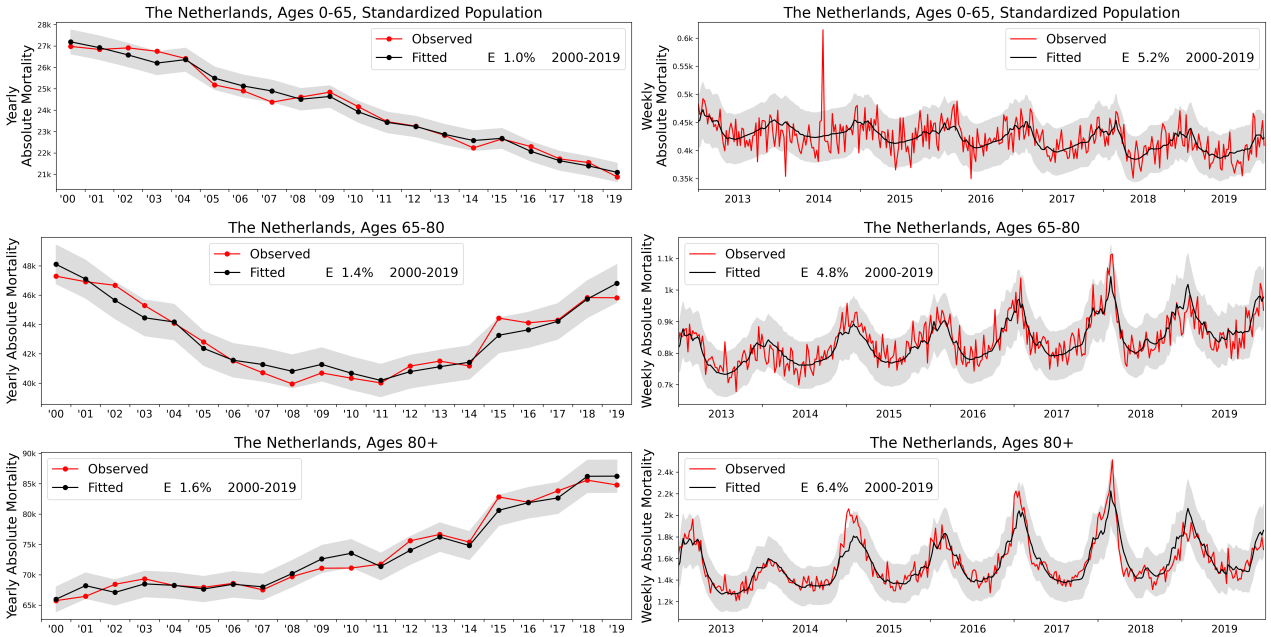

Figure 16: Prepandemic 2000-2019 model fits for age-groups in The Netherlands, datasets NL 0-65, 65-80 and 80+. The peak in 2014 for ages 0-65 is the MH17 airplane incident.

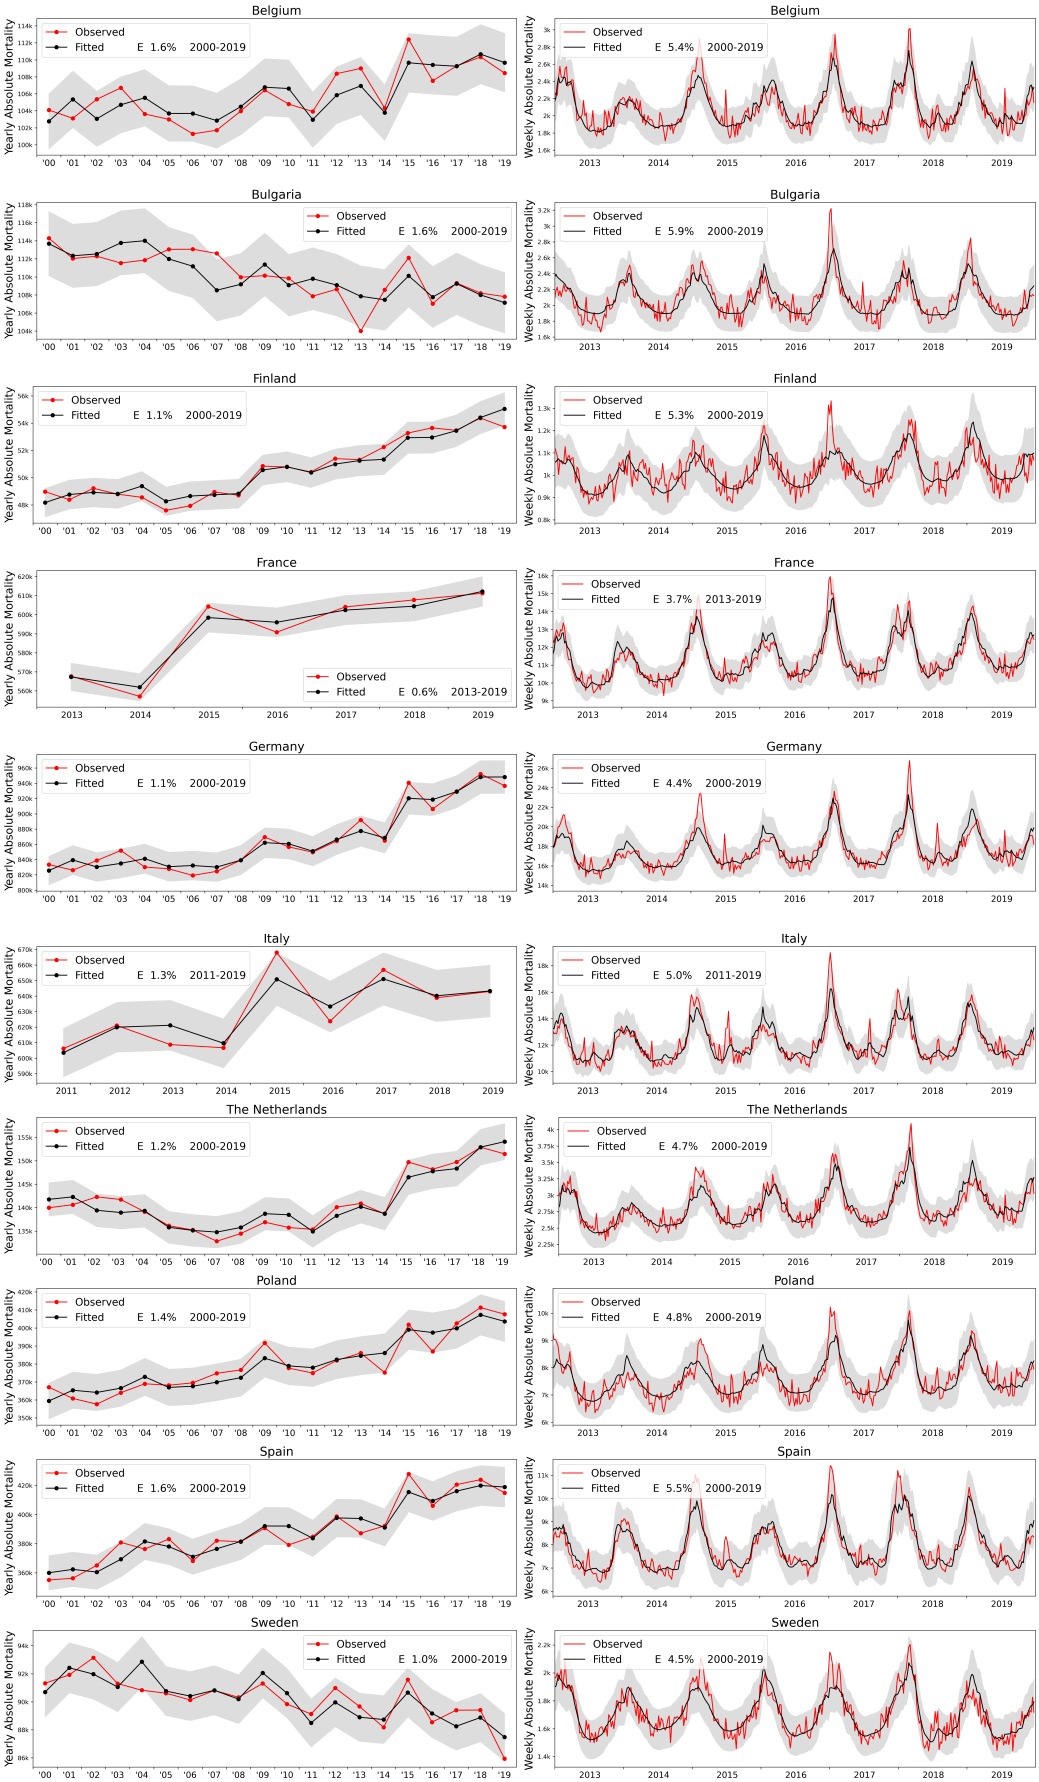

Figure 17: Prepandemic 2000-2019 model fits for 10-EU datasets, shown at derived yearly resolution, and native weekly (zoomed in 2013-2019).

3.4 Prediction bandwidth: fit to prediction amplification

Table 7 provides fit-to-prediction error amplification factors μ_A for 1-4 (winter)year extrapolation and interpolation during prepandemic periods. As no major unexpected events took place in this period, the fit-to-prediction error amplification A and its average μ_A are representative for prediction of normal mortality.

Type of Prediction 4 out of 20 calendar years or 19 winteryears		μ_A			
		Weekly	Monthly	Winter yearly	Yearly
Extrapolation					
year 1	$A_{\geq 15 \rightarrow 1}$	1.1	1.2	1.1	1.5
year 2	$A_{\geq 15 \overset{1}{\rightarrow} 1}$	1.2	1.3	1.2	1.9
year 3	$A_{\geq 15 \overset{2}{\rightarrow} 1}$	1.3	1.5	1.8	2.3
year 4	$A_{\geq 15 \overset{3}{\rightarrow} 1}$	1.4	1.6	2.4	3.2
all 4 years	$A_{\geq 15 \rightarrow 4}$	1.3	1.4	1.9	2.6
Interpolation					
year 1	$A_{15 \rightarrow 1 \overset{3}{\leftarrow} 1}$	1.2	1.3	-	1.6
year 2	$A_{15 \overset{1}{\rightarrow} 1 \overset{2}{\leftarrow} 1}$	1.1	1.2	-	1.4
year 3	$A_{15 \overset{2}{\rightarrow} 1 \overset{1}{\leftarrow} 1}$	1.2	1.3	-	1.4
year 4	$A_{15 \overset{3}{\rightarrow} 1 \leftarrow 1}$	1.4	1.6	-	1.5
all 4 years	$A_{15 \rightarrow 4 \leftarrow 1}$	1.2	1.4	-	1.7

Table 7: Fit-to-prediction error amplification A for extrapolation in the prepandemic period (years 2000-2019 or winteryears 2000/01-2018/19), for 12 datasets (all except time-limited FR and IT). Interpolation results are for reference (winteryear results not computed).

The μ_A for 4-year extrapolation and interpolation at yearly time resolution are 2.6 and 1.7 respectively. Clearly, interpolation with even a single year at the other end of the prediction period strongly reduces μ_A. In this report, interpolation results are illustrative and become relevant only when the first postpandemic year becomes available. Pandemic predictions made in the next subsection only use prepandemic extrapolation.

All extrapolative μ_A range from 1.1 to 3.2, and scale consistently with extrapolation distance from the fit period. Completely opposite to fit errors, μ_A are lowest for weekly resolution and higher for monthly, winteryearly and yearly resolution. For 4-year extrapolative predictions, yearly $\mu_A \approx 2.6$ is substantially above winteryearly's $\mu_A \approx 1.9$, while yearly $E_{fit} \approx$ 1.3 is far lower then winteryearly's $E_{fit} \approx$ 1.5 (Table 6). Their product $\mu_A E_{fit}$ is the relevant estimated error in prediction $E_{prediction;estimated-for-normal-mortality-only}$, ca. 3.4% for normal years and ca 2.8% for winteryears. The use of winteryears for mortality predictions provides a relative error reduction of ca. 18% over calendar years.

3.5 Prediction from prepandemic to pandemic years: excess mortality

Figure 18 shows the pandemic prediction by the proposed model versus the Dutch national baseline, both using only prepandemic mortality observations, all-ages, in The Netherlands. By now, an adapted baseline has been introduced that does include pandemic aspects; useful for examination of later endemic/postpandemic years but not for determination of pandemic excess mortality during 2020-2023.

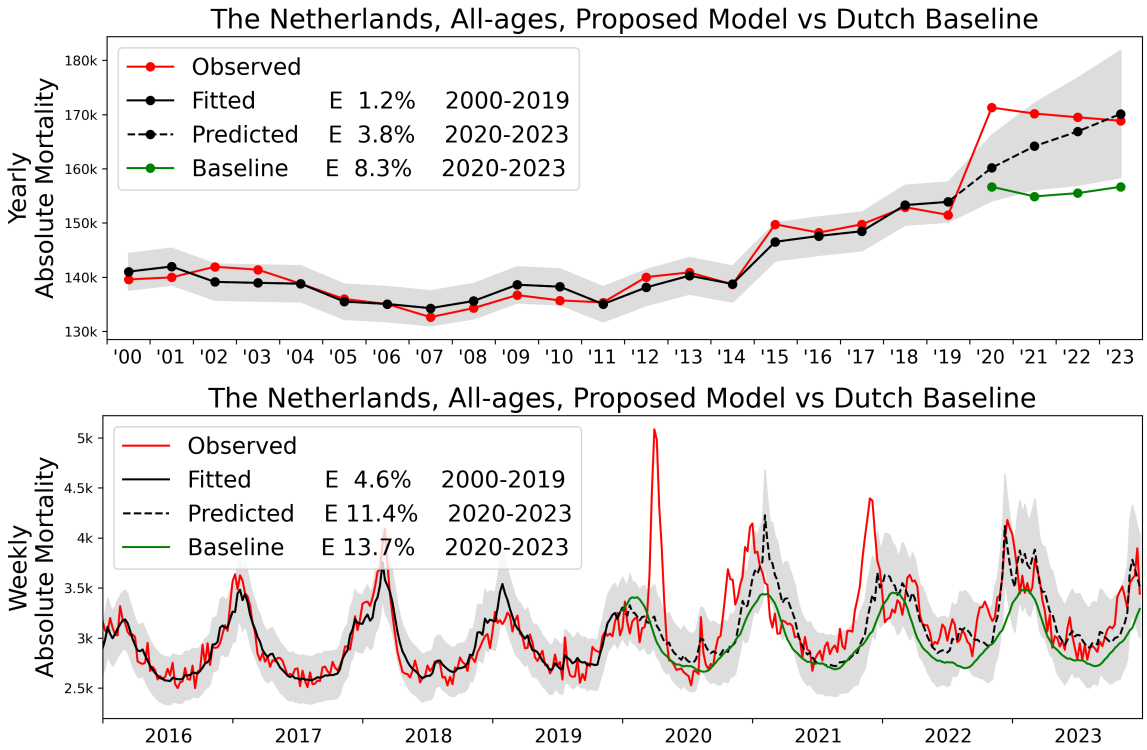

Figure 18: Comparison of proposed model (dataset NL ALL) and Dutch national baseline; both pandemic predictions are based on prepandemic mortality observations.

At yearly resolution, the baseline's long-term behaviour seems effectively based on a linear extrapolation of earlier years, or maybe even lower than that. This results in substantial derived excess mortality, similar to the linear trend results for Germany in Figure 1. My model uses parabolic extrapolation, so the yearly results diverge with every year further from 2020; clearly the proposed model is much closer to observations, even spot-on in 2023.

A close look at weekly resolution shows that the seasonal effect in the baseline is a fixed pattern, the same for every predicted year. The pattern is in fact a smoothed average seasonal pattern of the past five years (2015-2019), subsequently rescaled per year according to a complex long-term trend analysis [Cbs1,Cbs3]. My model uses the actual weekly temperature in prediction years and therefore varies every year, still missing some but also capturing some observed events.

In march 2020, the 1st mortality wave occurred, not predicted by both baseline and model. Right after the wave, mortality was close to baseline prediction, but lower-than-predicted by the proposed model. In this case, the latter prediction is better in the sense that deficits always follow excesses. In summer 2020, a heat wave occurred, which is partly captured by the proposed model but not by the baseline. In winter 2020/2021, the 2nd wave was missed by the baseline, but highly similar increased mortality was in fact predicted by the proposed model, albeit a few months later. In that perspective, the 2nd wave did accelerate mortality slightly causing loss of a few life-months per death, but its net excess mortality over the entire winter is considerably less than suggested by the baseline. At the second half of 2021, a 3rd mortality wave occurred that

was not predicted by baseline and model. In years 2022 and 2023, no clear mortality waves occurred and the new model just captured events more accurate than the baseline, both in terms of yearly trend and winterpeaks.

On yearly time resolution, the prediction errors for proposed model and baseline are 3.8% and 8.3% respectively during the pandemic; the model yields a more-than-double improvement in prediction error, or a ca. fivefold variance reduction (ca 79% lower).

Table 8 lists excess mortality numerically in The Netherlands according to the proposed model and the Dutch baseline. The differences are substantial, with total excess measured as 18k by the model and 56k by the baseline. The proposed model finds a highly significant excess of 11k in 2020 with $Z \approx 3.5$. Other years, or the entire pandemic period, do not show significant excess.

Note that the approximation $E_{prediction} \approx |\Delta M_{excess}/\mu_M|$ presented in section 2 holds very well for individual years; for multiple years the yearly variability adds to $E_{prediction}$.

Prediction period	$E_{prediction}$ [%]		Absolute mortality in period			Excess according to proposed model		
	Model	Baseline	Observed μ_M	Excess baseline	Excess model ΔM_{excess}	$\frac{\Delta M_{excess}}{\mu_M}$	$\mu_A E_{fit}$	Z
2020	6.5	8.5	171k	15k	**11k**	**+6.5%**	1.9%	**3.5**
2021	3.5	9.0	170k	15k	6k	+3.5%	2.4%	1.5
2022	1.5	8.3	169k	14k	2.6k	+1.5%	2.9%	0.5
2023	0.8	7.2	169k	12k	-1.3k	-0.8%	3.4%	-0.2
2020-2023	3.8	8.3	680k	56k	18k	+2.7%	3.1%	1.2

Table 8: Pandemic yearly predictions and resultant excess mortality in The Netherlands, dataset NL ALL. Significant excesses in bold.

Figure 19 illustrates excess mortality in The Netherlands. Notably, at all times, both in fit period 2000-2019 and prediction period 2020-2023, the model succeeds in following the overall trend, yielding fluctations that are substantially zero-sum over reasonably small time scales. The zoomed-in result at right shows two covid waves in 2020. Not only *after* the 1st wave, but also *before*, a deficit occurred. Throughout 2021 a remarkable linear slope occurred.

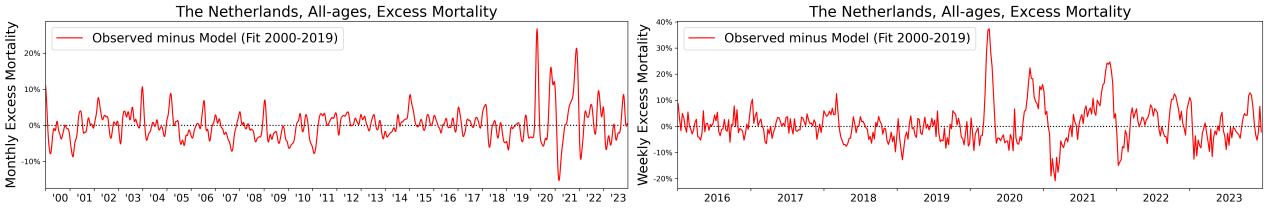

Figure 19: Excess mortality for NL ALL, in % of year mortality. Model fitted over 2000-2019, excess is observed minus fitted/predicted 2000-2023. Monthly and weekly (zoomed-in) time resolution.

Figure 20 shows the pandemic prediction by the proposed model for The Netherlands, ages 0-65, using relative mortality.

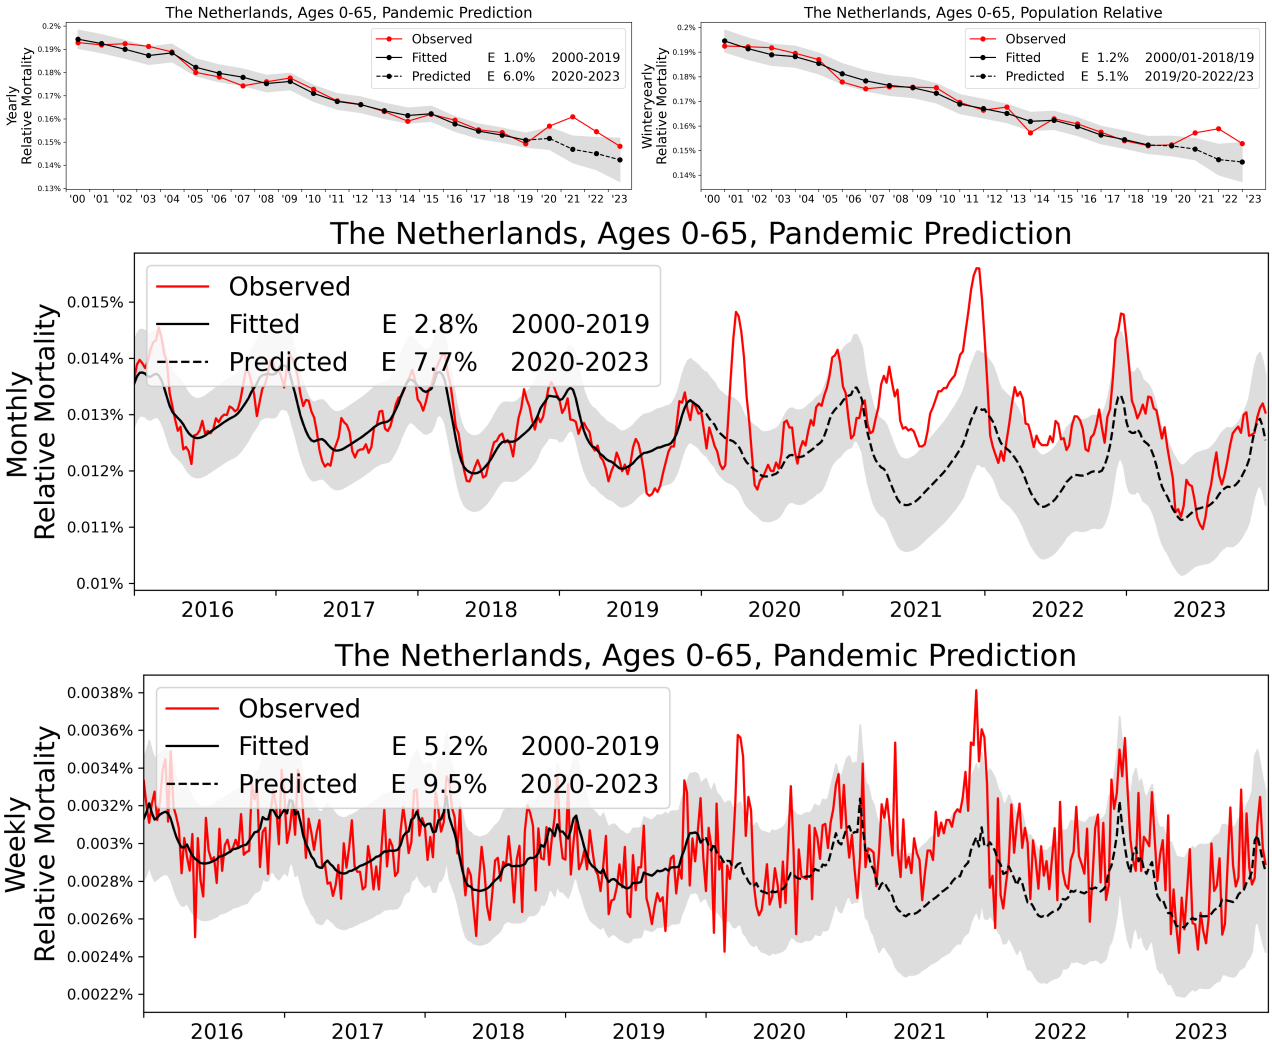

Figure 20: Extrapolation from prepandemic to pandemic years in The Netherlands, ages 0-65. All results based on population-relative mortality.

For extrapolative prediction of 4 pandemic calendar years at year resolution, the error made by the model in predicting the expected, normal part of mortality is $\mu_A E_{fit}$ = ca. 2.6 x 1.0% ≈ 2.6% (see Tables 6-7). The excess has yearly $E_{prediction}$ of 6.0%, and is thus $Z \approx 2.3$ sigma significant. For the 4 winteryears 2019/20-2022/23, $\mu_A E_{fit}$ = ca. 1.9 x 1.2% ≈ 2.3% with $E_{prediction}$ of 5.1%; the excess is $Z \approx 2.2$ sigma significant. Note that the excess in winteryears vs years is similar in significance, but different in magnitude (5.1% vs 6.0%). This is expected, as winteryears include the 2nd half of prepandemic 2019 instead of the 2nd half of pandemic 2023.

Winteryear 2019/20, which includes the 1st covid wave, does *not* show any excess mortality. Weekly and monthly results show two mortality deficits (autumn 2019 and slightly before 1st wave), that precompensated the wave. As the deficit preceeds the excess, there is no loss of lifetime but a saving, and covid marked the end of it. At the end of 2020, the 2nd wave occurs, and a few months into 2021, a strong and continuous excess wave starts peaking strongly at end of 2021, and declining but continuing into 2023.

Table 9 quantifies excesses of Figure 20. Note that the excess in 2023 is still higher by value than in 2020, and similar for excess in winteryears 2022/23 vs 2020/21. The later excesses are only less significant due to the wider bandwidths of later predictions.

Year	2016	2017	2018	2019	2020	2021	2022	2023
$\Delta M_{excess}/\mu_M$	+1%	0.4%	0.7%	-1%	**3.4%**	**8.7%**	**6.1%**	4.0%
$\mu_A E_{fit}$	1.0%	1.0%	1.0%	1.0%	1.5%	1.9%	2.3%	3.2%
Z	1.0	0.4	0.7	-1.0	**2.3**	**4.6**	**2.6**	1.3

Winteryear	2016/17	2017/18	2018/19	2019/20	2020/21	2021/22	2022/23
$\Delta M_{excess}/\mu_M$	+0.7%	-0.3%	-0.2%	0.3%	**4.2%**	**7.9%**	4.9%
$\mu_A E_{fit}$	1.2%	1.2%	1.2%	1.4%	1.5%	2.2%	2.8%
Z	0.6	-0.3	-0.2	0.2	**2.8**	**3.6**	1.7

Table 9: Prepandemic and pandemic excess mortality for The Netherlands, dataset NL 0-65. Significant excesses in bold.

Figure 21 shows results for other age categories in The Netherlands (with absolute mortality).

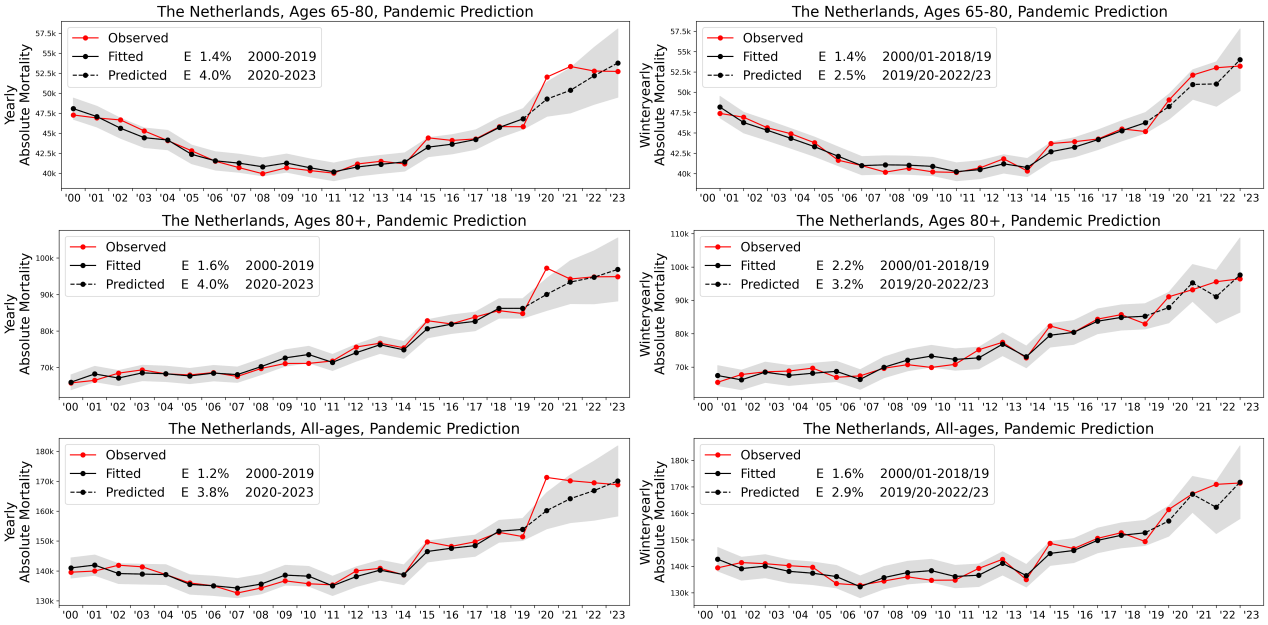

Figure 21: Extrapolation from prepandemic to pandemic years in The Netherlands, ages 65-80, 80+, and all-ages. Left) years, Right) winteryears.

From all results in The Netherlands, several observations can be made:

- Persistent excess mortality is seen *only* for ages 0-65, reaching significance $Z \approx 4.6$ in 2021
- Ages 65+ have just-significant excesses in 2020 and 2021 ($Z \approx 2$), but *none* in *any* winteryear
- No *matching* deficit after any excess has occurred in any age category
- The year 2023 shows signs of a possible deficit starting for ages 65+

FPD predicts that excesses fade in a few years even under a sustained health pressure upon the frail, and that excess-matching deficits occur once the health pressure is released (Figure 7 top-left). Overall, two conclusions can be drawn:

- For ages 65+, the moderate, declining excesses and absence of deficits suggest the pandemic provided a very modest but persistent health pressure, with a possible release starting the earliest by end of 2023

- For ages 0-65, the strong excesses and absence of deficits suggest the pandemic provided a strong and persistent health pressure, that is either growing upon the few frail within the age category, or remains persistently applied (also) upon the healthy, starting only *after* the 1st wave and with no signs of release in 2023

These are strong and disturbing signals, especially as covid mortality is known to focus on the old and frail [Axf,Ghi,Pez], the 0-65 excess rose strongest in 2021 and not in 2020, covid was no longer omnipresent in 2023, and vaccines should have eliminated the health pressure during 2021 or even 2022.

Table 10 and Figure 22 provide winteryear results for the 10-EU datasets. Significant excess mortality during pandemic winteryears occurs in:

- Germany and The Netherlands : *never*
- Belgium, Finland, Spain and Sweden : 1-2 winteryears, $Z \approx$ 2-4
- Bulgaria and Poland : 2 winteryears, $Z \approx$ 10-12
- France and Italy : unreliable excess overestimation, μ_A are not applicable

	Z							
Dataset	2015/16	2016/17	2017/18	2018/19	2019/20	2020/21	2021/22	2022/23
BE	-1.4	0.1	-0.5	-0.3	**2.9**	1.0	1.1	-0.5
BG	-0.4	0.5	-0.8	1.0	-1.2	**10**	**7.0**	-0.6
DE	-1.1	-0.2	0.3	-0.5	-0.7	-0.5	1.3	1.2
ES	-0.5	0.6	0.2	-0.3	**4.3**	**2.0**	0.8	0.6
FI	0.2	1.4	-0.6	-1.2	-0.2	-0.4	**3.2**	**2.3**
FR*	*-1.7*	*0.1*	*0.6*	*-0.3*	***3.7***	***7.3***	***4.4***	***3.5***
IT*	*-0.8*	*0.9*	*-0.7*	*-0.1*	***4.4***	***7.7***	***2.5***	*1.8*
NL	0.3	0.2	0.4	-1.3	1.3	-0.2	1.8	-0.2
PL	-1.6	-0.2	0.6	1.1	0.6	**12**	**5.1**	0.1
SE	-1.4	1.1	0.1	-0.8	**3.5**	0.6	0.6	0.8

*Table 10: Pandemic excess significance in 10 EU countries. *) France and Italy have time-limited mortality data (from 2011, 2013 up), invalidating the use of amplification factors A from Table 7.*

The results for France and Italy overestimate excess significance due to non-applicable error amplification factors μ_A. The μ_A's used for the bandwidths in Table 10 and Figure 22 were based on 16-year model fits $A_{16\to4}$ and used as approximation for $A_{20\to4}$, while France and Italy datasets have much less years for the model fit (8 and 10 years). The latter lowers FR and IT's model fit errors *and* increases their real prediction errors, while estimated prediction errors scale down with model fit error. For France and Italy, proper amplification errors $A_{10\to4}$ (and/or $A_{8\to4}$) should have been computed, which I did not.

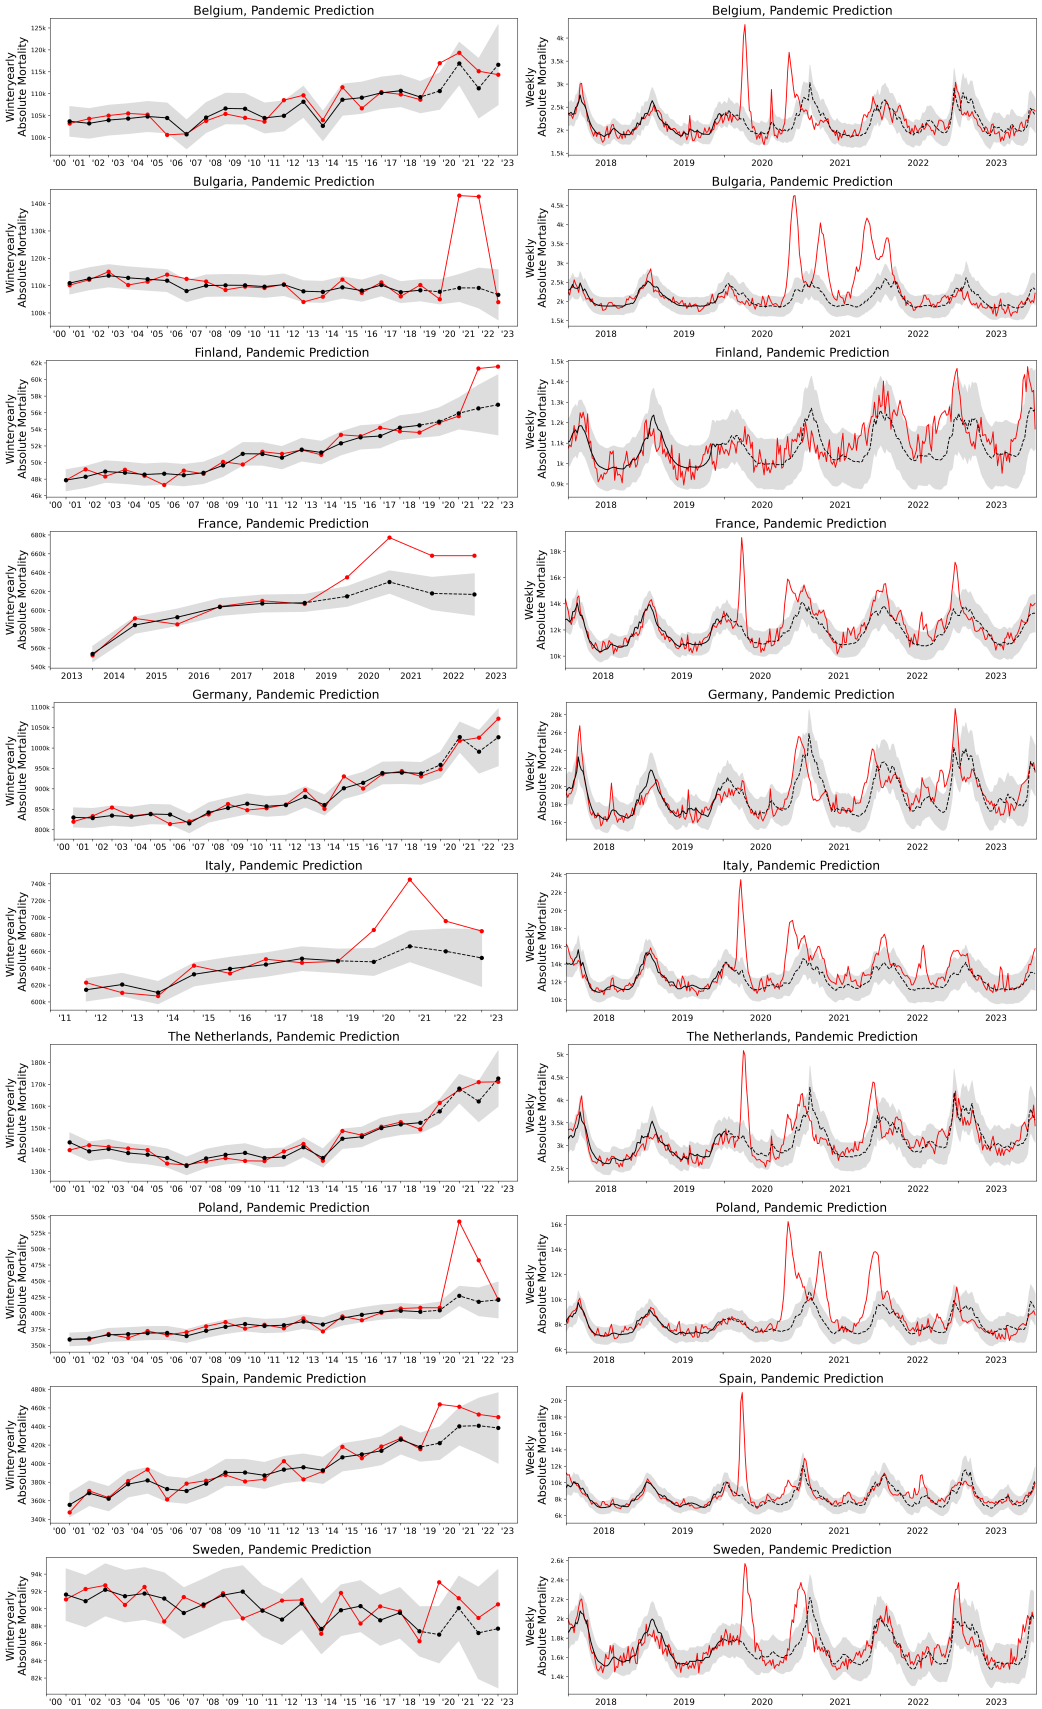

Figure 22: Prediction of pandemic mortality by extrapolation from prepandemic mortality, 10-EU datasets, winteryearly and weekly. France and Italy have time-limited mortality data (from 2011, 2013 up), invalidating the use of amplification factors A from Table 7.

The proposed model does not confirm the findings in [Kuh] of high excess deaths in Germany during pandemic years (see Figure 1). Instead, it measures *no* significant excess mortality in

Germany during the entire pandemic. This is not due to a high significance threshold; ΔM_{excess} is in fact *negative* in the first two pandemic winteryears, and starts being positive only in 2021/22. Mortality in winter 2020/21 was truly high, and covid may have slightly accelerated deaths. Nonetheless, the proposed model predicted these deaths due to winter temperatures, if only a few weeks later.

Figure 23 further shows that, according to the proposed model, all net excess in Germany concentrated in calendar year 2022.

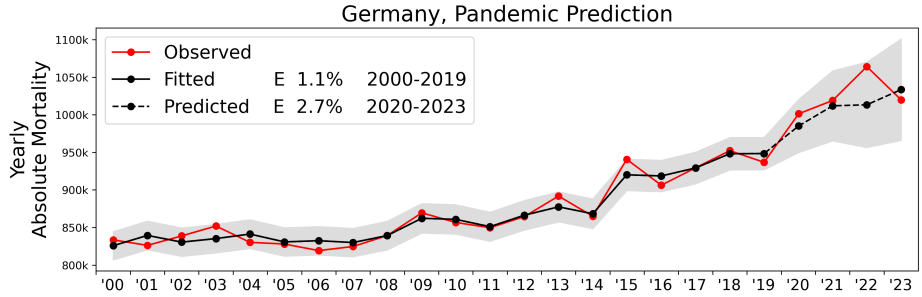

Figure 23: Prediction of pandemic mortality by extrapolation from prepandemic mortality, Germany. Compare with Figure 1.

3.6 Pandemic model: determinants for covid and vaccination campaigns

From here on, the pandemic mortality model is used, including covid test and vaccine dose determinants, to explain pandemic excesses and deficits. The model is fitted to the entire period 2000-2023. The model error is evaluated over the entire fit period as well as the pandemic subperiod 2020-2023. Data-limited FR and SE are not used in some experiments, or used while excluding the 1st half of 2020 (denoted by 2020.5-2023 in figure legends).

Figure 24 illustrates the substantial improvement by the pandemic model for The Netherlands, all-ages. The fact that the fit error over the entire period (4.6%) equals prepandemic fit error, and fit error for the pandemic subperiod (4.8%) is slightly higher, suggests that the pandemic model is both applicable and not overmodeling/overfitting.

Table 11 shows pandemic subperiod fit errors averaged over 12 datasets (all minus FR and SE), for normal model and partial/full pandemic models. Relative to the normal mortality model, the addition of covid mortality (C and CFR) improves the model most, with an average relative error reduction of 47%. From this CFR-only model, a further reduction is achieved of ca. 10% by including the vaccine's fatal effects (V and VFR), and ca. 10% by the vaccine's protective effects (VE). The full pandemic model $CFR + VE + VFR$ provides a 21% error reduction over the CFR -only model, very close to the sum of individual contributions of VFR and VE, suggesting the latter two contribute independently to model accuracy. With respect to the normal mortality model, a 58% error reduction is achieved. The full pandemic model has an average fit error of 5.4% on the pandemic subperiod, slightly above the average fit error of the normal model on prepandemic periods of 5.2% (on the same 12 datasets).

Clearly, the main determinant of pandemic excess mortality is covid. Further, the vaccine's fatal effects appear to match its protective effects. This resonates with findings in [Ben], that mRna vaccines (as used mostly in The Netherlands) have a net zero effect on all-cause-mortality.

Figures 25-26 show results for all datasets. As expected, the yearly fit accuracy is better during pandemic years than prepandemic, due to the assymetric distribution of model parameters over prepandemic and pandemic years. The model fits most datasets well, except FR and SE where the first covid wave is not well represented by test data C.

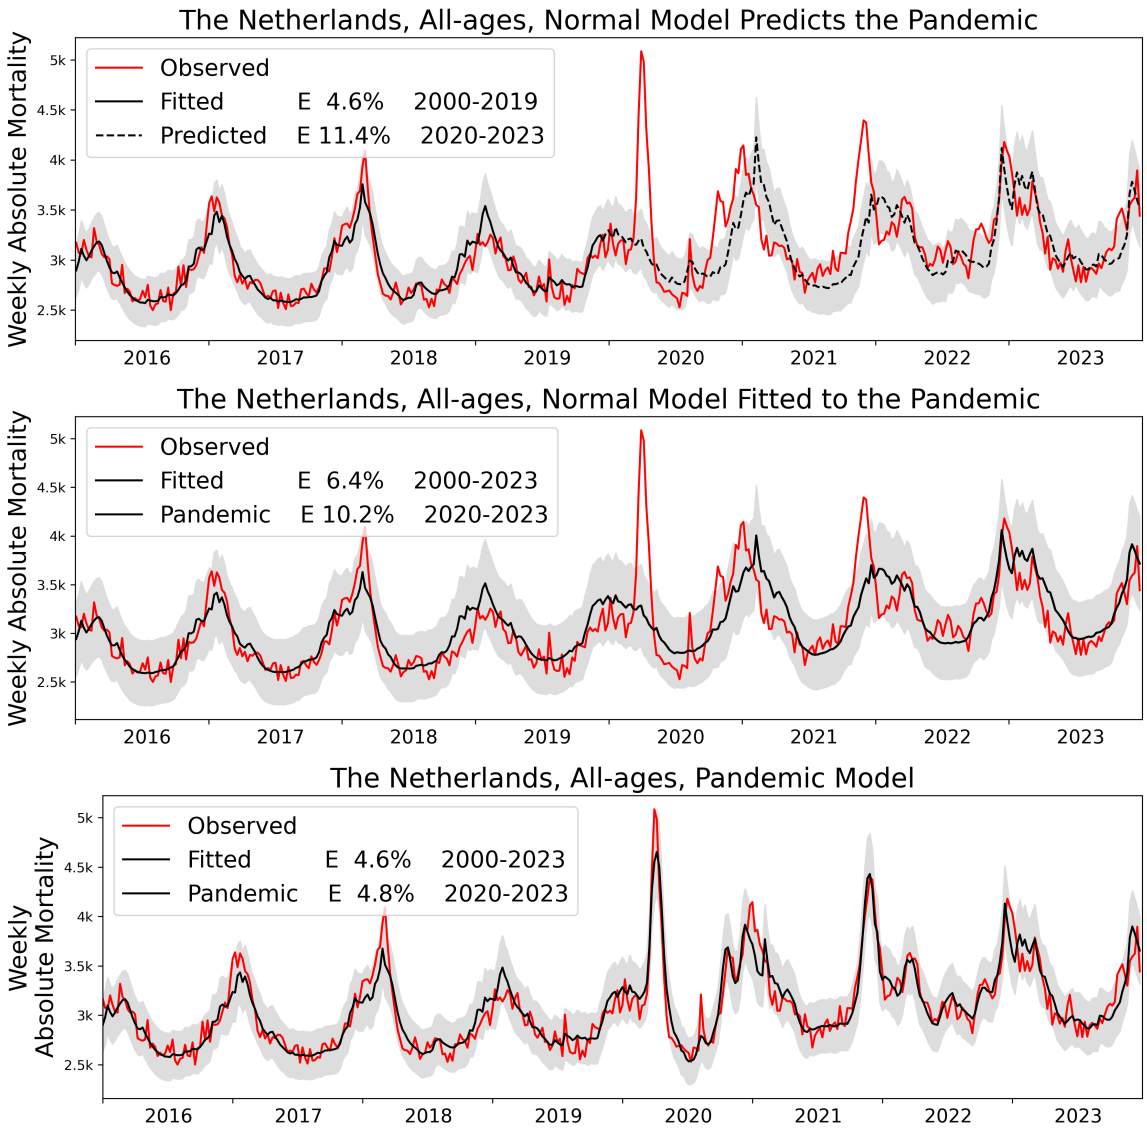

Figure 24: Top) Normal model with prepandemic fit and pandemic prediction, Middle) Normal model fitted to entire period incl pandemic, Bottom) Pandemic model fitted to entire period. Dataset is NL ALL. Fit error within the pandemic subperiod is more than halved from 10.2% to 4.8% (over fourfold variance reduction) and close to the prepandemic fit error (4.6%).

	Pandemic subperiod error $E_{2020-2023}$ [%] for Various Models fitted to entire period 2000-2023						**Prepandemic 2000-2019 E_{fit} [%]**
	Normal model	CFR only	VFR only	CFR and VFR	CFR and VE	Pandemic model CFR+VE+VFR	Normal model
Average 12 datasets	12.7	6.8	11.8	6.1	6.1	5.4	5.2
Improvement by relative E reduction		vs normal **-47%**	vs normal **-7%**	vs CFR-only **-10%**	vs CFR-only **-10%**	vs CFR-only (vs normal) **-21% (-58%)**	

Table 11: Pandemic models from none to full. The model is fitted on 12 datasets (all except FR and SE) over all years 2000-2023, and the average partial fit error is shown for pandemic years 2020-2023. At right for reference, the prepandemic fit error (2000-2019) by the normal model.

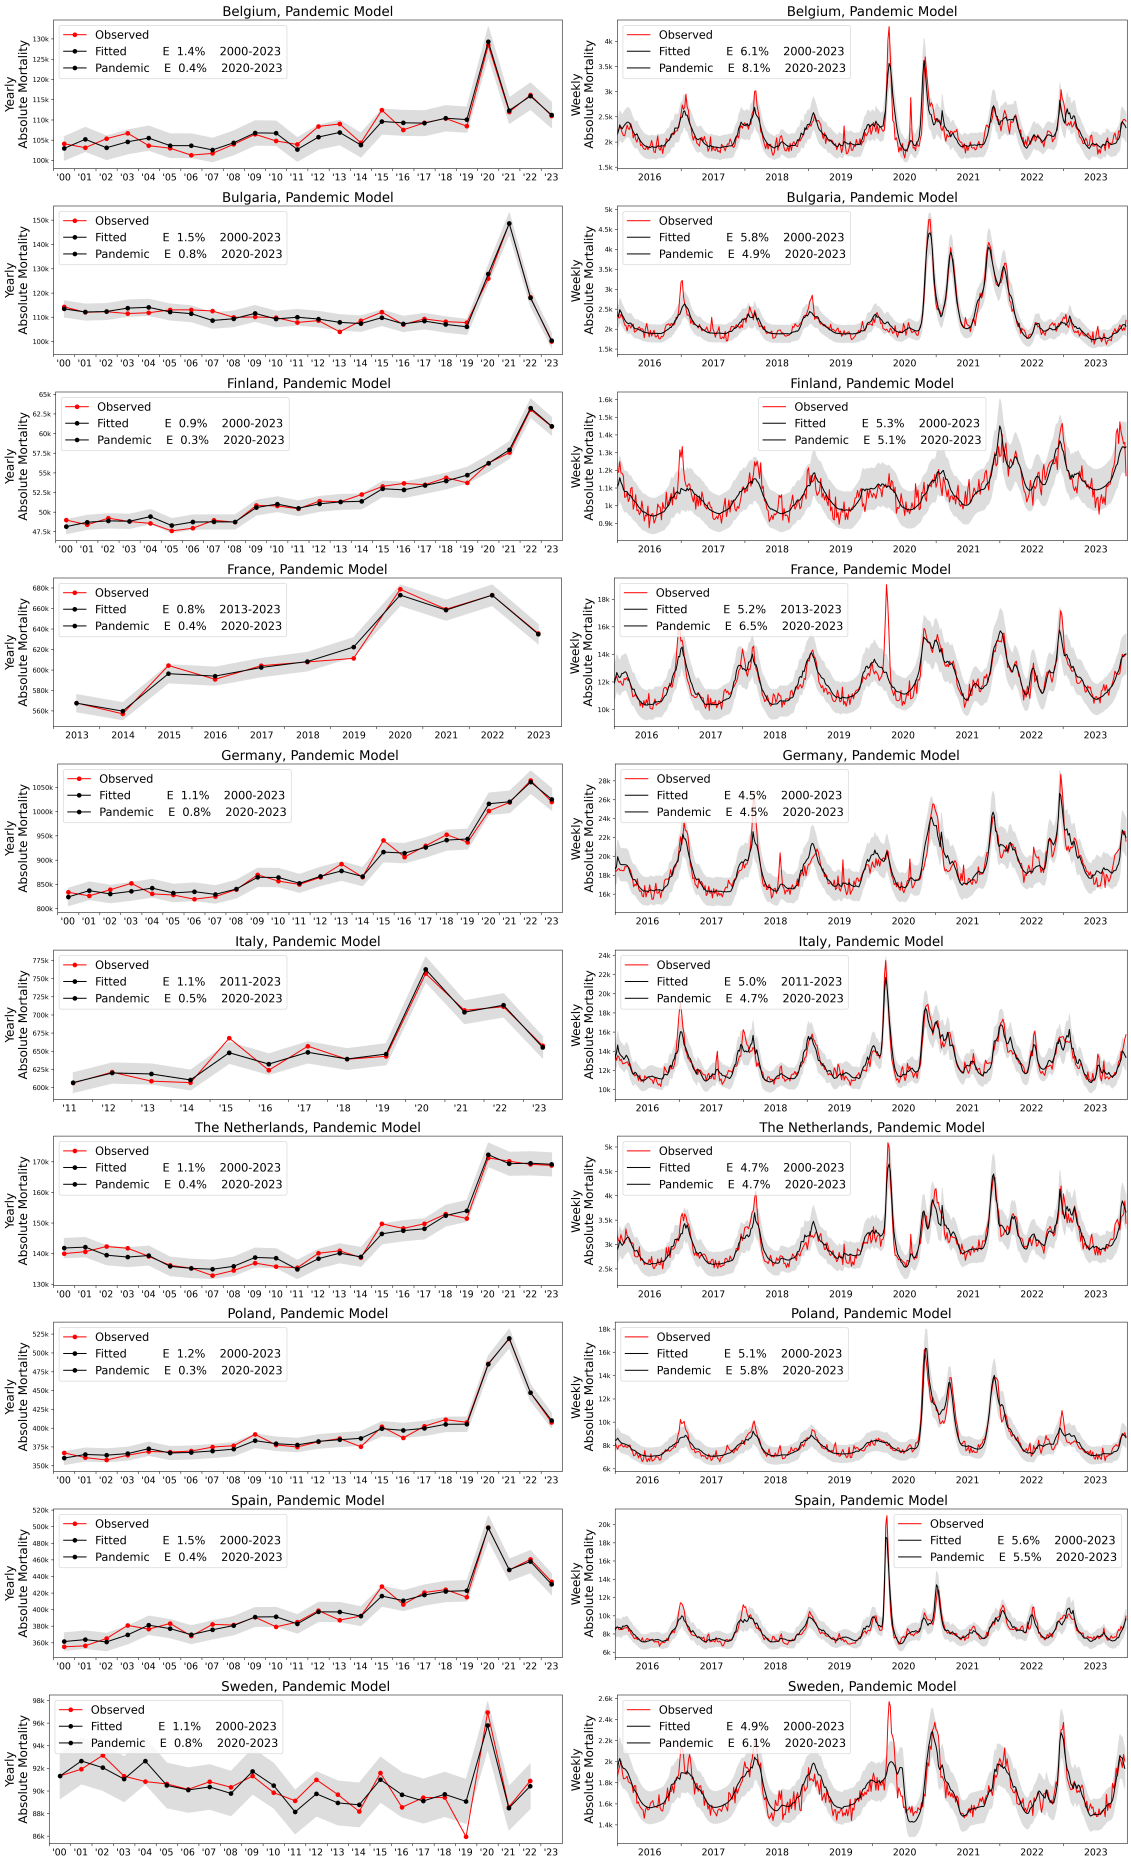

Figure 25: Pandemic all-age results for 10-EU datasets. The model fails to explain the first wave in France and Sweden (FR and SE), as their covid test data does not represent the wave well.

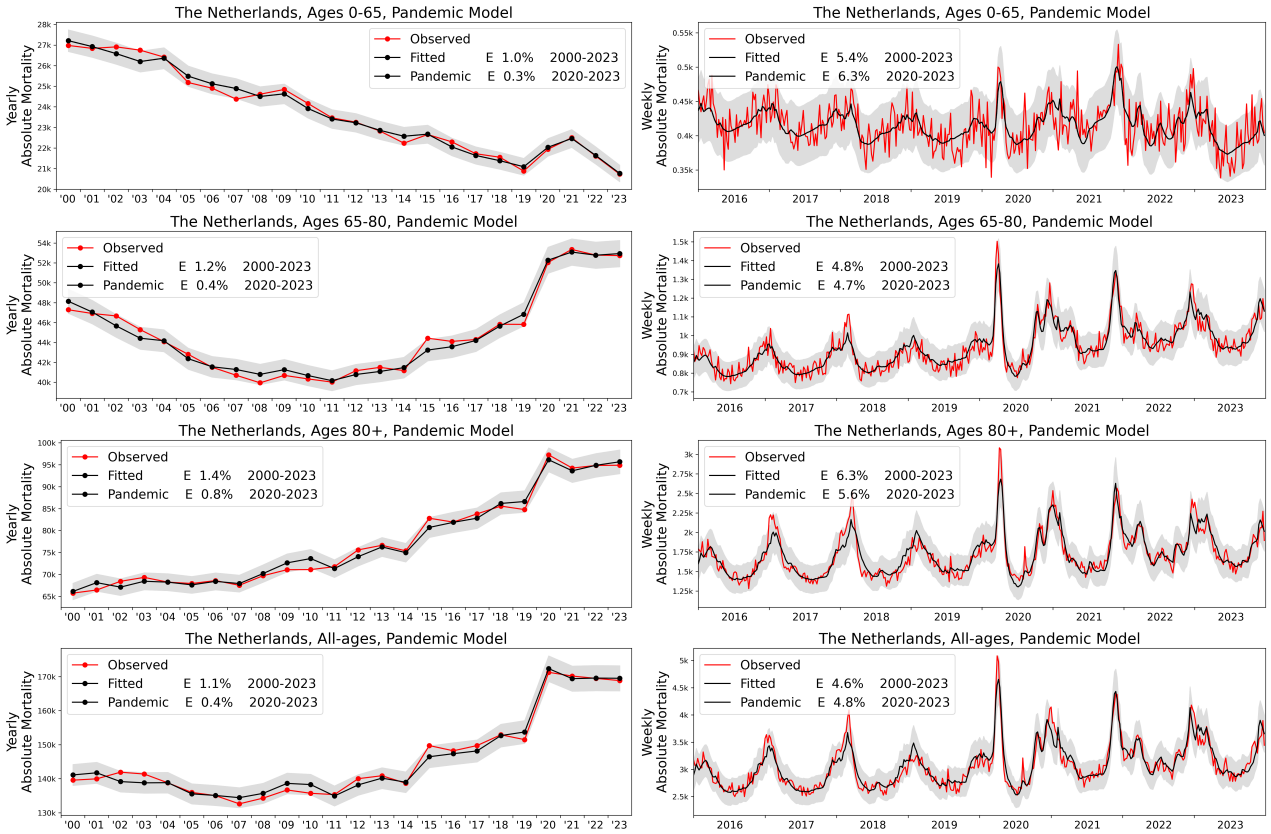

Figure 26: Age-stratified results for the pandemic in 4-NL datasets.

Figure 27 shows FR and SE results where the 1st half of 2020 is excluded from the model fit. Native weekly fit errors are now at the same level as prepandemic fit errors (FR 3.7%, SE 4.5%, see Table 6).

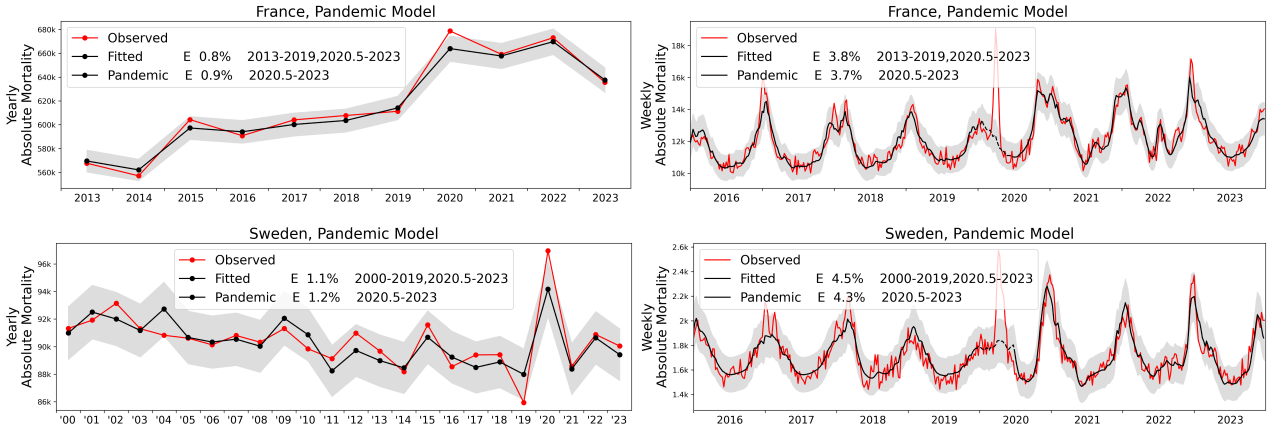

Figure 27: Pandemic results for datasets FR and SE, excluding the 1st half of 2020 in the model fit.

Table 12 shows results for pandemic nonlinear parameters. The μ_{VFR} is an average over all campaigns weighted by number of doses. The CFRs (or μ_{CFR}) are not shown as they are irrelevant on their own, as are test cases C. Only combined as $CFR \cdot C$, they gain relevance, and together with $\tau_{FPD;C}$ they relate to lifetime.

Not unexpectedly given the low quality of pandemic determinants, accuracy of several parameters is very low and far from ideal. This holds especially for FPD time scales, which are used to determine lifetime lost/saved. Also, several parameters vary considerably across

datasets. This may reflect some real differences among populations, but also insufficiencies in both source data and my brute-force optimization algorithm.

	Mortality Natural Variability	**Covid Mortality**		**Vaccine Effectivity** against covid mortality			**Vaccine-dose Mortality** fatal adverse events		
Data set	α [%]	$\tau_{HPA;C}$ [w]	$\tau_{FPD;C}$ [w]	VE_0 [%]	$\tau_{VE-100\%}$ [w]	τ_{VE} [w]	$\tau_{HPA;V}$ [w]	$\tau_{FPD;V}$ [w]	μ_{VFR} [%]
NL 0-65	2.8±0.3 (2.3-3.3)	1.4±0.3 (1.1-3.6)	>48	100	>22	>15	>40	>120	.05±.02 (.02-.09)
NL 65-80	3.4±0.1 (3.2-3.6)	1.8±0.2 (1.4-2.2)	20±4 (13-30)	>58	0	>23	36±15 (12-75)	>44	.13±.04 (.06-.22)
NL 80+	5.6±0.1 (5.4-5.9)	2.5±0.3 (2.0-3.3)	10±2 (7.1-14)	100	11±0.7 (10-13)	9.2±2.1 (5.3-13)	4.9±2.6 (1.8-13)	>16	.32±.11 (.20-62)
NL ALL	4.2±0.1 (4.0-4.4)	1.8±0.2 (1.5-2.2)	23±4 (16-33)	100	16±0.6 (15-17)	21±8 (13-37)	15±6 (8-33)	>100	.09±.02 (.06-.15)
BE	5.4±0.1 (5.1-5.7)	2.9±0.2 (2.5-3.3)	42±6 (32-58)	100	9.8±1.5 (7-13)	5.3±1.4 (3.2-9.0)	12±6 (6-30)	>81	.07±.02 (.04-.10)
BG	5.4±0.1 (5.2-5.6)	3.1±0.3 (2.6-3.7)	47±15 (26-88)	100	>3.7	>2.6	39±24 (25-82)	>240	2.8±0.4 (2.3-4.0)
DE	4.3±0.1 (4.1-4.5)	5.1±0.6 (3.9-6.1)	6.7±0.8 (5.5-8.9)	100	15±0.8 (14-17)	>92	>63	>120	.22±.07 (.12-.41)
ES	5.1±0.1 (4.9-5.4)	1.7±0.1 (1.5-1.9)	74±18 (49-120)	100	19±3 (13-22)	>15	>9	>190	.13±.06 (.08-.33)
FI	4.3±0.1 (4.0-4.6)	9.0±1.4 (6.6-12)	12±2.3 (9-17)	>38	0	45±7 (33-55)	>350	>480	.38±.09 (.19-.55)
FR*	3.6±0.1 (3.4-3.9)	6.9±1.8 (4.0-10)	23±11 (11-52)	>50	0	55±11 (38-81)	5.3±2.1 (2.2-10)	6-760	.03±.01 (.01-.06)
IT*	4.8±0.1 (4.6-5.1)	1.3±0.1 (1.1-1.5)	>350	>45	<8	4.3±2.2 (1-9)	2.7±1.3 (1-6)	>8	.07±.01 (.06-.09)
NL	4.2±0.1 (4.1-4.4)	1.9±0.2 (1.6-2.3)	22±3 (16-30)	100	16±0.6 (15-17)	14±3 (10-23)	15±6 (8-36)	>73	.10±.03 (.07-.18)
PL	4.7±0.1 (4.6-5.0)	1.9±0.1 (1.7-2.2)	>360	>3	(0-1)	(1-13)	12±5 (6-23)	>75	.09±.01 (.06-.11)
SE*	3.7±0.1 (3.9-4.3)	5.8±0.6 (4.8-7.0)	7.7±0.9 (6.4-9.8)	100	26±2 (21-29)	37±5 (29-45)	>2	>130	.03±.01 (.01-.04)
All sets	4.4±0.8	3.4±2.4	26±20	100	13±8	24±19	16±13	-	0.3±0.7

Without outlier BG: 0.13±.11

*Table 12: Pandemic nonlinear parameters by $\mu \pm \sigma$ and/or 95% (CI), >CLB or <CUB. Statistics over all datasets are computed over μ only. *) FR and SE have unreliable covid test data in the 1st half of 2020, this half year was excluded from their analysis. FR and IT have time-limited prepandemic data.*

Despite low accuracies, several inter-dataset average results match well with other existing findings:

- Covid mortality occurs quickly after a positive test, with $\tau_{HPA;C} \approx$ 3w
- Lifetime lost per person dying from covid is $\tau_{FPD;C} \approx$ 26w, ca. half of temperature's τ_{FPD}
- Vaccine effectivity starts at ca. 100%, remains so for ca. 13w, wanes in ca. 24w
- Fatalities can occur many months after vaccination, by $\tau_{HPA;V} \approx$ 16w, and a year for 0-65
- Lifetime lost per person dying from the vaccine, $\tau_{FPD;V}$, ranges from years to decades
- Vaccine-dose fatality rates μ_{VFR} are ca. 0.13% (all-age, ignoring outlier BG), and rise with age

Time to death after a positive covid test is found as ca. 3 weeks without any clear age dependance given the 4-NL results. This is in line with the 18 days regardless of age reported in [Mar2].

The fact that $\tau_{FPD;C}$ is ca. *half* that of temperature's τ_{FPD}, means that covid targets people *more frail* than those targeted by natural causes represented by temperature. These people are not just nearing their natural end of life*;* they are the lucky few that managed to live beyond their natural end by surviving the previous seasonal cycle. These findings resonate with reported ages at death in The Netherlands [Cbs5], which was ca. 82 year for deaths by covid (averaged over males/females and suspected/sure covid cases) and ca. 79 year for all-cause-mortality (estimated from 10-year cohorts).

The high VE_0 and fast waning are in line with many studies in vaccine effectivity against covid mortality, e.g. [Riv3].

The $\tau_{FPD;V}$ is considerably higher than temperature's τ_{FPD}, up to many years. Vaccine fatal adverse events thus target much healthier people than natural causes do. The mean time $\tau_{HPA;V}$ between dose and fatality (if any), is ca. 16 weeks (all-ages, all datasets), but rises for younger people to ca. >40 weeks (ages 0-65).

The μ_{VFR}'s found here of ca. 0.3% over all datasets, or 0.13% without outlier BG, are comparable to those found in my earlier study [Red3] which are 0.13%±0.04% (0.05-0.21%) in NL and 0.35%±0.10% (0.15-0.55%) in EU (both refer only to campaigns in 2022). This study used a *very* different approach based on explicit causal modeling of weekly differential mortality and vaccine uptake, also taking into account positive tests as a confounder.

For all-ages, the found μ_{VFR}'s are of the same order as covid IFR [Axf,Ghi,Pez]. With VE_0 at near 100%, the vaccine's net effect on all-cause-mortality would be low, in line with findings in [Ben]. The net effect of covid vaccinations appears to be a change in cause-of-death, from covid to whatever the fatal adverse event manifests as (heart attack, stroke, "natural death", etc).

Table 13 shows derived lifetime lost/saved by covid and vaccines. Similar to pandemic parameters, accuracy is low. For every dataset, covid costs and vaccination saves considerable lifetime. Also, *for every dataset* a substantial amount of lifetime is lost by vaccinations. By order of magnitude, the losses match the savings. If vaccines are safe and fatal adverse events are rare, one expects LTL_V to be zero, or at least very low.

As explained in the pandemic model section, my model enforces non-negativity on CFR and VFR parameters by a zero-lower-bound, thereby biasing the average of zero-mean random fluctuations into a positive net value. I have repeated the entire analysis without zero-lower-bound, allowing any real value for CFR and VFR, including negative. This result is shown in Table 13 in grey. While non-physical (both some CFR and VFR actually become negative), the result is free from this positive bias; and it is essentially still the same.

Dataset	**Covid Mortality** Lifetime lost to covid LTL_C [ky]	**Vaccine Effectivity** against covid mortality Lifetime saved by vaccine LTS_V [ky]	**Vaccine-dose Mortality** fatal adverse events Lifetime lost to vaccine LTL_V [ky]	no zero lower bound
NL 0-65	1-52	>18	13-94	52-120
NL 65-80	4-7	6-290	16-65	22-100
NL 80+	5-9	10-43	5-34	2-20
NL ALL	14-28	33-120	50-500	100-620
BE	24-40	13-34	20-290	27-390
BG	48-140	2-100	420-1400	1-610
DE	36-50	2800-5000	450-4600	570-2300
ES	150-360	370-9700	470-1900	840-4100
FI	3-8	14-42	670-2800	220-1400
FR*	68-140*	150-470	7-600	> 8400
IT*	1100-3800	1-1400	180-1500	150-1400
NL	15-25	30-78	50-400	41-320
PL	1000-3100	1-150	67-1100	97-1800
SE*	5.2-7.6*	29-100	3-120	210-750

*Table 13: Lifetime lost/saved in kiloyears, by 95% CI or CLB. *) FR and SE have unreliable covid test data in the 1st half of 2020. This half year was excluded from their analysis, so their covid mortality does not include the 1st wave. FR and IT have time-limited prepandemic data.*

4 Discussion

Study limitations: As a new and different approach to mortality modeling, my examination could not possibly cover all aspects thoroughly; preliminary results were used on many occasions.

My model uses only population-aggregated data and no individual data, due to data availability to independent researchers. A limitation on results but a strength in terms of data requirements.

For all-age populations, absolute mortality performs as good as population-relative mortality, thereby circumventing the need for population data. Still, applying the model on population-relative data will per definition provide better results when used with highly-stratified datasets (e.g. 5 year age bins at weekly time resolution), but such datasets are not (publicly) available.

I used data of only ten countries; starting with 4 datasets I added more during the work to avoid data exhaustion, but more datasets are needed. Given the conclusions for the young and healthy, more experiments with age-stratified datasets are certainly justified.

Several experiments with determinants alternative or additional to temperature did not make it into this report, as results where too inconclusive. These determinants include other weather features (absolute humidity, indoor and outdoor relative humidity, precipitation, sunlight), an optimism/pessimism indicator (AEX stock index [Fdb]), and a datafree determinant (fixed sine wave) as also used in the recently renewed Dutch baseline [Riv1]. Here, I would like to mention a few preliminary results that may incentivise follow-up research. The AEX determinant performed nearly as bad as a random-number determinant. Several of the other weather features performed similar to temperature or even better, but this was only tested on NL datasets as my EU datasets lacked determinant data. The datafree sine determinant outperformed temperature on 4 NL

datasets, but underperformed on 8 EU datasets. Further, systematic delay between several determinants and mortality was measured, i.e. “system dead-time” where health pressure does not yet lead to any effect on mortality at all. For temperature, half a week of delay was measured, but this was not taken into account in the proposed model. An advantage specifically for the sine determinant is that the delay can be represented by additional linear parameters (via cosine terms), where all other determinants need an explicit non-linear parameter. Unexpectedly, incorporation of the delay in the model *increased* prediction errors for several determinants.

The higher-than-expected variability of mortality [Red4] is still not explained by my model, using temperature or any of the other determinants mentioned above. The mystery of nationwise correlation among deaths remains interesting as it may relate to an important end-of-life process we are currently unaware of. If this process is identified (assuming it can be), it will probably not only explain mortality’s variability but also provide more accurate predictions. Given that the variability has been present for decades, new potential causes can be ruled out (e.g. 5G radiation).

There are quite some options to explore further, namely everything shared by many people over many decades: air, food, water, cosmic radiation, yet other weather aspects as UV light exposure, weekly optimism/pessimism inflicted on people via mass-media, etc. Maybe brightness and/or emotion levels of TV-shows watched by frail elderly have strong effects in the final days of life. Maybe health care workers’ behaviour is synchronized by traffic intensities and daily prices of electricity or medicines. Maybe some randomness is implicitly injected into death registrations, e.g. for privacy purposes. And quite possibly, temperature is the main driver after all and my model is just not capturing it well.

My linear FPD model is a simplification of the canonical mortality model, which is a simplification of reality. A first upgrade to FPD would be to use the canonical model, where the frail pool is not uniform but has several subpools binned by frail age, that is, stratified by health. This supports other health pressure interactions that are more deterministic, e.g. “this event kills every frail person with health level 16% or lower”. This leads to a variable FPD time depending on health pressure strength and duration, and presumably better results. Model optimization, however, will require more brute force computing and exclude the use of linear regression.

There are many tiny details that affect results. In data preparation, I did not use population density to weigh temperature readings from different weather stations. I used only quadratic health pressures by temperature, but other functions may be more suitable as a 3rd degree polynomial clearly showed promise. In the pandemic model I averaged the first two vaccine doses into a single “full” dose, but one could also use the original two doses plus a more realistic protection saturation model. The hidden universe [Bre] offers infinitely many such alternative model options, and I tried almost none of them.

VE biases and FPD: In at least three areas of VE measurements, the concepts of FPD and frail age may turn out useful. First, a generic weakness in common VE analyses is that lives-saved are reported rather than lifetime-saved. The former has meaning only for as far as it reflects the latter, and in FPD they are related by a single, measurable constant.

Secondly, FPD may assist in the quantification of the so-called healthy vaccinee effect (HVE) [Høe]. This effect is caused by the sensible policy to exclude people from vaccination when they are deemed to frail to survive the vaccine’s health pressure. This biases natural deaths systematically towards unvaccinated status, which on its turn biases apparent VE against all-cause-mortality (VE-ACM) optimistically if the policy is not compensated for. In a simplified HVE model, the policy effectively excludes people within *N* weeks from natural death. The $\tau_{FPD;V}$ defined in this report is a related property of vaccination for which *N* acts as a lower bound, which for ages 65+ in The Netherlands is roughly >35 weeks. For ease of argument, if one simplistically assumes *N* = 35 and all eligible 65+ people get vaccinated, then all natural deaths in the next 8

months are unvaccinated. This results in a very high apparent VE regardless of any true vaccine effect for 8 months.

Thirdly, another bias that relates to HVE and thus to FPD, is what I adopt as vaccine-death mislabeling (VDM). It builds on the common malpractice of intentionally mislabeling people as being unvaccinated for 2-4 weeks after being vaccinated [Cbs6]. Table 14 shows a simplified representation of the mislabeling, which is typically assisted by different and complex vaccination status definitions per report, per institute, per country, per vaccine brand, etc. The table also shows a proposal for nomenclature disambiguation, where the bar means “non” and the dot means “recently”.

People that	Currently labelled as	Formerly known as	Proposed new label
never received a jab	Unvaccinated $\bar{V}$		Non-jabbed $\bar{J}$
received a jab	Ambiguous $\bar{V}$ & V	Vaccinated V	Jabbed J
received a jab < 2-4w ago	Unvaccinated $\bar{V}$	Vaccinated V	Recently-jabbed $\dot{J}$
received a jab > 2-4w ago	Vaccinated V		Nonrecently-jabbed $\bar{\dot{J}}$

Table 14: Current medical practice of intentionally mislabeling people that received a vaccine dose.

Undisputedly, the immune system needs some weeks to respond to the vaccine’s health pressure. While this justifies the *intention* to take that time into account in VE assessment, it does not justify the means chosen of *mislabeling*. If a vaccine-death occurs within the mislabeling period, VDM labels that death as unvaccinated: the more deaths, the higher VE-ACM.

VDM relates to HVE and FPD via the small but everpresent group of very frail people balancing on the edge of the exclusion policy. Inevitably, a certain part will get vaccinated and die quickly as a result, labeled as recently-jabbed and thus unvaccinated. In a campaign, the frailest are vaccinated first, who are most susceptible to VDM; this biases VE-ACM mostly in the first month.

The bias investigated in [Nei2] relates also to the recently-jabbed, but not to FPD. It is based on *natural* death mislabeling (NDM): recently-jabbed people $\dot{J}$ are assigned vaccinated status V when they live but unvaccinated $\bar{V}$ when they die, from whatever cause. NDM differs from VDM in that it does not depend on vaccine action; it turns even a placebo into effectiveness. NDM biases VE-ACM for as long as the campaign runs, creating the perceived need for an endless stream of booster campaigns. NDM operates similar in all age categories.

While beyond the scope of this report, the new nomenclature may help describe intricacies in VE biases in a simple and explicit way:

$$VE_{TRUE} = 1 - \frac{N_{\dagger J} \,/\, N_J}{N_{\dagger \bar{J}} \,/\, N_{\bar{J}}} \qquad VE_{BIASED} = 1 - \frac{N_{\dagger V} \,/\, N_V}{N_{\dagger \bar{V}} \,/\, N_{\bar{V}}} \qquad VE_{NDM} = 1 - \frac{N_{\dagger J - \dot{J}} \,/\, N_J}{N_{\dagger \bar{J} + \dot{J}} \,/\, N_{\bar{J}}} \tag{27}$$

$N_{\dagger J}$ Total number of deaths among the jabbed during some period P

N_J Average number of jabbed people during P

(and similar for other jab/vaccine statuses, P small so average N_J has few variance)

The above shows the absurd assymetry of NDM, where 2 of the N’s are biased. The equation for VDM is more complex involving 3 biased N’s and not shown here.

Pandemic analysis: The results of my pandemic model have substantial inter-dataset variability and wide confidence intervals of parameter values. Determinants are of low data quality

for various reasons, which seems not easy to improve on. NPI stringency indices were not included.

My proprietary brute-force optimization and sampling algorithms are far from perfect, and may have found local sub-optima, or failed to measure parameter inaccuracies well. Similar issues may play a role as well in other studies that use default techniques. Default is not the same as free-of-issues; it often just means free of *discussion*. Cox Proportional Hazard (CPH) models are widely used for VE analysis. The model's inventor (sir Cox) was openly critical about its use in biology and preferred Accelerated Failure Time models [Wik]. From [Wik]: "There is strong basic science that … AFT models are the correct model for biological survival processes". I am unfamiliar with AFT models but my model's frailty and HPA ingredients seem related.

My pandemic results need to be taken with care and require more research. Like all science, medical in particular [Bre,Ioa]. In this context, Dutch national institutes used CPH models to claim the primary vaccination campaign had VE-ACM up to 50-75% in all age categories [Cbs1], see Figure 28. Notice the bump in the first month, and the sharp decline after 8 months at ages 65+; these findings probably suffer from one or more VE biases.

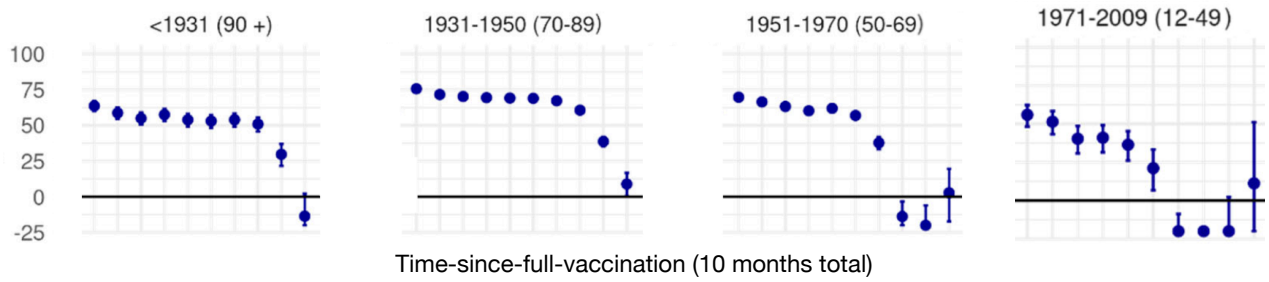

Figure 28: Vaccine effectiveness in % against all-cause-mortality in The Netherlands for the primary vaccination in 2021, source [Cbs1].

If above VE's bare any relation with reality, one would expect that my results would show at least something remotely similar. Each of my four studies ([Red1,Red2,Red3] and the current, covering very different model approaches) finds the opposite: more vaccinations mean more mortality. The only similarity with Figure 28 is located at the end.

The finding of positive lifetimes lost to vaccination LTL_V, is so disturbing that I have examined possible biases in my model. I identified one and have temporally removed it, resulting in a model that is non-physical, without bias, and essentially the same outcome. So many differences among countries, age groups, long-term trends, covid waves, campaign periods, random fluctuations; still, all LTL_V are positive and of similar order as LTS_V. I find it hard to believe that these results are so consistent over all datasets just by happenstance.

5 Conclusions

Objective: Present a model for mortality prediction dedicated to more objective quantification and localization in time of excess mortality. Secondary objective is to analyse the role of covid and vaccines in recent pandemic excesses.

Method: The new model is designed following rules of simplicity, explicitness and source data reliability. The entire model is just one equation, *M = W*atT*, with 12 parameters representing trends over decades, seasonal/weekly variability, health pressure accumulation (HPA) relating to delay between cause and death, frail pool dynamics (FPD) relating to the delay between excess and subsequent deficit which on its turn determines lifetime lost/saved (LTL/LTS), and the natural variability (NV) of remaining random mortality fluctuations.

Normal and excess mortality are explicitly defined in a generic, non-mechanistic way. Input source data is weekly mortality and outside temperature, both of which are objective, reliable and available worldwide. Only aggregated absolute mortality numbers are required, but population-

relative numbers can be used as well. FPD redistributes mortality uncertainties towards births, away from deaths in an age-independent manner using the concept of “Frail Age”. LTL/LTS is more relevant than lives lost/saved which is zero-sum per definition on the long term. NV is modeled to match reported observations and differs from the often-used Poisson model.

Bayesian probabilistics and brute-force numerical optimization are used to fit the model. A simpler 9-parameter model is provided that requires only default linear regression. Prediction accuracy is explicitly defined and parametrized.

Additional determinants can be integrated into the model, to measure their relation to observed excesses. Pandemic excesses are examined via determinants of positive covid tests and vaccination doses administered. Measured parameters are covid Case Fatality Rate (CFR), Vaccine-dose Fatality Rate (VFR), Vaccine Effecivity (VE) against covid mortality, and LTL/LTS by covid and vaccination.

Experiments: Results with 10 EU countries (344M people) yield a prepandemic model fit over 2000-2019 with accuracies relative to average mortality of ca. 1%, 4% and 5% at resp. yearly, monthly and weekly resolution. HPA and FPD times vary over countries and age. For the Netherlands (NL), HPA time (heat/cold stress to death) is found as 4.5 weeks (w) with 95% confidence interval (4.0-5.0w), and FPD time (between excess and deficit, LTL per death) as 44w (33-65w), matching theoretical expectations based on a recent finding that frail lifetime is ca. 4±1 years. Prediction accuracy is measured ranging from ca. 1 to 3x the model fit accuracy, depending on prediction scheme (interpolation/extrapolaton, weekly/yearly resolution, distance between fit and prediction).

Pandemic predictions for 2020-2023 based on prepandemic mortality have error variance ca. 5 times lower than the Dutch national baseline (79% reduction). Highest excess is found in Bulgaria (10σ) and Poland (12σ), while Germany (DE) and The Netherlands (NL) show no significant excess in any pandemic winteryear. For ages 0-65 in NL, substantial excess is found starting from winteryear 2020/21, peaking in 2021 (5σ) and enduring still. FPD predicts mortality deficits in the frail to follow and match excesses within a year, but none are observed in any country or age group. This suggests the pandemic health pressure has been fully persisting throughout 2020-2023, and is significantly affecting the young and healthy.

Pandemic determinants are far less reliable/accurate than temperature, limiting pandemic measurements’ accuracy. Covid HPA time is found at ca. 3±2w over all countries, and 1.8w (1.5-2.2w) for NL, matching known mean times between positive test and death. Covid FPD time is ca. 26±20w for all countries, 23w (16-33w) for NL, ca. half of temperature’s FPD time. This suggests covid targets people more frail than natural mortality causes do, in line with findings of higher average age at death by covid compared to all-cause-mortality.

VE against covid mortality is found at ca. 100% initially, and wanes very fast in months to a year, matching existing studies. VFR is 0.13±0.1% death/dose for all countries, 0.09% (0.06-0.15%) for NL and 0.22% (0.12-0.41%) for DE. Vaccine HPA time between dose and death ranges from months to a year. Vaccine’s FPD time was measured an order above temperature’s, ranging from years to decades. This suggests vaccine fatal adverse events target much healther people than natural mortality causes do.

In all 10 countries examined, measured LTL and LTS have low accuracy and wide 95% CIs. Still, results suggest vaccines are responsible for a loss of lifeyears comparable to lifeyears saved. LTL by vaccinations in DE is 450-4600ky (kiloyear) and in NL 50-500ky for all-ages. Age-stratified NL results are 13-94ky (ages 0-65), 16-65ky (ages 65-80) and 5-34ky (ages 80+).

Conclusions: The proposed mortality model is simple, uses objective observations of mortality and temperature, fits and predicts mortality over decades, and outperforms the Dutch national baseline by a 79% prediction error variance reduction. Additional mortality determinants can be easily integrated. While pandemic determinants’ low reliability limits model accuracy, several results are significant and in line with other recent findings. Vaccine’s fatal adverse events appear

to target much healthier people than natural causes do, while covid targets people much frailer. Lifetime lost due to fatal adverse events is of the same order as lifetime saved by vaccines.

Message: The proposed model is very different from the state of the art and many options for further research are outlined. It has all kinds of flaws unknown to me, to all [Bre]. In its current simple form, the temperature-based model is very promising for generic mortality prediction. The pandemic model provides extremely disturbing findings on the mortal effect of covid vaccination campaigns in the young and healthy population that warrant more research in vaccine safety.

Bio

I am André Redert, born in The Netherlands (1971), received MSc (1995) and PhD (2000) degree from the Information and Communication Theory group at Delft University of Technology. I worked as senior scientist at Philips Research Eindhoven (2000-2010) on 3DTV, music analysis, health and identification technology, and as guest lecturer at Delft University (2005-2012). I am (co-)author of 25 conference papers, 7 journal papers and 29 patent applications.

Currently I am an indy app developer at Qneo (qneo.net, 2009-) and Rodotti (rodotti.nl, 2013-), focusing on entertainment, music and utilities. Among the apps are the classic cult app “Blower”, featured on CNN and the Ellen Degeneres show, that makes your mobile phone blow out air using an assymetry in the physics of sound production; “Hover Engine”, that synthesizes V8 engine sounds live when riding a hoverboard; and “Voice Synth”, a professional vocoder app, the only that can regenerate the beautiful voices from Cylon robot antagonists in the 1978 TV series Battlestar Galactica. On request, I make tools/prototyping apps for research purposes.

The name Rodotti is derived from prodotti, the Italian word for products. It started as a joke nickname, but stuck over the years. My parents took the Rodotti logo photo in 1975 at Rotterdam Blijdorp Zoo. The red tie was made by my grandmother, still have it.

In spare time, I make music and investigate worldwide excess mortality since 2020 as independent researcher. If you want to support my work, feel free to visit rodotti.nl/support or buymeacoffee.com/andreredert. It is very much appreciated. Thank you.

References

[Axf] C. Axfors and J.P.A. Ioannidis, “Infection fatality rate of COVID-19 in community-dwelling elderly populations”, doi.org/10.1007/s10654-022-00853-w

[Ben] C.S. Benn et al., “Randomized clinical trials of COVID-19 vaccines: Do adenovirus-vector vaccines have beneficial non-specific effects?”, doi.org/10.1016/j.isci.2023.106733

[Bre] N. Breznau et al., “Observing Many Researchers Using the Same Data and Hypothesis Reveals a Hidden Universe of Uncertainty”, doi.org/10.31222/osf.io/cd5j9

[Cbs1] Dutch Central Bureau for Statistics (CBS), cbs.nl/nl-nl/longread/statistische-trends/2022/sterfte-oversterfte-en-covid-19-sterfte-in-2020-en-2021/2-ontwikkelingen-in-de-totale-sterfte-en-oversterfte and a supplement for VE: https://www.cbs.nl/-/media/_pdf/2022/25/sterfte_oversterfte_bijlage-4.pdf

[Cbs2] Statline of CBS, opendata.cbs.nl

[Cbs3] cbs.nl/nl-nl/nieuws/2021/02/bijna-4-duizend-mensen-overleden-in-eerste-week-van-2021/oversterfte-en-verwachte-sterfte

[Cbs4] cbs.nl/nl-nl/dossier/dossier-asiel-migratie-en-integratie/hoeveel-immigranten-komen-naar-nederland

[Cbs5] cbs.nl/nl-nl/visualisaties/dashboard-bevolking/bevolkingsgroei/overlijden and https://www.cbs.nl/nl-nl/nieuws/2020/40/10-duizend-coronadoden-tijdens-eerste-golf-van-de-pandemie

[Cbs6] cbs.nl/nl-nl/longread/rapportages/2023/oversterfte-en-doodsoorzaken-in-2020-tot-en-met-2022, cbs.nl/nl-nl/longread/rapportages/2024/covid-vaccinatiestatus-en-sterfte, see below *)

[Eur] Eurostat data browser, Mortality "Deaths by week and sex", ec.europa.eu/eurostat/databrowser/view/DEMO_R_MWK_TS/default/table, and Positive covid tests and vaccination doses in Europe, ecdc.europa.eu

[Fdb] Financiëel dagblad, https://beurs.fd.nl, Dutch AEX stock index history

[Ger] D. Ger et al., "Late-Life Decline in Well-Being Across Adulthood in Germany, the UK, and the US: Something is Seriously Wrong at the End of Life", doi.org/10.1037/a0017543

[Ghi] S. Ghisolfi et al., "Predicted COVID-19 fatality rates based on age, sex, comorbidities and health system capacity", British Medical Journal, Global Health, https://gh.bmj.com/content/bmjgh/5/9/e003094.full.pdf

[Høe] T.B. Høeg et al., "Potential 'Healthy Vaccinee Bias' in a Study of BNT162b2 Vaccine against Covid-19", nejm.org/doi/full/10.1056/NEJMc2306683

[Hul] N. Hulscher et al., "A Systematic Review of Autopsy findings in deaths after covid-19 vaccination", doi.org/10.1016/j.forsciint.2024.112115

[Ioa] J.P.A. Ioannidis, "Why Most Published Research Findings Are False", doi.org/10.1371/journal.pmed.0020124

[Knm] KNMI, "Klimatologie - Metingen en waarnemingen", knmi.nl/nederland-nu/klimatologie-metingen-en-waarnemingen

[Kuh] C. Kuhbandner, M. Reitzner, "Estimation of Excess Mortality in Germany During 2020-2022", https://www.cureus.com/articles/149410-estimation-of-excess-mortality-in-germany-during-2020#!/

[Lev] M. Levitt et al., "Excess death estimates from multiverse analysis in 2009-2021", doi.org/10.1101/2022.09.21.22280219, and "Comparison of pandemic excess mortality in 2020-2021 across different empirical calculations", http://doi.org/10.1016/j.envres.2022.113754

[Low] A.C. Lowen et al., "Influenza Virus Transmission Is Dependent on Relative Humidity and Temperature", doi.org/10.1371/journal.ppat.0030151

[Mar1] L.C. Mar et al., "Mechanistic insights into the effect of humidity on airborne influenza virus survival, transmission and incidence", doi.org/10.1098/rsif.2018.0298

[Mar2] I.C. Marschner, "Estimating age-specific COVID-19 fatality risk and time to death by comparing population diagnosis and death patterns: Australian data", bmcmedresmethodol.biomedcentral.com/articles/10.1186/s12874-021-01314-w

[Mee] R. Meester et al., "A likelihood analysis of COVID-19 mRNA vaccine safety in the third booster campaign in The Netherlands", researchgate.net/publication/368691068_A_likelihood_analysis_of_COVID-19_mRNA_vaccine_safety_in_the_third_booster_campaign_in_The_Netherlands

[Moo] J.R. Moon et al., "Short- and long-term associations between widowhood and mortality in the United States: longitudinal analyses", doi.org/10.1093/pubmed/fdt101

[Nei1] M. Neil at al., "Discrepancies and inconsistencies in UK Government datasets compromise accuracy of mortality rate comparisons between vaccinated and unvaccinated", dx.doi.org/10.13140/RG.2.2.32817.10086

[Nei2] M. Neil at al., "The extent and impact of vaccine status miscategorisation on covid-19 vaccine efficacy studies", dx.doi.org/10.13140/RG.2.2.15216.67846

[Noa] National Centers for Environmental Information, ncdc.noaa.gov/cdo-web/search

[Pez] A.M. Pezzullo, “Age-stratified infection fatality rate of COVID-19 in the non-elderly population”, doi.org/10.1016/j.envres.2022.114655

[Poi] scholar.google.nl/scholar?q=mortality+poisson+model

[Red1] A. Redert, “Covid-19 vaccinations and all-cause mortality - a long-term differential analysis among municipalities”, dx.doi.org/10.13140/RG.2.2.33994.85447

[Red2] A. Redert, “Short-term Vaccine Fatality Ratio of booster and 4th dose in The Netherlands”, dx.doi.org/10.13140/RG.2.2.29841.30568

[Red3] A. Redert, “Causal effect of covid vaccination on mortality in Europe”, researchgate.net/publication/368777703_Causal_effect_of_covid_vaccination_on_mortality_in_Europe

[Red4] A. Redert, “Towards objective excess mortality measurements: Mortality’s natural weekly variability”, dx.doi.org/10.13140/RG.2.2.21994.90561

[Riv1] Rijksinstituut voor Volksgezondheid en Milieu (RIVM), “Monitoring sterftecijfers Nederland”, rivm.nl/monitoring-sterftecijfers-nederland, see below *)

[Riv2] Rijksinstituut voor Volksgezondheid en Milieu, Vaccination doses given in The Netherlands, fall/winter campaign 2023, rivm.nl

[Riv3] Rijksinstituut voor Volksgezondheid en Milieu (RIVM), “Eindrapport CBS en RIVM - sterfte en oversterfte in 2020 en 2021", rivm.nl/documenten/eindrapport-cbs-en-rivm-sterfte-en-oversterfte-in-2020-en-2021

[Sch] M. Schmeling et al., “Batch-dependent safety of the BNT162b2 mRNA COVID-19 vaccine”, doi.org/10.1111/eci.13998

[Sto] L. Stoeldraijer et al. “Bevolkingsprognose 2012-2060: model en veronderstellingen betreffende de sterfte”, cbs.nl/nl-nl/achtergrond/2013/26/bevolkingsprognose-2012-2060-model-en-veronderstellingen-betreffende-de-sterfte

[Wik] Wikipedia, “Proportional hazards model”, https://en.wikipedia.org/wiki/Proportional_hazards_model

[Wu] N. Wu et al., “Long-term effectiveness of COVID-19 vaccines against infections, hospitalisations, and mortality in adults: findings from a rapid living systematic evidence synthesis and meta-analysis up to December, 2022”, doi.org/10.1016%2FS2213-2600(23)00015-2

[Xia] Y. Xia et al., “Effects of ambient temperature on mortality among elderly residents of Chengdu city in Southwest China, 2016–2020: a distributed-lag non-linear time series analysis”, bmcpublichealth.biomedcentral.com/articles/10.1186/s12889-022-14931-x

[Zon] Dutch government subsidy institute ZonMW, zonmw.nl/nl/onderzoek-naar-redenen-en-oorzaken-van-oversterfte-de-coronacrisis

*) ad [Riv1], translated from Dutch:

*Every year in the first week of July, the RIVM calculates the expected mortality for the coming year. For this we use the mortality figures of the **past five years**. To avoid distortion of expected mortality, we **do not include previous peaks**. These peaks often coincide with cold and heat waves or outbreaks of infectious diseases. This concerns the **25% highest mortality rates** over the past five years and the 20% highest mortality rates in July and August. The calculation uses a **linear regression model with a linear time trend** and sine/cosine terms to describe possible seasonal fluctuations.*

*) ad [Cbs6], translated from Dutch:

*Vaccination status '**vaccinated' is defined as** 'fully vaccinated' (i.e. **two weeks after two approved vaccinations**, or a positive test at least 56 days before at least one approved*

vaccination, or four weeks after vaccination where one vaccination counts as fully vaccinated according to the vaccination certificate, or when a booster or repeat injection has been given without a known basic series) possible with boosters and repeat injections. Vaccination status ***'unvaccinated' is defined as*** *no known vaccination or only* ***one known vaccination*** *without previously reported infection (with the exception of the vaccine where one vaccination counted as fully vaccinated).*