European Actuarial Journal
https://doi.org/10.1007/s13385-023-00372-2

ORIGINAL RESEARCH PAPER

Forecasting, interventions and selection: the benefits of a causal mortality model

Snorre Jallbjørn[1,2] · **Søren F. Jarner**[1] · **Niels R. Hansen**[1]

Received: 16 June 2022 / Revised: 7 June 2023 / Accepted: 1 November 2023

Abstract

Integrating epidemiological information into mortality models has the potential to improve forecasting accuracy and facilitate the assessment of preventive measures that reduce disease risk. While probabilistic models are often used for mortality forecasting, predicting how a system behaves under external manipulation requires a causal model. In this paper, we utilize the potential outcomes framework to explore how population-level mortality forecasts are affected by interventions, and discuss the assumptions and data needed to operationalize such an analysis. A unique challenge arises in population-level mortality models where common forecasting methods treat risk prevalence as an exogenous process. This approach simplifies the forecasting process but overlooks (part of) the interdependency between risk and death, limiting the model's ability to capture selection-induced effects. Using techniques from causal mediation theory, we quantify the selection effect typically missing in studies on cause-of-death elimination and when analyzing actions that modify risk prevalence. Specifically, we decompose the total effect of an intervention into a part directly attributable to the intervention and a part due to subsequent selection. We illustrate the effects with U.S. data.

Keywords Mortality modelling · Risk factors · Cause elimination · Interventions · Causality

✉ Snorre Jallbjørn
snorrejall@math.ku.dk

Søren F. Jarner
soren@jarner.dk

Niels R. Hansen
niels.r.hansen@math.ku.dk

[1] Department of Mathematical Science, University of Copenhagen, Copenhagen, Denmark

[2] Danish Labour Market Supplementary Pension Fund, Kongens Vænge 8, Hillerød, Denmark

Published online: 20 December 2023

1 Introduction

Soaring life expectancies throughout the industrialized world have prompted a rapid increase in research on mortality modelling and forecasting. Most models aim to produce accurate forecasts of all-cause mortality, relying solely on age, calendar time, and birth-cohort as predictors, e.g. [4, 24, 34]. While these models are effective at predicting future death rates, they do not incorporate information about the causal mechanisms underlying past trends. Instead, they assume secular linear trends at an aggregate level—an assumption often at odds with reality [17, 19].

To foster a deeper understanding of mortality and its drivers, recent studies have sought to enrich traditional models by incorporating the impact of risk behaviour on health. The goal is to disentangle the effects of general health improvements from those of risk behaviour change over successive generations. Booth and Tickle [3] characterize models exploiting relations between behavioural risks and death as explanatory models but warn that such relationships are still imperfectly understood. Nonetheless, various advances to make more precise and better substantiated forecasts by integrating health and lifestyle related trends have been made in recent years, e.g. [11, 18, 23, 33, 41].

Still, little work exists on the applicability of explanatory models to assess the effectiveness of health interventions aimed at reducing mortality rates in population-level forecasts. Studying how human longevity can be improved by curing or reducing the prevalence of existing diseases, or by modifying risk behaviour, is of great interest with many potential applications across numerous disciplines, including demography, actuarial risk management, and health economics. In practice, such analyses are challenging due to their inherent causal nature, which necessitates explicit assumptions about the data generating process and parameters that can be adjusted to represent interventions.

The difficulty in carrying out these analyses is twofold. First, causal models require careful consideration of the underlying mechanisms that lead to changes in mortality. This involves identifying the causal relationships between various risk factors and mortality, and inferring their strength. We elaborate on this issue in Sect. 2. Second, current population-level mortality frameworks have limited ability to take varying individual responses into account. Population heterogeneity creates a two-way feedback mechanism between risk prevalence and mortality, wherein the risk composition among the surviving population changes over time in accordance with the risk composition among those who die. This selection-induced mechanism is ignored by explanatory models that treat risk prevalence as an exogenous process. However, real-world consequences of interventions can sometimes only be correctly predicted if the feedback mechanism is incorporated into the analysis [22, 39]. This is the main theme of the paper.

The purpose of this paper is to discuss how mortality forecasts are affected by interventions and to show by example the assumptions and data needed for such an analysis to be operationalized. To facilitate this discussion, we define a generic causal mortality model in Sects. 3 and 4 using the framework of potential outcomes, and we emphasize when explicit modelling of the feedback mechanism is needed. In Sect. 5 we propose a method based on causal mediation theory that decomposes the total

effect of an intervention into a part directly attributable to the action and a part due to selection. The overall message is that if we want to make accurate statements about the effect of an intervention, that is if we are to quantify both the direct *and* the indirect effects of an intervention, risk prevalence must be endogenous to the mortality model. We consider this work a first step towards population-level forecasting of mortality under interventions.

2 When do we need a causal mortality model?

In general, modelling and forecasting (cause-specific) mortality is a problem of prediction, for which a probabilistic model suffices. That is, if we wish to predict mortality given the prevalence of concurrent health risks, π say, it suffices to specify a distribution over a lifetime, X, through a model of the conditional death rate $(x, \pi) \mapsto \mu(x|\pi)$. In theory, it is also possible to use such a model to study the impact of interventions by perturbing the distribution of variables that have been conditioned on as is done in conventional stress tests. That is, by computing $\mu(x|\pi^a)$ for some alternative risk distribution π^a. But this analysis will only give a realistic picture of the consequences of an intervention if the conditional distribution of mortality does not change when the predictors change.

For example, suppose that we are interested in studying the effect of the number of cigarettes smoked daily, S, on the risk of death. There is an abundance of evidence in the literature to support the conclusion that the more one smokes, the higher one's risk of death becomes. Suppose we focus on the regression task of learning $s \mapsto \mathsf{E}[X|S = s]$. It is then tempting to interpret $\mathsf{E}[X|S = s] - \mathsf{E}[X|S = 0]$ as the expected increase in survival for a smoker, who quits smoking. But such an interpretation is invalid in the presence of confounding factors. A confounding factor is any variable that influences both risk exposure and response. Here we could imagine the relationship depicted in Fig. 1. A lack of commitment to not smoke tends to co-occur with a higher susceptibility to being obese through various underlying social factors. Since associations between '*Smoking*' and the outcome may arise not only through the number of cigarettes smoked but also through underlying factors that determine general risk behaviour, the parameters associated with $\mathsf{E}[X|S = s]$ may have no causal meaning.

In comparison, a causal model is able to update its prediction in response to changing conditions in a way that is intrinsic to the model. This is because, in a causal model, the distribution of a variable given its causes is stable under interventions that only affect other variables. That is, the generative mechanism for a variable not targeted by an intervention is left intact. Several comprehensive books have been devoted to

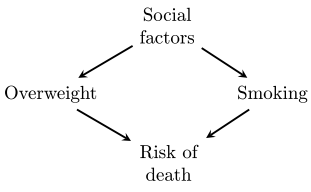

Fig. 1 Example of confounding. The arrows represent causal relations

the topic of causal inference and discovery, e.g. [14, 15, 31, 32, 38], and we will thus not give an in-depth account of the concepts and methods here. In the present paper we focus on decomposing the effects of interventions in population-level forecasts of mortality with a particular emphasis on the role of selection, that is how differences in risk can cause the composition of the population (and thereby death rates) to change over time.

2.1 Inferring cause–effect relationships is challenging

To give precise answers to causal questions, we need to invoke restrictive assumptions about the data generating process. A formal account of the ones needed can be found in, e.g., [14, 35, 37]. These may hold by virtue of study design, for instance in a randomized controlled trial, but generally we can only identify causal effects from observational data when there is no unmeasured confounding. Because no regard is paid to confounding factors in a standard regression analysis, these estimates cannot be endowed with a causal interpretation. This issue is also well recognized in the mortality forecasting literature, for instance by [23] who carefully remark: "*Indeed, none of our results should be seen as claims about the causal effects of obesity, smoking, or any other factor.*".

Estimating causal effects is an ambitious task requiring specialized methods, subject matter expertise, and detailed individual-level health data. To operationalize a causal mortality model, dose-response relationships must be based on epidemiological evidence from the literature, supported by trial and cohort data. Recent advances in the field of epidemiology, spearheaded by the Global Burden of Disease (GBD) initiative, may assist in bridging this gap between mortality and its determinants. In particular, [28] gives a standardized and comprehensive account of how 87 risk factors interact and affect different causes of death, covering in total 560 risk-outcome pairs based on a systematic review of partial studies. Using data from the GBD, it is possible to construct a causally interpretable mortality model with the aim of forecasting country specific mortality under varying scenarios. We take up this task in Sect. 6.

The modelling choices we make are close in spirit to the seminal work of [11], who also build an explanatory model based on the GBD estimates with the aim of substantiating mortality forecasts and exploring alternative health scenarios. While [11] does evaluate better/worse scenarios, these are made using a more conventional stress testing procedure, where the improvement rates of the risk factors are varied. They stress that such scenarios are to be understood as "[…] *a signal on the scope for policy change*", rather than actual alternative scenarios. In Sects. 3–5, we expand on their method of analysis and explain which additional model components are needed to evaluate actual interventions.

2.2 Two types of interventions

We aim to explicate the causal consequences of two types of interventions, namely on actions that target the death rates directly by cause-of-death elimination, and on

actions that modify behavioural risk factor prevalence thus targeting the death rates indirectly.

The impact of eradicating certain causes of death is a topic widely debated in the actuarial literature, dating back to Bernoulli's discussion in 1760 of a hypothetical world without small-pox [21]. The pivotal assumption made by Bernoulli was that individuals "saved" from the eliminated cause were as susceptible to dying from the non-eliminated causes as the general population, an assumption that still permeates most cause-deleted lifetable calculations today. Indeed, the prevailing methodology is to directly manipulate the cause-specific death rates of interest, while leaving remaining rates unaffected. This is commonly referred to as cause elimination under an assumption of independent competing risks, an approach that typically overstates the actual effect because it fails to account for subsequent selection, see also the discussions in [22, 39].

Some papers recognize the issue of dependence among competing causes in their estimates of cause-deleted life tables, e.g. [1, 8, 20, 25–27], but they do not explain the pathways through which dependence originates. Here, we provide an explanation by linking individual risk behaviour to cause-specific mortality. It is this link that allows us to explicate the consequences of selection following interventions.

As a motivating example, suppose that we are able to prevent all deaths due to lung cancer by some unusually successful targeted laser therapy. In this hypothetical world there will, at least initially, be fewer deaths compared to the world where the cause still operates. But since everyone eventually dies, deaths are ultimately redistributed among remaining causes. What is left is then to quantify how soon those "saved" die from something else, and what they die from instead. Because individuals who die from lung cancer are predominantly smokers, improvements in the treatment of lung cancer will indirectly affect (and most likely increase) the mortality rates for other tobacco-attributable causes such as heart diseases, following the progressive build-up of smokers in the population. The initial decrease in the aggregate death rate due to lung cancer being eradicated is thus partly offset by a subsequent "harvesting" of the "saved" smokers. We elaborate on this feedback mechanism in the next section, and formally decompose interventions into their direct and indirect effects in Sect. 5.

3 The feedback mechanism

To define interventions and their consequences on mortality, we will conceptualize how risk mechanisms at the level of individuals transfer to the level of populations in the framework of potential outcomes. We establish the basic relations in this section and use them to emphasize that the necessity of modelling the feedback mechanism, whereby mortality and risk prevalence influence each other in population-level models, is an inherent consequence of aggregation. We then give an instructive example that demonstrates the role of the mechanism in a scenario of cause-of-death elimination.

For ease of exposition, we consider the dynamics of a single ageing (birth) cohort followed until some maximum attainable age $\omega \in (0, \infty)$. Since age and calendar time advance synchronously in this case, we omit dependence on time in the following. Consider $i = 1, \ldots, n$ independent lives endowed with an age-varying vector of

categorical[1] covariates Z_i and life times X_i. For the purposes of this paper, we think of the covariates as modifiable lifestyle risk factors (e.g., smoking, weight, etc.), although disease history, socio-economic indicators and other variables that predict mortality could be included.

Denote by an overbar the history of the covariate process up to age x, that is $\bar{Z}_i(x) = (Z_i(u) : 0 \leq u \leq x)$. With probability one, we assume that $[0, \omega] \ni x \mapsto Z_i(x)$ only has finitely many jumps. Let $\bar{z}(x)$ be a possible (fixed) covariate trajectory and define $X_i^{\bar{z}}$ as an individual's *potential* life time had covariate exposure been $\bar{z}$ (possibly different to what was observed) with $\bar{z}$'s dependence on x suppressed. For each possible trajectory of $\bar{z}$ the distribution of $X_i^{\bar{z}}$ is completely characterized by the hazard rate

$$\mu^{\bar{z}}(x) := \lim_{dx \searrow 0} \mathsf{P}\left(x \leq X_i^{\bar{z}} < x + dx \mid X_i^{\bar{z}} \geq x\right)/dx. \tag{1}$$

The superscript identifies that $\mu^{\bar{z}}$ is the hazard function for $X_i^{\bar{z}}$. Of course, for a given individual we only observe one outcome and not all potential outcomes. We can relate the observed outcome to the potential outcomes through the assumption of consistency, see e.g. [14], namely $X_i = X_i^{\bar{z}}$ when $\bar{Z}_i(x) = \bar{z}(x)$ for all $x \in [0, X_i]$. In other words, the two outcomes coincide for the observed covariate trajectory.

3.1 A cause-specific relative risk framework

All-cause mortality is decomposed by considering $k \in \{1, \ldots, K\}$ mutually exclusive and exhaustive causes of death. This is a situation of competing risks, where different causes compete to end the life of an individual and occurrence of one event precludes occurrence of the remaining. The cause-specific hazard $\mu_k^{\bar{z}}$ characterizes the instantaneous rate of death from cause k in the presence of competing causes and is constructed such that $\sum_{k=1}^{K} \mu_k^{\bar{z}}(x) = \mu^{\bar{z}}(x)$ for all x.

To facilitate estimation and inference, some structure must be imposed on the hazard. In epidemiological and biostatistical applications where the inferential goal is to establish causal explanations for the etiology of disease and death, mortality is often studied in the relative risk framework. We adopt this framework throughout and assume that the death rate is related to $\bar{z}$ in a multiplicative fashion adhering to the form[2]

$$\mu_k^{\bar{z}}(x) = \mu_{0k}(x) R_k(\bar{z}(x), x), \quad k = 1, \ldots, K, \tag{2}$$

where μ_{0k} is a baseline death rate common to all individuals, and R_k is a known function governing the effect of risk exposure. Equation (2) can be classified as a marginal structural model in the sense that it provides a structural (and thereby *causal*)

[1] Covariates are, especially for large cohort studies, often reported as categorical variables even when the underlying exposure is continuous.

[2] One could also rewrite (2) to not condition on the entire covariate history of an individual but only on concurrent exposure. Prior behaviours could then be incorporated by making them explicit levels of the categorical covariates. We adopt this methodology in Sect. 6.

description of the marginal distribution of $X_i^{\bar{z}}$. This interpretation will be important when arguing about the effects of interventions.

3.2 Mortality with individual- and population-level information

To understand the interdependence between an individual's exposure to risk and probability of death, we associate to each lifetime, X_i, the multivariate counting process $N_i(x) = (N_{i1}(x), \ldots, N_{iK}(x))$ where $N_{ik}(x) = \mathbb{I}(X_i \leq x, \delta = k), k \in \{1, \ldots, K\}$, is a counting process registering whether or not individual i has died before or at age x from cause k, with δ designating cause. For every i, the process $(N_i(x))_x$ is adapted to the filtration $(\mathcal{F}_{ix})_x$ with

$$\mathcal{F}_{ix} = \sigma(N_i(s), Y_i(s), Z_i(s) : 0 \leq s \leq x) \tag{3}$$

being the σ-algebra generated by the internal history of the counting process augmented by the history of the covariate process Z_i and the at-risk process $Y_i(x) = \mathbb{I}(X_i \geq x)$. Assuming sufficient regularity, $N_i(x)$ has intensity process $\lambda_i(x) = (\lambda_{i1}(x), \ldots, \lambda_{iK}(x))$ given by

$$\lambda_{ik}(x) = \mu_{0k}(x) R_k(\bar{Z}_i(x), x) Y_i(x), \quad k = 1, \ldots, K, \tag{4}$$

with respect to $\mathcal{F}_{ix}$. The interpretation in infinitesimal terms is that the expected (local) change in the death process

$$\mathsf{E}[dN_{ik}(x) \mid \mathcal{F}_{ix-}] = \mathsf{E}\left[N_{ik}((x + dx)-) - N_{ik}(x-) \mid \mathcal{F}_{ix-}\right] = \lambda_{ik}(x)dx, \tag{5}$$

is a function of the individual's covariates and their survivorship status. We note that Y_i depends on the set of causes operating, but we suppress this in the notation for now.

In demographic and actuarial studies of mortality the focus is on the aggregate age-specific death rate. To relate the individual level model to the population level we marginalize over surviving individuals. Consider the aggregated counting process $N_\bullet(x) = (N_{\bullet 1}(x), \ldots, N_{\bullet K}(x))$ given by

$$N_{\bullet k}(x) = \sum_{i=1}^{n} N_{ik}(x), \quad k = 1, \ldots, K, \tag{6}$$

which has intensity process

$$\lambda_{\bullet k}(x) = \sum_{i=1}^{n} \lambda_{ik}(x) = \mu_{0k}(x) \sum_{i=1}^{n} R_k(\bar{Z}_i(x), x) Y_i(x), \quad k = 1, \ldots, K, \tag{7}$$

with respect to $\mathcal{F}_x = \bigvee_{i=1}^{n} \mathcal{F}_{ix}$. Since the covariates, Z, are categorical with sample paths for each individual taking only a finite number of values, we can assume a grouping of the individuals based on their covariate configurations. Let G be the finite

number of possible subgroups of covariate histories, and denote these configurations by z_g for $g \in \{1, \ldots, G\}$. The proportion of surviving individuals in group g at age x is $\pi_g(x) = Y_{\bullet g}(x)/Y_{\bullet\bullet}(x)$ where $Y_{\bullet g}(x) = \sum_{i=1}^{n} Y_i(x)\mathbb{I}(\bar{Z}_i(x) = z_g)$ is the number of individuals at risk in group g while $Y_{\bullet\bullet}(x) = \sum_{g=1}^{G} Y_{\bullet g}(x)$ is the total number at risk at age x. With $R_{kg}(x) = R_k(z_g, x)$ being the combined relative risk of group g for cause k, we can write (7) as a weighted average

$$\lambda_{\bullet k}(x) = \mu_{0k}(x) \sum_{i=1}^{n} \sum_{g=1}^{G} R_{kg}(x)\mathbb{I}(\bar{Z}_i(x) = z_g)Y_i(x) = Y_{\bullet\bullet}(x)\mu_{0k}(x) \sum_{g=1}^{G} R_{kg}(x)\pi_g(x). \tag{8}$$

Usually when working with population-level data, we do not have access to $\mathcal{F}_x$. Instead, data are often aggregated and thus only available in the form of tables of frequencies or in the form of histograms. Suppose therefore that the observed data available at age x is the aggregated history

$$\widetilde{\mathcal{F}}_x = \sigma(N_\bullet(s), Y_\bullet(s) : 0 \le s \le x), \tag{9}$$

where $Y_\bullet(x) = (Y_{\bullet 1}(x), \ldots, Y_{\bullet G}(x))$ is the number of individuals at risk in the G groups. Since the aggregated history is nested in the individual-level history, that is $\widetilde{\mathcal{F}}_x \subseteq \mathcal{F}_x$ for all x, we can apply the innovation theorem of [2] to (8). It follows that the intensity process of $N_\bullet$ with respect to $(\mathcal{F}_x)_x$ is also an intensity process of $N_\bullet$ with respect to $(\widetilde{\mathcal{F}}_x)_x$. Thus, the aggregated history yields the same population-level death rate

$$\mu_k^\pi(x) := \mu_{0k}(x) \sum_{g=1}^{G} R_{kg}(x)\pi_g(x), \quad k = 1, \ldots, K, \tag{10}$$

where $\pi(x) \in \{p \in [0,1]^G \mid \sum_{g=1}^{G} p_g = 1\}$ is the *potential* risk factor composition identified by the superscript π.

From a modelling perspective, when full information is available at the level of individuals, forecasting risk exposure prior to mortality is an admissible strategy. Since it is possible to calculate the aggregate death rate by marginalizing over surviving individuals, we can implement a modularized forecasting procedure by first determining the covariate dynamics conditionally on survival, and then the death rate given the covariate trajectory. However, the same level of modularity is not achievable with population-level models. If risk prevalence is exogenously given, as is common with explanatory mortality models, then there is no effect of mortality at age $s < x$ on risk prevalence at x. In this case, feedback due to selection is not part of the model, making the risk composition invariant to perturbations that cause a change in the death rates. Such perturbations will therefore not necessarily reflect the real-world consequences of an intervention. To build intuition for this problem, we conclude with an example of the role of the feedback effect. A more formal characterization of the feedback effect in a forecasting context is given in Sect. 5.

3.3 Changes in population composition due to feedback

Understanding mortality patterns produced by differential selection is a classical topic [39], and it is instructive to see how the feedback mechanism works when eradicating a cause-of-death in the simplest possible setting. Consider a closed population consisting of two homogeneous subgroups, differing only by their exposure to a two-level risk factor $Z \in \{1, 2\}$. Suppose that there are two causes operating in this world, governed by the individual-level model

$$\mu_k^z(x) = \begin{cases} \mu_{0k}(x), & \text{if } z = 1, \\ \mu_{0k}(x) R_k, & \text{if } z = 2, \end{cases} \tag{11}$$

with $R_k > 1$, $\mu_{0k}(x) > 0$ for all $x \geq 0$ and both causes-of-death $k \in \{1, 2\}$. At the population level, the cause-specific death rate is a weighted average of the healthy ($z = 1$) and the unhealthy ($z = 2$) subpopulations

$$\mu_k^\pi(x) = \pi_1(x)\mu_k^{z=1}(x) + \pi_2(x)\mu_k^{z=2}(x), \tag{12}$$

where $\pi_2(x)$ is the proportion of unhealthy individuals at age x, namely

$$\pi_2(x) = \frac{\pi_2(0)S^{z=2}(x)}{\pi_1(0)S^{z=1}(x) + \pi_2(0)S^{z=2}(x)} = \left[\frac{\pi_1(0)}{\pi_2(0)}\frac{S^{z=1}(x)}{S^{z=2}(x)} + 1\right]^{-1}, \tag{13}$$

with survival function $S^z(x) = \exp\{-\int_0^x (\mu_1^z(u)+\mu_2^z(u))\,\mathrm{d}u\}$, initial state $\pi_1(0)$, $\pi_2(0) \in [0, 1]$, and the requirement that $\pi_1(x) + \pi_2(x) = 1$ for all x.

Now, consider the "reference" world with both causes operating, $\mathcal{K} = \{1, 2\}$, and a hypothetical world in which cause 1 has been eradicated, $\mathcal{K}^* = \{2\}$. Because of competing risks, we need to be mindful that the cause k hazard is evaluated in the presence of other causes. We make this explicit in the notation now with π^{K} identifying the risk proportion and $S^{\mathsf{K},z}$ the survival function in a world where a specific (sub)set of causes $\mathsf{K} \subseteq \{1, 2\}$ are operating. By the assumptions above, we have that

$$\frac{S^{\mathcal{K},z=1}(x)}{S^{\mathcal{K},z=2}(x)} = \exp\left\{\int_0^x [\mu_{01}(u)(R_1 - 1) + \mu_{02}(u)(R_2 - 1)]\,\mathrm{d}u\right\} \geq \frac{S^{\mathcal{K}^*,z=1}(x)}{S^{\mathcal{K}^*,z=2}(x)}, \tag{14}$$

which, combined with (13), implies that $\pi_2^{\mathcal{K}}(x) < \pi_2^{\mathcal{K}^*}(x)$ for all $x > 0$. Thus, eradicating cause 1 weakens the selection mechanism resulting in a progressive build-up of unhealthy individuals, which makes the cause 2 death rate rise at the population level

$$\mu_2^{\pi^{\mathcal{K}^*}}(x) - \mu_2^{\pi^{\mathcal{K}}}(x) = \mu_{02}(x)(R_2 - 1)\left(\pi_2^{\mathcal{K}^*}(x) - \pi_2^{\mathcal{K}}(x)\right) > 0, \quad x > 0. \tag{15}$$

Equation (15) describes an indirect effect of cause removal brought about by a change to the risk composition through the feedback mechanism. We formalize the distinction

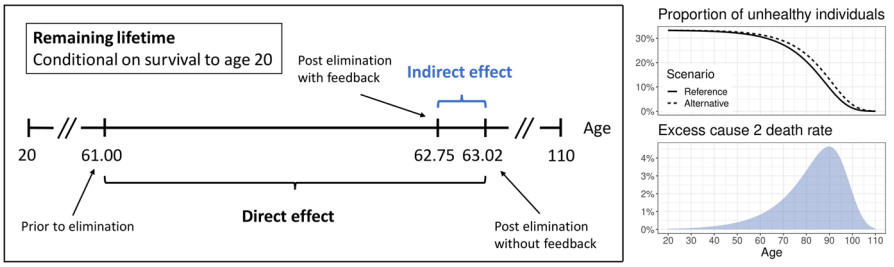

Fig. 2 Example effect of cause elimination for a cohort aged 20. The initial proportion of unhealthy individuals is $\pi_2(0) = 1/3$. Baseline mortality curves are given by $\mu_{01}(x) = \exp(-11.69 + 0.074x)$ and $\mu_{02}(x) = \exp(-10.58 + 0.088x)$ and relative risks by $R_1 = 5$ and $R_2 = 2.5$. The parameters are calibrated to reflect current death rates due to cancer and residual causes. The upper right panel visualizes the build-up of unhealthy individuals following cause-1 elimination, while the lower right panel pictures the subsequent harvesting. The left panel shows the remaining life expectancies prior and post elimination. In this example, the feedback effect reduces the life expectancy gained by about a quarter of a year

between direct and indirect effects of cause removal in Sect. 5. For now, notice that if the model ignored the feedback effect, that is if π^{K} was given exogenously, then the indirect effect would be zero since then $\pi^{\mathcal{K}}(x) = \pi^{\mathcal{K}^*}(x)$ for all x. The effects of cause removal are exemplified in Fig. 2.

4 A causal mortality model

For the remainder of the paper, we switch focus to a model spanning multiple birth cohorts and therefore consider a population defined in the rectangular age-period region

$$\mathcal{R}_{\text{data}} = \{(x,t) \mid x_{\min} \le x \le x_{\max}, t_{\min} \le t \le t_{\max}\}. \tag{16}$$

To give a proper justification for the declining mortality rates observed over the past centuries, one needs to model the influence of both individual and contextual factors on the risk of death. Contextual factors are the general living conditions to which all individuals are exposed, while individual factors may be divided into two types—observable and unobservable. We do not consider unobserved heterogeneity in the following, although this could be modelled using frailty theory as in [40]. Our focus is instead on individual differences relating to (observable) health and lifestyle related behaviour. Following the notation outlined in the previous section, we assume that the cause-specific death rate under a potential covariate trajectory $\bar{z}$ follows the relative risk model[3]

$$\mu_k^{\bar{z}}(x,t;C(t)) = \mu_{0k}(x;C(t))R_k(\bar{z}(x,t),x,t). \tag{17}$$

[3] The relative risk estimates of the GBD study [28] vary with age but not over time. Age-related changes are consistent with current epidemiological research which indicates that the relative effect of (most) risk exposures dissipate over the course of a life span. Time invariance is, however, only justifiable over short- to medium horizons as it renders the model unable to capture temporal changes in the effect of exposure.

Here, the pair $(x, t) \in \mathcal{R}_{\text{data}}$ identifies the cohort in question. The process $C(t)$ tracks the evolution of contextual variables like availability of food and water, access to healthcare services, improvements in medical technology, GDP change, and so on. It plays a dual role as a time-varying confounder process that must be controlled for to ensure that R_k has a causal interpretation, and acts as an effect modifier by stratifying the baseline death rate. The contextual variables relevant to us typically exhibit a secular trend over time, and disentangling their effect from the effect of calendar time can be complex and resource-intensive. In many cases calendar time is therefore used as a surrogate confounder by equating $C(t) = t$ for all t as a compromise.

The population-level equivalent of (17) under the assumption of categorical covariates as in Sect. 3 reads

$$\mu_k^{\pi}(x, t; C(t)) = \mu_{0k}(x, t; C(t)) \sum_{g=1}^{G} \pi_g(x, t) R_{kg}(x, t) = \mu_{0k}(x, t; C(t)) R_k^{\pi}(x, t), \tag{18}$$

where R_{kg} collects the risks of individuals in group g associated with each covariate and thus describes the combined relative risk of a given group. The aggregated model can also be viewed as a relative risks model in which the entire heterogeneous population has been collapsed into a single risk weighted individual with risk $R_k^{\pi}(x, t) = \sum_{g=1}^{G} \pi_g(x, t) R_{kg}(x, t)$.

5 Forecasting, interventions and selection

We turn to the impact of interventions on demographic mortality forecasts. Demographic forecasting is centered around the projection of population-level quantities such as the aggregate death rate. Predominantly, the stochastic processes considered are indexed by discrete time and standard time series methods are used for prediction. Even though we also view our data as time series in the following, all the points made can be extended to the continuous time case.

Our focus will be on time points in the forecast region $\tau \subseteq \mathbb{Z}$. We have three processes we need to consider jointly:

1. The mortality process $\mathcal{M} = \{\mu_k(x, t) : x_{\min} \leq x \leq x_{\max}, k = 1, \ldots, K\}_{t \in \tau}$,
2. The risk prevalence process $\pi = \{\pi(x, t) : x_{\min} \leq x \leq x_{\max}\}_{t \in \tau}$, and
3. The confounder process $C = \{C(t)\}_{t \in \tau}$.

To understand how interventions affect this system, we must describe how the processes influence each other, and, in particular, whether or not one process has predictive power over another. The concept of Granger causality [13] known from econometrics formalizes the notion of influence between processes, and is particularly useful for studying dynamic relationships in multivariate time series. We give a precise definition and explain how the concept is used to obtain graphical representations in Appendix A.

We can represent the above system by the graph shown in Fig. 3A. In the graph each process is represented as a single node with time being implicit. Two nodes are joined

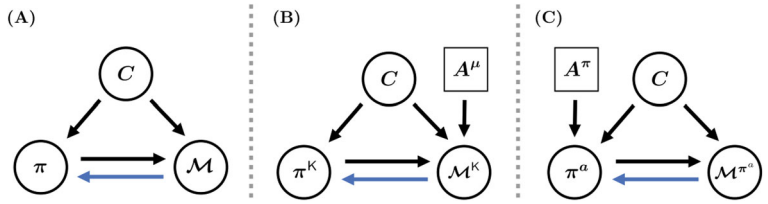

Fig. 3 Rolled graphs of models with and without interventions. The blue arrow encodes feedback. Panel **A** shows the unintervened setting describing the relationship between π, $\mathcal{M}$, and C. Panels **B** and **C** are graphs of models where the death process and the risk process are indexed by possible actions. In panel **B** the action is on the set of causes operating and in panel **C** the action is an intervention on the risk distribution (color figure online)

by a directed edge whenever a process at time t is predictive for another process at a future time $s > t$. For instance, in a model with feedback the cycle $\pi \to \mathcal{M} \to \pi$ represents a mutual dependence between $\mathcal{M}$ and π. The level of risk faced by the population at time t affects the death rate experienced between t and $t+1$, which in turn affects risk prevalence at time $t+1$. Conversely, the absence of an edge implies that a process is not predictive for another. Thus, in a model without feedback there is no arrow pointing from $\mathcal{M}$ to π because $\pi(t+h) \perp\!\!\!\perp \overline{\mathcal{M}}(t) \mid (\overline{\pi}(t), \overline{C}(t))$ for any $h \in \mathbb{N}_+$. This relation is asymmetric in the sense that the risk composition always predicts the death rate $\mathcal{M}(t+1) \not\perp\!\!\!\perp \overline{\pi}(t) \mid (\overline{\mathcal{M}}(t), \overline{C}(t))$. We remind the reader that an overbar denotes the history of the time series, cf. Sect. 3.

5.1 Cause-of-death elimination

Inspired by [9, 10], we consider a set of actions $A^\mu = (A_1, \ldots, A_K)$ that act on components of $\mathcal{M}(t)$ through all points in time $t \in \tau$. For our purposes, each A_k takes values in $\{0, 1\}$ describing two different regimes. Having $A_k = 0$ corresponds to no action on the k'th component, while $A_k = 1$ is an atomic intervention that forces $\mu_k(t)$ to be zero for all $t \in \tau$. More general interventions could also be considered but will not be pursued in the present paper.

We assume that intervening on the k'th component of $\mathcal{M}$ does not affect the remaining components or the remaining processes in the system other than through past variables that may develop differently depending on the intervention:

$$(C(t), \pi(t), \mathcal{M}_{-k}(t)) \perp\!\!\!\perp A_k \mid \left(\bar{C}(t-1), \bar{\pi}(t-1), \bar{\mathcal{M}}(t-1)\right), \quad t \in \tau, \quad (19)$$

where $\mathcal{M}_{-k}(t)$ denotes $\mathcal{M}(t)$ without the k'th component.[4] We can then represent an intervention on $\mathcal{M}$ by augmenting the graph in Fig. 3A with an additional source node A^μ pointing into $\mathcal{M}$ as shown in Fig. 3B. Because A^μ is a decision variable it is represented graphically by a box and indicates possible eradication of certain causes of death. The variable $\mathcal{M}^{\mathsf{K}}$ is a potential outcome indexed by the set of causes operating

[4] We note that A_k is not a stochastic variable thus altering slightly the meaning of the $\perp\!\!\!\perp$-symbol. Here, $\perp\!\!\!\perp$ expresses that the distribution of $(C(t), \pi(t), \mathcal{M}_{-k}(t))$ is the same regardless of the value of A_k, cf. [6].

following the action $A^{\mu} = a$. Risk prevalence π^{K} is likewise indexed by this set as it may develop differently depending on which components in $\mathcal{M}$ that are affected.

It follows that without feedback there is no causal effect of intervening in $\mathcal{M}$ on $\pi(t)$ for any $t \in \tau$. Thus the figure with the blue feedback edge removed represents a model where the action of cause removal only has a direct effect on the risk of death, because there is no indirect effect through variables between time points t and $t+h$, $h \in \mathbb{N}_+$. In other words, while the action does prevent specific types of death, thereby increasing the *absolute* number of deaths at later points in time because of competing risks, it changes neither the *relative* risk composition nor the death rates of non-eliminated causes. When the model includes feedback the risk process acts as an intermediate variable that mediates an additional effect through the loop $\mathcal{M} \to \pi \to \mathcal{M}$. In this case, cause-elimination weakens the selection mechanism and leads to larger (relative) concentration of high-risk individuals at future time points.

5.1.1 A decomposition of the death rate

The causal contrast of interest is the difference between the death rate in the reference world where all causes are operating compared to the rate in a world where only a subset of causes are operating. The all-cause death rate is

$$\mu^{\mathsf{K}}(x,t;C(t)) = \sum_{k\in\mathsf{K}} \mu_{0k}(x,t;C(t)) \sum_{g=1}^{G} \pi_g^{\mathsf{K}}(x,t) R_{kg}(x,t). \tag{20}$$

Here, μ^{K} is a single-world quantity where the K-index refers to both the set of causes entering the sum and the world in which π is evaluated. Examining the impact of an intervention $A^{\mu} = a^*$ that leaves only a subset of causes $\mathcal{K}^* \subsetneq \mathcal{K} = \{1,\ldots,K\}$ operating comes down to evaluating the difference

$$\mathrm{TE}(x,t) = \mu^{\mathcal{K}}(x,t) - \mu^{\mathcal{K}^*}(x,t), \tag{21}$$

which constitutes the total causal effect. We have left the conditioning on C implicit for readability. Comparing the total effect in the model without feedback to the total effect in the model with feedback does tell us something about how much of the effect is mediated via the risk process, but it does not give us a clean decomposition. Instead, we consider the standard definitions of natural direct and indirect effects from the mediation literature adapted to the present setup, cf. [30, 36].

We seek to measure the direct effect of the action $A^{\mu} = a^*$ associated with the arrow $A^{\mu} \to \mathcal{M}^{\mathcal{K}^*}$ separately from the indirect effect associated with the loop $\pi^{\mathcal{K}^*} \to \mathcal{M}^{\mathcal{K}^*} \to \pi^{\mathcal{K}^*}$. To this end, we introduce a cross-world model. Cross-world models specify a joint distribution of processes corresponding to different values of the action $A^{\mu} = a^*$. We introduce the cross-world quantity $\mu^{\mathcal{K}^*,\pi^{\mathcal{K}}}$ indexed by both $\mathcal{K}^*$ and $\mathcal{K}$ to denote the death rate in a world where causes $\mathcal{K} \setminus \mathcal{K}^*$ are eliminated, but where the risk process develops as if all causes were still operating. To achieve this, we need "to run" $\mathcal{M}^{\mathcal{K}}$ simultaneously to drive the risk process $\pi^{\mathcal{K}}$. Note that $\mu^{\mathcal{K},\pi^{\mathcal{K}}} = \mu^{\mathcal{K}}$ and $\mu^{\mathcal{K}^*,\pi^{\mathcal{K}^*}} = \mu^{\mathcal{K}^*}$.

The total effect (21) may now be decomposed into a part directly attributable to cause removal, i.e. the expected change in μ induced by replacing the set of causes $\mathcal{K}$ with $\mathcal{K}^*$ while keeping the "mediator" fixed at its reference value $\pi^{\mathcal{K}}$, and an indirect effect relayed through the mediating variable. We write

$$\mathrm{TE}(x,t) = \underbrace{\mu^{\mathcal{K},\pi^{\mathcal{K}}}(x,t) - \mu^{\mathcal{K}^*,\pi^{\mathcal{K}}}(x,t)}_{\stackrel{\text{def}}{=} \mathrm{DE(x,t)}} + \underbrace{\mu^{\mathcal{K}^*,\pi^{\mathcal{K}}}(x,t) - \mu^{\mathcal{K}^*,\pi^{\mathcal{K}^*}}(x,t)}_{\stackrel{\text{def}}{=} \mathrm{IE(x,t)}}, \tag{22}$$

where the natural direct (DE) and indirect (IE) effects are given by

$$\mathrm{DE}(x,t) = \sum_{k \in \mathcal{K}\setminus\mathcal{K}^*} \mu_{0k}(x,t;C(t)) \sum_{g=1}^{G} \pi_g^{\mathcal{K}}(x,t) R_{kg}(x,t), \tag{23}$$

$$\mathrm{IE}(x,t) = \sum_{k \in \mathcal{K}^*} \mu_{0k}(x,t;C(t)) \sum_{g=1}^{G} \left[\pi_g^{\mathcal{K}}(x,t) - \pi_g^{\mathcal{K}^*}(x,t)\right] R_{kg}(x,t). \tag{24}$$

We note that (23) marks the change in μ caused by simply subtracting the death rates of causes $\mathcal{K}\setminus\mathcal{K}^*$ from the all-cause rate without adjusting risk prevalence. This action coincides with the notion of cause removal in the setting of independent competing risks in which elimination does not alter the composition of the surviving population.

5.2 Alternative risk prevalence distributions

Another type of intervention deals with the effect on mortality brought about by changing risk prevalence from the reference distribution π to some alternative distribution π^a. We consider a set of interventions $A^\pi = \{A(t)\}_{t\in T}$ for a subset of time points $T \subseteq \tau$. Each $A(t)$ can be represented as a point in the $G-1$ dimensional probability simplex $\{p \in [0,1]^G \mid \sum_{g=1}^{G} p_g = 1\}$, augmented by an additional state $\emptyset$ that represents no action. We assume that an intervention on $\pi(t)$ is not predictive for earlier or remaining contemporaneous variables:

$$(\overline{\mathcal{M}}(t), \overline{\pi}(t-1), \overline{C}(t)) \perp\!\!\!\perp A(t), \tag{25}$$

and that future variables are unaffected by the intervention other than through past variables:

$$\{\overline{Q}(t+h)\}_{h\in\mathbb{N}_+} \perp\!\!\!\perp A(t) \mid \overline{Q}(t) \tag{26}$$

where $Q(t) = (\mathcal{M}(t), \pi(t), C(t))$. An intervention on π is then represented graphically as in Fig. 3C with π and $\mathcal{M}$ indexed by the action $A^\pi = a$.

Measuring again the total causal effect on the risk difference scale, we have

$$\mu_k^{\pi^{\emptyset}}(x,t;C(t)) - \mu_k^{\pi^a}(x,t;C(t)) = \mu_{0k}(x,t;C(t)) \sum_{g=1}^{G}\left[\pi_g^{\emptyset}(x,t) - \pi_g^a(x,t)\right] R_{kg}(x,t) \tag{27}$$

for cause k. The total effect can be decomposed in a similar manner to what we did when the intervention was on $\mathcal{M}$. We introduce the cross-world quantity π^{a,a^*} where the action on π is a, but the death rates behave as if it were a^*. The cross-world process $\pi^{a,\emptyset}$ is thus the risk process when the action is a but with the death process developing as if no intervention has been made. Omitting the dependency on C for readability, we can then write the total effect of the action $A^\pi = a$ on cause k as

$$\mu_k^{\pi^{\emptyset}}(x,t) - \mu_k^{\pi^a}(x,t) = \underbrace{\mu_k^{\pi^{\emptyset,\emptyset}}(x,t) - \mu_k^{\pi^{a,\emptyset}}(x,t)}_{\text{natural direct effect}} + \underbrace{\mu_k^{\pi^{a,\emptyset}}(x,t) - \mu_k^{\pi^{a,a}}(x,t)}_{\text{natural indirect effect}}. \tag{28}$$

Because all effects are effectively mediated via π itself, the decomposition is complicated to interpret. The direct effect is the effect as if there were no feedback. It describes the change in the death rate following one or more perturbations of the risk prevalence distribution in a world where mortality does not influence risk prevalence. In other words, the risk prevalence prediction is unaffected by the fact that the risk composition among those dying in the $\pi^{\emptyset}$-regime is different from the composition in π^a-regime.

The indirect effect describes a self-exciting change to the death process due to it developing differently within the π^a-regime. To better understand it, it is helpful to have a concrete example in mind. Suppose that we intervene on a "marginal" risk factor distribution. For example, consider a situation where the proportion of obese has been substantially increased. Because of competing risks, influencing the risk of one event affects the risk of all events (on a population level). As a result, we can expect smokers to die sooner on average than they would in the absence of the intervention, since their risk of dying from obesity-related causes has gone up. Conversely, we can expect the death rates for tobacco-related causes to decrease, even for diseases that are directly linked to smoking, such as chronic obstructive pulmonary disease. This phenomenon, where an increase in one risk factor appears to protect against another, is selection-induced false protectivity, and it is a consequence of the feedback mechanism.

5.3 Analyzing interventions in the presence of feedback

The rationale for a population-level model that explicitly accounts for the feedback mechanism lies in its ability to maintain internal consistency. By predicting risk prevalence endogenously based on a potential risk composition at projection jump-off, this type of model can accurately capture the complex and dynamic nature of selection.

In contrast, models that treat risk prevalence as an exogenous process are limited in their ability to produce consistent patterns of selection, as these patterns are implicitly given by the risk prevalence forecast.

In summary, to analyze the effect of interventions in the presence of selection-induced feedback we require data on cause-specific death counts, $D_k(x,t)$, with corresponding exposure to risk estimates, $E(x,t)$, and group-wise risk factor prevalence proportions, $\pi_g(x,t)$ for $(x,t) \in \mathcal{R}_{\text{data}}$. Furthermore, we need data such that the relative risk function $R_k(z_g,x,t)$ and the confounder process $C(t)$ in Eq. (17) are known. Consequently, only the baseline model $\mu_{0k}(x,t;C(t),\theta)$ in Eq. (18), parametrized in terms of some vector θ, needs to be specified and estimated (for each k) to model mortality in the data window. Next, to project mortality, a joint forecasting procedure of risk prevalence and mortality is required. In practice, this involves specifying the dynamics of any time-varying parameters in θ, and the dynamics of risk prevalence conditionally on survival. If calendar time is not used as a surrogate for C or if R_k is time-varying, then the time dynamics of these processes would also have to be specified. The chosen joint model is used to produce a reference scenario (e.g., a best-estimate projection) and any desired alternative scenario. Finally, the causal quantities of interest, e.g. risk or life expectancy differences, are computed.

6 An application to US data: illustrating the direct and indirect effects of cause-of-death elimination

We consider an application of the methodology outlined in the previous sections to U.S. risk and mortality data. To keep the exposition concise we restrict the analysis to the SNAP risk factors: smoking, poor nutrition, excess alcohol consumption, and insufficient physical activity. These four modifiable lifestyle related risks are associated with most causes of death.

6.1 Data sources

We use the relative risk estimates of [28], part of the Global Burden of Disease initiative, to describe the link between risk exposure and mortality. The estimates are reported as time homogeneous quantities by sex and 5-year age groups. To get single age estimates we perform linear interpolation with the age bucket centroids as fixed points, see Appendix C.

For smoking the risk-outcome relationship is listed by either current number of cigarettes smoked daily or by pack-years. Pack-years collapse smoking intensity and duration into a single variable, so that we do not have to condition on the entire smoking history of an individual. One pack-year is the equivalent of having smoked one pack of cigarettes (20) a day for a year. For a given sex, age, and risk-outcome pair the exposure category is listed in jumps of 10. We use natural cubic spline interpolation between categories to obtain a continuous dose-response curve, see Appendix C.

As an indicator for nutritional status we use the Body Mass Index[5] (BMI), which is the dominant metric for categorizing individuals in terms of weight excess or deficiency. The relative risk is reported per five-unit change in BMI with 20 to 25 $\mathrm{kg/m^2}$ being the baseline category. Risks for alcohol consumption are reported directly in terms of grams consumed per day while the relative risk for physical activity is measured in metabolic equivalents (METs) with one MET being the rate of energy expenditure at rest.

Cause of death data is extracted from [5] and contains U.S. specific mortality and population data through the years 1999–2018. The data is based on death certificates on which a single underlying cause of death is registered. Matching the data with risks from the GBD study of [28], we consider in total 35 causes of death known to be influenced by the risk factors. These causes make up about two-thirds of the total age-specific deaths in the population above the age of 35. To obtain an exhaustive list of causes such that the sum of the cause-specific rates equals the all-cause rate, remaining causes are collected and aggregated into a 'residual' category and assigned a relative risk of one for all risk factors.

Risk prevalence data is collected from the IPUMS National Health Interview Survey (NHIS) database [16]. The NHIS is a large cross-sectional survey conducted annually by the U.S. government and contains comprehensive health and behaviour data at the level of individuals. The IPUMS NHIS data relies on sampling weights to produce representative estimates. Each unit of study can thus be inflated such that the sum of the weighted units constitutes the entire U.S. population. The present analysis is based on adult individuals covering ages 20–84 and years 1999–2018. Observations with missing data are placed into the baseline category. Pack-years of smoking exposure is constructed by assuming that the amount someone currently smokes has not changed since they began smoking. Exposure among former smokers is estimated using years since cessation and average cigarette consumption of the respective cohort. Figure 4 shows the evolution of risk prevalence over time.

6.2 Baseline model

For modelling purposes we assume that the (true) hazard rate μ_k is constant over the squares $[x, x+1) \times [t, t+1)$ for integer ages x and calendar years t. We consider data on the form given in Sect. 5.3. We can estimate the parameters associated with μ_k in (18) via maximum likelihood using the customary Poisson assumption

$$D_k(x,t) \mid E(x,t), R_k(x,t), C(t) \overset{\text{indep.}}{\sim} \text{Pois}\left(E(x,t)R_k(x,t)\mu_{0k}(x,t;C(t),\theta)\right). \tag{29}$$

Contrary to all-cause mortality that is generally well-behaved as a function of age, cause-specific mortality may exhibit several structural changes over the age span. Sudden rapid increases, periods of constancy, and even declines are not unusual. We

[5] Body Mass Index $:= \frac{\text{weight in kilograms}}{(\text{height in meters})^2}$. A BMI below 18.5 is considered underweight and a BMI in the range 25–29.99 is considered overweight. A BMI of 30 or above is classified as obese, subdivided into three categories: 30–34.99 is Class I, 35–39.99 is Class II, and 40 or greater is Class III.

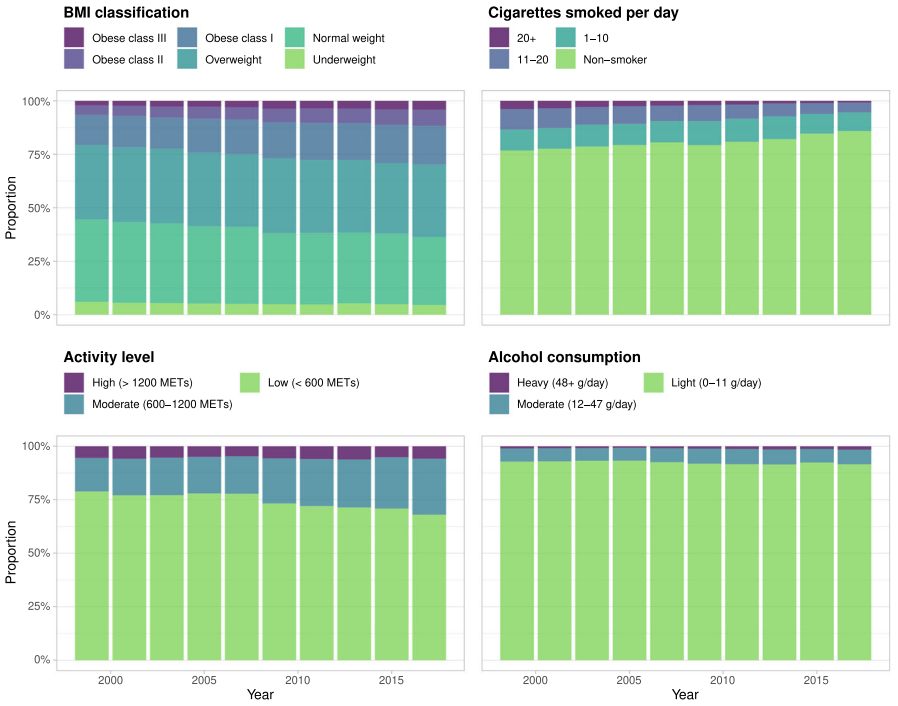

Fig. 4 U.S. risk proportions of BMI, smoking, alcohol consumption, and physical activity based on IPUMS data for both sexes and ages 20–84. The data shown in the figure is aggregated for the purpose of visual presentation. More granular data is used in the application

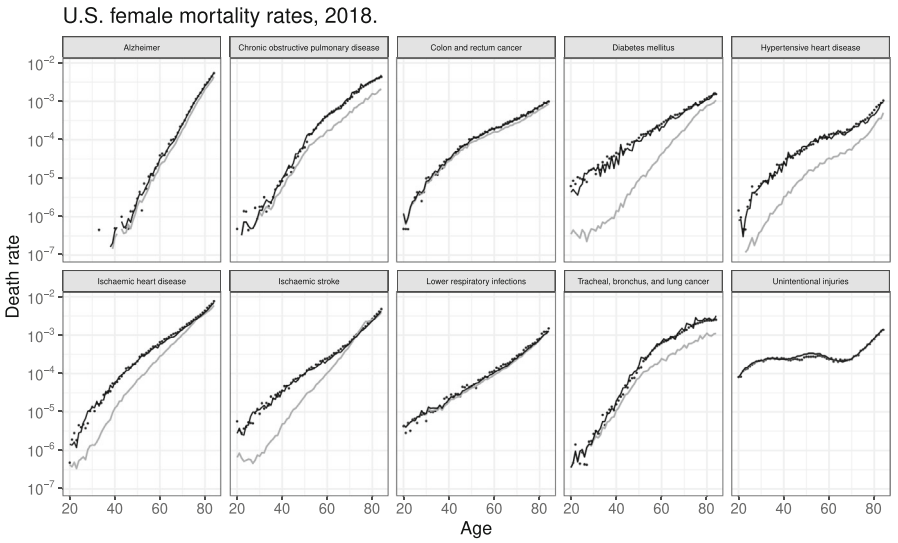

Fig. 5 Empirical (dotted), fitted (solid black), and baseline (solid grey) death rates for the top 10 leading causes of death in the dataset. The gap between the black and grey lines expresses the excess risk faced by the population due to deviations from baseline levels of exposure in the risk factors considered

could in principle use different functional forms to model μ_{0k} depending on cause, however stating just a single parametric form that generalizes well to most settings might be preferable in terms of interpretability. A simple yet widely used parametric form is the log-linear model

$$\mu_{0k}(x,t;C(t),\theta_k) = \exp\left(\theta_{k0x} + \theta_{k1x}t\right) = \mu_{0k}(x,t;\theta_k), \tag{30}$$

which has been applied in settings similar to ours, for instance by [11, 23]. The model is easy to estimate (see Appendix B), flexible enough to capture the different shapes associated with cause-specific mortality, and reflects that age is generally the most important driver of mortality regardless of risk exposure. We use (30) as the baseline model in what follows. Figure 5 shows the empirical female cause-specific death rates for the last year in the estimation period with fits of the aggregate and baseline rates superimposed.

6.3 Joint forecasting

Generally speaking, straightforward extrapolative approaches for forecasting risk prevalence are not recommended as they lead to an unabated continuation of historical trends and likely poor out-of-sample performance. Many researches resort to models that are specifically tailored to project the risk prevalence distributions in question, but existing methods are confined to working on the marginals and do not capture selection-induced feedback effects either. Developing a scalable joint forecasting procedure is an important topic of research, but it is beyond the scope of this paper. For the demonstration we have in mind we make do with a somewhat elementary state-transition model.

We aim at extrapolating the cohorts available in our sample until they reach age $x_{\max}$. We assume that there is no migration in or out of the composite population. Define the (one-step) survival probabilities

$$p_g(x,t) = \exp\left(-\sum_{k=1}^{K}\mu_{0k}(x,t)R_{kg}(x,t)\right), \tag{31}$$

for group g and denote by $m_{ij}(x,t)$ the probability of the cohort aged x at time t changing its "risk-group membership" from i to j. We employ a cohort state-transition model

$$Y(x+1,t+1) = M(x,t)Y(x,t), \tag{32}$$

with $Y(x,t) \in \mathbb{N}^G$ being the number of individuals in the G groups and $M(x,t) \in [0,1]^{G\times G}$ a matrix of transition probabilities with elements $M_{ij}(x,t) = p_i(x,t)m_{ij}(x,t)$. Note that as a consequence of this formulation, migration rates are only applied to the surviving population. The estimated transition probabilities m_{ij} are stated in Appendix D. To aid in understanding the output of

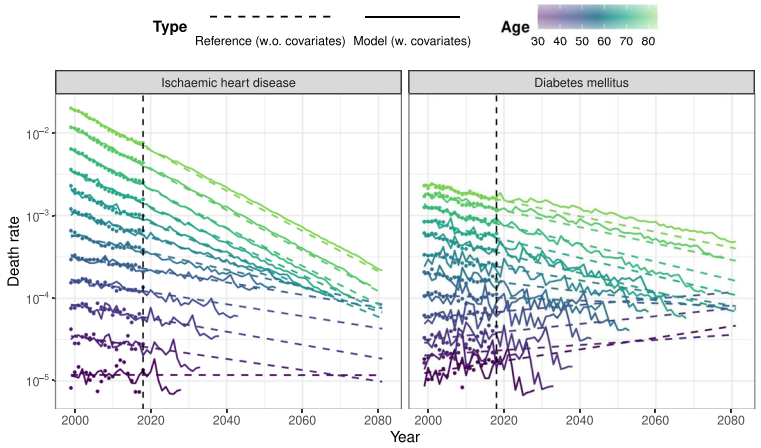

Fig. 6 Historical (dotted) and estimated and projected (lines) female rates for diabetes mellitus and ischaemic heart disease. The dashed lines are reference projections using the baseline model fitted without additional covariates

the model (32), Fig. 10 in Appendix D shows the mean trajectory of the continuation of the compositions from Fig. 4.

6.3.1 Example death rate forecast

Figure 6 shows the empirical and forecasted female rates for ischaemic heart disease, the leading cause of death in the dataset, and diabetes for select ages. To gauge the effect of including additional covariates in the forecast, the superimposed dashed lines are reference projections using the baseline model (30) fitted without additional covariates.

Mortality projections are usually based on empirical regularities such as smooth age profiles and small incremental mortality improvements, but the present projection depends heavily on the cohort specific exposures causing it to exhibit a rather erratic behaviour.[6] Figures 5 and 6 suggest that this is particularly so for diabetes as a large proportion of mortality is attributable to obesity and cohorts evidently differ substantially in their exposure.

On the other hand, the inclusion of covariates caters for the fact that (baseline) mortality levels ought to be consistently declining over time. While all-cause mortality adheres to this pattern, historical cause-specific rates may have actually increased over time—even recently as Fig. 6 shows—for some causes and ages. This is an ever-present issue widely acknowledged in cause-specific forecasting. The problem is that increasing rates generally do not express that health care and treatment options have worsened, but that mortality improvements have been substantially offset by changes

[6] For practical applications some smoothing is warranted. The type of smoothness violations seen here prompted the Bayesian modelling approach developed by [12], applied by e.g. [23] and [11], that down-weighs risk factor information if contradicted by observed empirical patterns. These papers also use smoothed prevalence estimates, whereas we simply apply the raw data.

to the risk prevalence distribution. A purely extrapolative model is not able to explain this development and will simply continue the observed trend unabated as seen in the reference forecast for the youngest age groups in Fig. 6. In the long run this may result in the aggregate all-cause forecast being dominated by the causes that have increased historically.

The causal model with covariate information is, in contrast, equipped to analyze the historic evolution of death rates at a granular level and may disentangle the effects of lifestyle related habits changing from generation to generation from general health care improvements. Indeed, the model successfully separates risk prevalence and mortality in this case by yielding negative slopes for the baseline for all age groups considered in Fig. 6. This shows that baseline mortality has improved, despite the immediate trend in the raw rates suggesting otherwise, and hints at why cause-specific modelling without additional information or assumptions ought to be avoided.

6.4 Cause-of-death elimination: what happens if cancer were eradicated?

We now seek to answer the central question posed in the beginning of the paper. If certain causes of death are eradicated, how soon will the individuals "saved" die from something else and what will they die from instead? We illustrate this query using the U.S. dataset by considering an elimination of deaths due to cancers. To answer the questions precisely, we look at the cumulative incidence

$$F_k(u+x \mid x) = \frac{1}{S(x)} \int_0^u S(v+x)\mu_k(v+x)\,\mathrm{d}v, \tag{33}$$

i.e., the probability of dying from cause k before or at age $u+x$ conditionally on being alive at age x.

Figure 7 shows as an example how cumulative incidence for the cohort aged 60 in 2018 ($t_{\max}$) is affected by the intervention. The figure explains the probability of dying before or at age 84 ($x_{\max}$). The cumulative incidences add up to the total probability of death which is 38.7 percent prior to elimination. Cancers make up roughly a quarter of all deaths with the cumulative incidence being 9 percent. An elimination therefore initially causes the total death count to decrease to 29.7 percent of the original cohort, while adjusting for the subsequent redistribution of deaths brings it up to 31.6 percent. The figure decomposes the redistribution into a part attributable to competing risks and a part due to feedback, using the framework developed and calibrated over the previous subsections.

The decomposition allows us to compare the model with feedback to its non-feedback alternative, namely the same model but with the feedback mechanism disengaged. Without feedback, individuals saved from cancer die according to the rates observed in the population prior to the intervention. Because of competing risks this leads to a rise in the cumulative incidence for every non-eliminated cause (grey bars). This change occurs despite the fact that the corresponding death rates are unaltered. Because the risk prevalence distribution among individuals who previously died from cancer is not the same as that of the general population, there is an additional

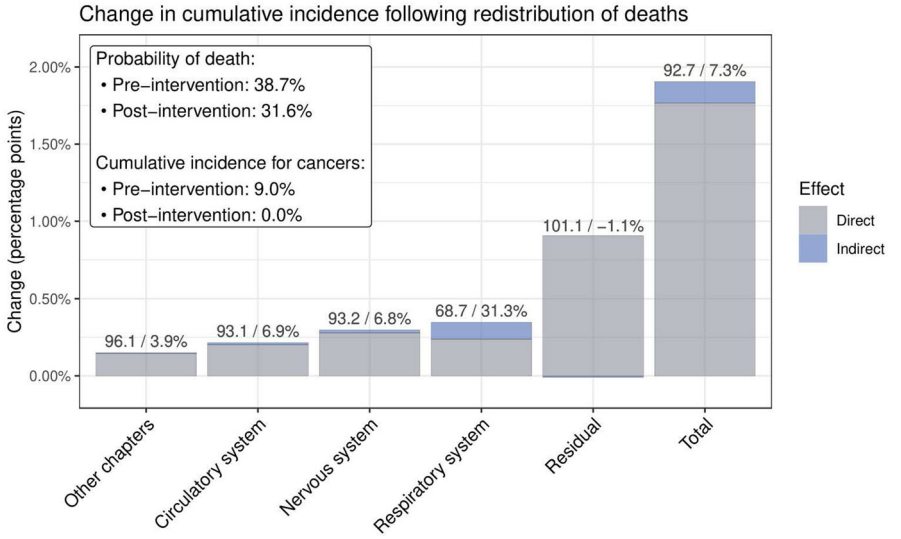

Fig. 7 Change to the cumulative incidence of select disease chapters following the elimination of deaths due to neoplasms affected by smoking and/or obesity for the female cohort aged 60 in 2018. The bars explain how the 9 percent of the cohort who previously died from cancer are redistributed into other categories. The percentages listed on top of the bars differentiate the part of the change due to direct and indirect effects respectively

effect (blue bars). In the model with feedback, higher-than-average-risk individuals are carried forward in the system, causing an increase in the death rates among remaining causes attributed to the SNAP risks. The number of deaths due to diseases of the respiratory system is particularly amplified.

Overall, a simple deletion of cancer death rates that disregards feedback will understate the total probability of death by 7.3 percent among those that are saved and by more than 30 percent at the cause-specific level. These percentages are naturally bounded by the number of risk factors included in the model. As additional risk factors are introduced, and as the departure from population homogeneity becomes more pronounced, selection-induced effects will carry even more weight.

6.4.1 Comparison to other non-feedback alternatives?

One might also take interest in comparing the method we have applied here to other non-feedback alternatives. Such a comparison is, however, beside the point that we are trying to make. We do not claim that our model is superior in predicting the reference scenario compared to other models. In fact, other models with carefully "sculptured" risk prevalence projections likely have better out-of-sample performance compared to the method we have used.

Our focus has been on finding and discussing the magnitude of second-order effects. We have shown that non-feedback models are unable to quantify these, making a comparison between the effect of an intervention based on our proposed method and an alternative somewhat fruitless. Often—especially when it comes to policy making—it

is tacitly assumed that second-order effects are small and can be more or less deliberately ignored. However, to reveal the extent of such an assumption, we need a method that allows us to realistically and consistently analyze the impact of interventions. We have detailed how to do so in this paper.

7 Concluding remarks

In this paper we discussed how mortality forecasts are affected by interventions in structural models that link individual risk behaviour to cause-specific mortality. We saw that when these risk mechanisms were specified at the level of populations, the model's ability to relay selection effects hinged on a feedback mechanism controlling how risk prevalence changed in response to differential mortality. We made the point that perturbations of the system only conform with real-world consequences of interventions when risk prevalence is endogenous to the model.

We considered how death rates changed following the eradication of certain causes of death. The prevalent approach directly manipulates the death rates of interest, with little or no regard for subsequent effects on non-eliminated rates. However, since individuals "saved" cannot be expected to follow the same pattern of mortality as that observed in the population prior to the intervention, these methods are too generous in their estimate of mortality reduction—but by how much? To disentangle and quantify the magnitude of indirect effects we applied techniques from causal mediation theory. This method gave us a straightforwardly interpretable decomposition of the total effect of cause-elimination with a part directly attributable to death rate deletion and a part due to disrupting the selection mechanism. The latter effect is, however, only quantifiable when risk prevalence is endogenous to the mortality model.

From a methodological perspective, our analysis of indirect effects is limited to those induced by changes in behavioural risks. Other health indicators, such as existing or developing medical conditions that have an impact on the length of life, may also be important contributing factors. To give one example, consider the COVID-19 vaccine which is highly effective at preventing serious disease, hospitalization, and death. Those who would have died from COVID without the vaccine may instead have a milder disease course although they could potentially still suffer from "long COVID". Such a delayed effect on their risk of death could be modelled using Barker frailty [29]. Those who would have survived even without the vaccine potentially never even contract the disease, thus producing a feedback that raises the vitality of the population. Further, by avoiding overcrowded hospitals due to COVID-related admissions, there could also be an effect on the general access to health care. This example shows that assessing all higher-order effects of an intervention can be extremely challenging and requires a comprehensive modelling framework.

From a practical perspective, mortality models with integrated epidemiological information are still in their infancy. A major challenge when building explanatory mortality models is the substantial data demand. Data is typically not available at a sufficient granular level to warrant a model at the level of individuals, and it is in fact rarely the case that a single authoritative source contains a complete set of the covariate distributions of interest, not even at an aggregate level. Instead, researchers

often have to collect (aggregate) prevalence data of marginal distributions from multiple sources. In time, however, as the availability and quality of detailed risk data continues to improve, causal models will inevitably gain a footing and contribute to more precise and better substantiated long-term projections of mortality. Moreover, the ability to formulate scenarios of interest in a straightforward and verbal manner is key to engaging non-specialist and making results accessible to a wider audience.

Acknowledgements The authors would like to thank Christian Bressen Pipper and two anonymous referees for their valuable input which helped improve the manuscript. The work was partly funded by Innovation Fund Denmark under File No. 9065-00135B.

Data availability The data used in the paper are available at https://mortality.org, https://wonder.cdc.gov/, https://nhis.ipums.org, https://ghdx.healthdata.org/gbd-2019.

Declarations

Conflict of interest The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A: Granger causality

In the following we give a brief overview of Granger causality and its use for describing conditional (in)dependence relations. For an in-depth account of causal reasoning in (graphical) time series models, we refer the interested reader to [9, 10] for the discrete time case and [7] for the continuous time analogue.

Granger causality was introduced by [13] and is a popular tool not only within its origin of econometrics, but also for causal time series analysis. Consider a multivariate time series $Q = \{Q(t)\}_{t\in\mathbb{Z}}$ with $Q(t) = (Q_1(t), \ldots, Q_d(t))^\top$. Let $V = \{1, \ldots, d\}$ be the index set and define for any $U \subseteq V$ the subprocess $Q_U(t) = (Q_u(t) : u \in U)$. Further, denote by an overbar $\overline{Q}_U(t) = \{Q_U(s)\}_{s\leq t}$ the history of the series. Let A and B be two disjoint subsets of V. We say that Q_A is *Granger non-causal* for Q_B up to horizon $h \in \mathbb{N}$ (w.r.t. Q) if

$$Q_B(t+l) \perp\!\!\!\perp \overline{Q}_A(t) \mid \overline{Q}_{V\setminus A}(t), \quad \forall\, l \in \{1, \ldots, h\}, t \in \mathbb{Z}. \tag{34}$$

Here the $\perp\!\!\!\perp$-symbol denotes independence. The formulation (34) of Granger causality tacitly assumes that all relevant variables for predicting Q are available in Q. This differs from the original formulation in which the information available is that of the "entire universe".

If (34) holds for $h = 1$, we say that Q_A is Granger non-causal for Q_B and we write $Q_A \nrightarrow Q_B$. Thus, a process Q_A is Granger non-causal for another process Q_B if the past of Q_A up to time t does not give a better prediction of Q_B at time $t+1$ given all information available up to time t but without that of Q_A. If Q_A is Granger non-causal for Q_B at all horizons we write $Q_A \overset{(\infty)}{\nrightarrow} Q_B$.

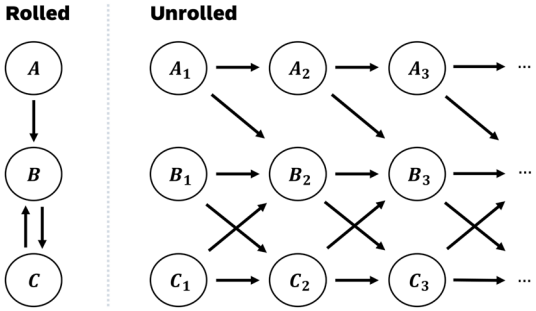

Fig. 8 An example of a rolling graph with three nodes A, B, and C and a corresponding unrolled version. Since, e.g., $A \to B$ in the rolled version, the unrolled version could contain edges from A_t to B_s for any $s > t$

A.1 Graphical representation

We can use the concept of Granger (non-)causality to obtain a graphical representation of the conditional independence relations of the time series. A graph $\mathcal{G} = (V, E)$ consists of a finite set of *nodes* V and a finite set of *edges* E. We only consider graphs containing *directed* edges, that is $E \subseteq V \times V$ is a subset of ordered pairs of nodes. We allow for multiple edges between two nodes if they are of different orientation in which case there is a *loop*.

Instead of a full time graph in which time is made explicit, we are primarily interested in a "rolled" version, also sometimes called a summary graph. To construct such a graph based on the time series Q, we partition the index set V into mutually disjoint subsets $A_1, \ldots, A_q$, $q \leq d$, and associate with the corresponding sub-processes the nodes $V = \{1, \ldots, q\}$. We join two nodes $a, b \in V$ by a directed edge if Q_{A_a} is Granger causal for Q_{A_b} at *some* horizon. Conversely, the absence of an edge implies that $Q_{A_a} \overset{(\infty)}{\nrightarrow} Q_{A_b}$. Although not explicitly shown in the graphs, we assume that all nodes have self-loops. Figure 8 gives an example of a rolled graph and a corresponding unrolled version.

Appendix B: Estimation of baseline parameters

Suppose we have data on cause-specific death counts, D_k, exposure-to-risk estimates, E, and relative risk coefficients, R_k, each of dimension $d_x \times d_t$ with $d_x = x_{\max} - x_{\min} + 1$ being the length of the age span and $d_t = t_{\max} - t_{\min} + 1$ being the length of the time span. The model (29) with baseline hazard function (30) is

$$D_k(x,t) \mid E(x,t), R_k(x,t) \overset{\text{indep.}}{\sim} \text{Pois}\left(E(x,t) R_k(x,t) \exp(\theta_{0x} + \theta_{1x} t)\right). \tag{35}$$

Since the predictor is linear we have the entire machinery of generalized linear models at our disposal. Using a Poisson error structure, canonical logarithmic link function, and stacking data into column vectors, that is, $d_k = \text{vec}(D_k)$, $e = \text{vec}(E)$ and $r_k = \text{vec}(R_k)$, we have that

$$\log \mathbb{E}[d_k \mid e, r_k] = \eta + \log(e \circ r_k) \tag{36}$$

where the latter term on the right-hand side is treated as an offset while $\eta = M\theta$ is the linear predictor with θ being the vector containing the parameters and

$$M = \left[1_{d_t} : (t_{\min}, \ldots, t_{\max})^\top\right] \otimes \mathsf{I}_{d_a} \tag{37}$$

being the model matrix. In the above, 1_d a d-dimensional vector of ones, I_d a d-dimensional identity matrix, $\circ$ the Hadamard product, and $\otimes$ the Kronecker product.

Appendix C: Interpolation of relative risks: examples

The relative risk estimates of [28] are reported as risk-outcome pairs by sex, age category, and exposure category. All quantities are time homogeneous.

For every risk-outcome pair the age category is listed in the groups 20–24, 25–29, ..., 90–94, 95–120. To obtain relative risk estimates for every (integer) age, we perform linear interpolation with the age bucket centroids $\mathcal{X} = \{20, 22, 27, \ldots, 92, 107.5, 120\}$ as fixed points. Thus, for fixed risk-outcome pair and sex and an age $x \in [x_0, x_1)$ where x_0 and x_1 are two consecutive numbers in $\mathcal{X}$, the relative risk at age x given by

$$\mathrm{R}(x_0) + (x - x_0)\frac{\mathrm{R}(x_1) - \mathrm{R}(x_0)}{x_1 - x_0},$$

where $\mathrm{R}(\cdot)$ supplies the relative risk estimate available in the data. An example is given in the left panel of Fig. 9.

For the risks "Number of cigarettes smoked daily" and "Pack years" the exposure categories are listed in jumps of 10, specifically by 0, 10, 20, 30, 40, 50, or 60 cigarettes per day and 0, 10, ..., 90, or 100 pack years. To obtain a dose-response curve for any (integer) number of cigarettes smoked per day or number of pack years, we extend

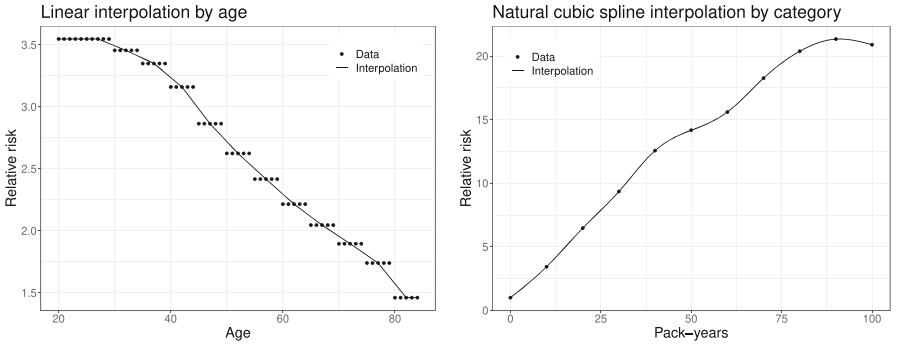

Fig. 9 Left panel: Female relative risk for diabetes mellitus by age (dots) using linear interpolation with age-bucket centroids as fix points (solid line). Right panel: Age 60 female relative risk for tracheal, bronchus, and lung cancer by pack-years (dots) using natural cubic spline interpolation (solid line)

Table 1 One-year smoking summary transition matrix

(pct.)	Smoking transition matrix			
	Non-smoker	1–10	11–20	20+
Non-smoker	100	0	0	0
1–10	2.35	97.65	0	0
11–20	4.36	0	95.64	0
20+	9.21	0	0	90.79

Table 2 One-year BMI summary transition matrix

(pct.)	BMI transition matrix				
	Underweight	Normal weight	Overweight	Obese class I	Obese class II–III
Underweight	94.55	5.55	0	0	0
Normal weight	0	97.60	2.40	0	0
Overweight	0	0	98.25	1.75	0
Obese class I	0	0	0	98.23	1.77
Obese class II–III	0	0	0	0	100

the exposure categories using natural cubic spline interpolation for fixed sex, age, and cause-of-death. A similar approach is used in the appendix of [28]. An example is given in the right panel of Fig. 9.

Appendix D: Transition matrices

Figure 4 shows a secular trend in the prevalence distributions for smoking, obesity and physical activity, whereas alcohol consumption remains roughly constant over the period. We opt for a migration model that captures the main effect, namely the net migration flow, using just the data at hand. We construct the number of net migration events (NM) by balancing the equation of population change

$$Y(x+1, t+1) = Y(x, t) - D(x, t) + NM(x, t),$$

and by imposing that transitions occur only between neighbouring categories of BMI and physical activity (in one year), while transitions for smoking are described in terms of the probability of cessation. In the above, $Y(x, t)$ is the number of individuals alive at age x and time t while $D(x, t)$ is the number of deaths. We estimate transitions for each risk factor independently of one another and for both sexes and all ages and calendar years combined. The resulting transition matrices for the surviving population, i.e. estimates of m_{ij} in Sect. 6.3, are shown below. To aid in understanding the output of the model, Fig. 10 shows the mean trajectory of the continuation of the compositions from Fig. 4.

Table 3 One-year physical activity summary transition matrix in terms of MET-minute-categories given by the cut points 0, 600, 1200, 1800, 2400, 3000, 3600, and 4200

(pct.)	Physical activity transition matrix							
	Cat. 1	Cat. 2	Cat. 3	Cat. 4	Cat. 5	Cat. 6	Cat. 7	Cat. 8
Cat. 1	95.37	4.63	0	0	0	0	0	0
Cat. 2	0	100	0	0	0	0	0	0
Cat. 3	0	21.83	78.17	0	0	0	0	0
Cat. 4	0	0	43.43	56.57	0	0	0	0
Cat. 5	0	0	0	19.73	80.27	0	0	0
Cat. 6	0	0	0	0	14.52	85.48	0	0
Cat. 7	0	0	0	0	0	3.69	95.73	0.58
Cat. 8	0	0	0	0	0	0	0	100

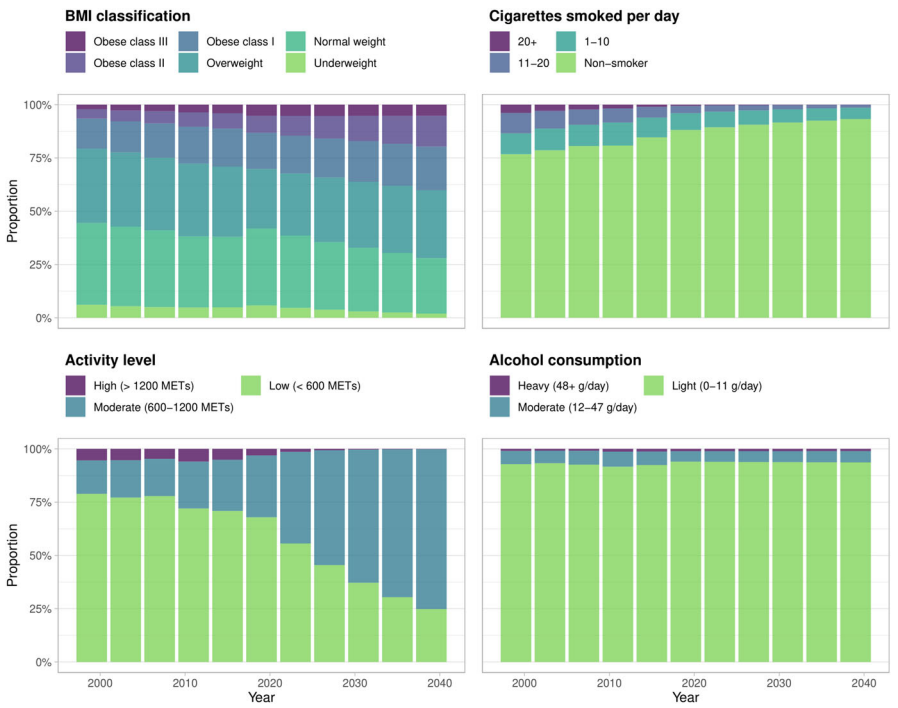

Fig. 10 U.S. risk proportions from Fig. 4 and their forecasted mean trajectory using (32) with migration rates from Tables 1, 2, 3.

References

1. Alai DH, Arnold S, Sherris M (2015) Modelling cause-of-death mortality and the impact of cause-elimination. Ann Actuar Sci 9(1):167–186. https://doi.org/10.1017/S1748499514000027X
2. Andersen PK, Borgan O, Gill RD, Keiding N (1993) Statistical models based on counting processes. Springer, New York

3. Booth H, Tickle L (2008) Mortality modelling and forecasting: a review of methods. Ann Actuar Sci 3(1–2):3–43. https://doi.org/10.1017/S1748499500000440
4. Cairns A, Blake D, Dowd K (2006) A two-factor model for stochastic mortality with parameter uncertainty: theory and calibration. J Risk Insur 73(4):687–718. https://doi.org/10.1111/j.1539-6975.2006.00195.x
5. CDC WONDER: Centers for Disease Control and Prevention, National Center for Health Statistics. Underlying Cause of Death 1999–2018 on CDC WONDER Online Database, released in 2020. Data are from the Multiple Cause of Death Files, 1999–2018, as compiled from data provided by the 57 vital statistics jurisdictions through the Vital Statistics Cooperative Program. http://wonder.cdc.gov/ucd-icd10.html
6. Dawid AP (2002) Influence diagrams for causal modelling and inference. Int Stat Rev 70(2):161–189. https://doi.org/10.1111/j.1751-5823.2002.tb00354.x
7. Didelez V (2000) Graphical models for event history analysis based on local independence. Logos, Berlin
8. Dimitrova D, Haberman S, Kaishev V (2013) Dependent competing risks: cause elimination and its impact on survival. Insur Math Econ 53(2). https://doi.org/10.1016/j.insmatheco.2013.07.008
9. Eichler M, Didelez V (2007) Causal reasoning in graphical time series models. In: Proceedings of the 23rd conference on uncertainty in artificial intelligence (UAI). AUAI Press, pp 109–116
10. Eichler M, Didelez V (2010) On Granger causality and the effect of interventions in time series. Lifetime Data Anal 16(1):3–32. https://doi.org/10.1007/s10985-009-9143-3
11. Foreman K et al (2018) Forecasting life expectancy, years of life lost, and all-cause and cause-specific mortality for 250 causes of death: reference and alternative scenarios for 2016–40 for 195 countries and territories. Lancet 392(10159):2052–2090. https://doi.org/10.1016/S0140-6736(18)31694-5
12. Girosi F, King G (2008) Demographic forecasting. Princeton University Press, Princeton
13. Granger CWJ (1969) Investigating causal relations by econometric models and cross-spectral methods. Econometrica 37:424–438. https://doi.org/10.2307/1912791
14. Hernán MA, Robins JM (2020) Causal inference: what if. Chapman & Hall/CRC, Boca Raton
15. Imbens GW, Rubin DB (2015) Causal inference in statistics, social, and biomedical sciences. Cambridge University Press, Cambridge
16. IPUMS, Blewett LA, Drew JAR, King ML, Williams KCW (2019) IPUMS health surveys: national health interview survey, version 6.4 [dataset]. IPUMS, Minneapolis. https://www.nhis.ipums.org
17. Janssen F, Kunst A (2007) The choice among past trends as a basis for the prediction of future trends in old-age mortality. Popul Stud 61(3):315–326. https://doi.org/10.1080/00324720701571632
18. Janssen F, Wissen L, Kunst A (2013) Including the smoking epidemic in internationally coherent mortality projections. Demography 50(4):1341–1362. https://doi.org/10.1007/s13524-012-0185-x
19. Jarner SF, Jallbjørn S (2022) The SAINT model: a decade later. ASTIN Bull J IAA 52(2):483–517. https://doi.org/10.1017/asb.2021.37
20. Kaishev V, Dimitrova D, Haberman S (2007) Modelling the joint distribution of competing risks survival times using copula functions. Insur Math Econ 41:339–361. https://doi.org/10.1016/j.insmatheco.2006.11.006
21. Karn MN (1931) An inquiry into various death-rates and the comparative influence of certain diseases on the duration of life. Ann Eugen 4(3–4):279–302. https://doi.org/10.1111/j.1469-1809.1931.tb02080.x
22. Keyfitz N (1977) What difference would it make if cancer were eradicated? An examination of the Taeuber paradox. Demography 14(4):411–418. https://doi.org/10.2307/2060587
23. King G, Soneji S (2011) The future of death in America. Demogr Res 25:1–38. https://doi.org/10.4054/DemRes.2011.25.1
24. Lee RD, Carter LR (1992) Modeling and forecasting U.S. mortality. J Am Stat Assoc 87(419):659–675. https://doi.org/10.2307/2290201
25. Li H (2019) Lu Y (2019) Modeling cause-of-death mortality using hierarchical Archimedean copula. Scand Actuar J 3:247–272. https://doi.org/10.1080/03461238.2018.1546224
26. Mackenbach JP, Kunst AE, Lautenbach H, Oei YB, Bijlsma F (1999) Gains in life expectancy after elimination of major causes of death: revised estimates taking into account the effect of competing causes. J Epidemiol Community Health 53(1):32–37. https://doi.org/10.1136/jech.53.1.32
27. Manton KG, Poss SS (1979) Effects of dependency among causes of death for cause elimination life table strategies. Demography 16(2):313–327. https://doi.org/10.2307/2061145

28. Murray CJL, Aravkin AY, Zheng P, Abbafati C, Abbas KM, Abbasi-Kangevari M et al (2020) Global burden of 87 risk factors in 204 countries and territories, 1990–2019: a systematic analysis for the Global Burden of Disease Study 2019. Lancet 396(10258):1223–1249. https://doi.org/10.1016/S0140-6736(20)30752-2
29. Palloni A, Beltrán-Sánchez H (2017) Discrete Barker Frailty and warped mortality dynamics at older ages. Demography 54(2):655–671. https://doi.org/10.1007/s13524-017-0548-4
30. Pearl J (2001) Direct and indirect effects. In: Proceedings of the 17th conference on uncertainty in artificial intelligence (UAI). Morgan Kaufmann Publishers Inc., San Francisco, pp 411–420
31. Pearl J (2009) Causality. Cambridge University Press, Cambridge
32. Peters J, Janzing D, Schölkopf B (2017) Elements of causal inference: foundations and learning algorithms. The MIT Press, Cambridge
33. Preston S, Stokes A, Mehta N, Cao B (2014) Projecting the effect of changes in smoking and obesity on future life expectancy in the United States. Demography 51(1):27–49. https://doi.org/10.1007/s13524-013-0246-9
34. Renshaw AE, Haberman S (2006) A cohort-based extension to the Lee-Carter model for mortality reduction factors. Insur Math Econ 38(3):556–570. https://doi.org/10.1016/j.insmatheco.2005.12.001
35. Robins JM (1998) Structural nested failure time models. Encycl Biostat 6:4372–4389. https://doi.org/10.1002/9781118445112.stat06059
36. Robins JM, Greenland S (1992) Identifiability and exchangeability for direct and indirect effects. Epidemiology 3:143–155. https://doi.org/10.1097/00001648-199203000-00013
37. Robins JM, Greenland S, Hu FC (1999) Estimation of the causal effect of a time-varying exposure on the marginal mean of a repeated binary outcome. J Am Stat Assoc 94(447):687–700. https://doi.org/10.1080/01621459.1999.10474168
38. Spirtes P, Glymour CN, Scheines R, Heckerman D (2000) Causation, prediction, and search. The MIT Press, Cambridge
39. Vaupel JW, Yashin AI (1985) Heterogeneity's ruses: some surprising effects of selection on population dynamics. Am Stat 39(3):176–185. https://doi.org/10.1080/00031305.1985.10479424
40. Vaupel JW, Manton KG, Stallard E (1979) The impact of heterogeneity in individual frailty on the dynamics of mortality. Demography 16(3):439–454. https://doi.org/10.2307/2061224
41. Wang H, Preston S (2009) Forecasting United States mortality using cohort smoking histories. Proc Natl Acad Sci USA 106:393–398. https://doi.org/10.1073/pnas.0811809106