DISCUSSION PAPER PI-0512

Stochastic Mortality Under Measure Changes

Enrico Biffis, Michel Denuit and Pierre Devolder

September 2009

ISSN 1367-580X

The Pensions Institute
Cass Business School
City University
106 Bunhill Row London
EC1Y 8TZ
UNITED KINGDOM

http://www.pensions-institute.org/

Stochastic Mortality Under Measure Changes

ENRICO BIFFIS[‡], MICHEL DENUIT[*†] & PIERRE DEVOLDER[†]

[‡] Imperial College Business School
Imperial College London
South Kensington Campus
London SW7 2AZ, United Kingdom

[*]Institut de Statistique
Université Catholique de Louvain
Voie du Roman Pays, 20
B-1348 Louvain-la-Neuve, Belgium

[†]Institut des Sciences Actuarielles
Université Catholique de Louvain
Grand-Rue, 54
B-1348 Louvain-la-Neuve, Belgium

First draft: March 2005.
This version: September 27, 2009.
Final version to appear in *Scandinavian Actuarial Journal*.

Abstract

We provide a self-contained analysis of a class of continuous-time stochastic mortality models that have gained popularity in the last few years. We describe some of their advantages and limitations, examining whether their features survive equivalent changes of measures. This is important when using the same model for both market-consistent valuation and risk management of life insurance liabilities. We provide a numerical example based on the calibration to the French annuity market of a risk-neutral version of the model proposed by Lee and Carter (1992).

Key words and phrases: stochastic mortality, Lee-Carter model, mortality risk premium, fair valuation, mortality-linked securities.

1 Introduction and motivation

In recent years stochastic mortality models have received growing attention for several reasons. First, the extent of mortality improvements has posed severe challenges to the insurance and pension industry, making clear the importance of demographic trends forecasting. Second, the introduction of market-consistent accounting and risk-based solvency requirements has called for the integration of mortality risk analysis into stochastic valuation models. Third, mortality-linked securities have attracted the interest of capital market investors, who in turn demand transparent tools to price demographic and financial risks in an integrated fashion.

Some of the models that try to address simultaneously these different objectives rely on a continuous time specification of the intensity of mortality or of survival/death probabilities. The approach provides a natural extension of static mortality models to a dynamic setting, and offers computational tractability in pricing and reserving applications. For example, the risk-neutral pricing machinery easily extends to mortality-contingent cashflows, and mortality risk can be immediately quantified as a spread on mortality-insensitive returns, exactly as in the case of credit or liquidity risk. In this paper, we review the most relevant assumptions and implications of this approach. In particular, we show how typical models rely, tacitly or explicitly, on the doubly stochastic assumption for the insureds' death times. This means that deaths are conditionally Poisson, given the evolution of the risk factors driving the instantaneous death probabilities. While this setup offers undubious computational advantages, it builds on restrictive assumptions that may be at odds with empirical evidence. Hence, care should be taken in using this class of models under the real-world measure. Even when the computational benefits outweigh the shortcomings, the joint use of the doubly stochastic setting under different probability measures may be problematic. We examine conditions under which this is possible and study whether they are restrictive for practical applications. All in all, the main drawback is that the observation of deaths in an insurance portfolio cannot be used to revise the conditional valuation of the cashflows to which the remaining survivors are entitled. That means some incomplete market pricing approaches can only be used at the price of computational tractability.

We examine quadratic hedging as an example.

Understanding the behavior of stochastic mortality models under equivalent probability measures is often important in practice. For example, an insurer may wish to use a reliable stochastic mortality model both under the physical measure, for solvency analysis or mortality projections, and under a pricing measure, for market consistent pricing of, or reserving for, insurance liabilities. On the other hand, insurers may adopt different approaches to fair valuation (state-price deflators, best estimates plus risk margins, or risk-neutral valuation), making it hard for regulators and investors to compare the results. A necessary condition for comparability is understanding the mortality risk premia supported by the modeling framework. We provide an analysis of such premia without postulating that the date of death of an individual attracts a zero risk premium (an assumption usually justified on the grounds of 'diversification'), and allow for a richer parametrization of mortality risk premia. There are three main reasons why this may be useful. First, the extent to which one can benefit from the law of large numbers is limited in the case of stochastic death rates, even in the case of conditionally i.i.d. death times (e.g., Schervish, 1995; Milevsky, Promislow and Young, 2006), and a risk premium for portfolio size may materialize more often than not.[1] Second, diversification and risk pooling are different concepts. While diversification can be achieved by dividing a *fixed* amount of capital across a large number of policies (and in this case the risk associated with the timing of death is indeed diversifiable; see Jarrow, Lando and Yu, 2005), risk pooling involves writing a larger and larger number of policies. Scaling of the exposure, however, does not yield diversification (e.g., Samuelson, 1963). Finally, even if the insurer can benefit from economies of scale, the insured is faced with her own risk of death only; this may be an aspect to consider when pricing some insurance contracts and guarantees. We provide several examples to illustrate these different aspects of mortality risk.

As a reference mortality model, throughout the paper we consider a continuous time version of the model proposed by Lee and Carter (1992) (LC henceforth), whose description of the

[1] This may be the case even with deterministic death rates: large portfolios may reduce to classes of very few policies once contracts are disaggregated by relevant risk characteristics; in secondary markets, portfolios that are very large in value may contain very few homogeneous contracts (e.g., Life Settlements portfolios).

secular change of mortality as a function of a time index has proved quite effective in forecasting mortality trends (Lee, 2000; Girosi and King, 2008). While popular among demographers and insurance practitioners, the use of the model is typically confined to realistic valuations, without a suitable translation under a pricing measure. Given the interest of regulators and practitioners in translating under a common framework the mortality assumptions used in different contexts, we introduce a class of equivalent probability measures under which the structure of the LC model, and the tractability of the doubly stochastic setup, are preserved.

The paper is organized as follows. Section 2 provides a constructive description of a stochastic mortality model rooted in the insurance literature (e.g., Norberg, 1987, 1989; Steffensen, 2000) and widely used in credit risk modeling (see Bielecki and Rutkowski, 2001, and references therein). The link between the two fields has been emphasized, for example, by Artzner and Delbaen (1995); Milevsky and Promislow (2001); Biffis (2005) (for related literature focusing on the parallel with term structure models, see Dahl, 2004; Bauer, Benth and Kiesel, 2009, or the review paper by Cairns, Blake and Dowd, 2006b). We examine the key assumptions leading to the doubly stochastic setup, and emphasize when modeling of the hazard rate process is equivalent to modeling of the intensity process. Section 3 introduces the equivalent changes of measure supported by the model and provides conditions preserving the most appealing features of the setup. In particular, we describe a parameterization of the Radon-Nikodym density ensuring that LC stochastic intensities remain of the LC type under equivalent measures. We then examine some limitations of the model by looking at an example of partial information and filtering of the intensity dynamics. Section 4 focuses on pricing applications (see Møller and Steffensen, 2007, for an overview) and on the structure of mortality risk premia. Section 5 offers a numerical example based on the calibration of the LC model to the French annuity market. We estimate the different mortality risk premia ensuring consistency of a risk-neutral LC model (calibrated to an annuity pricing basis) with an LC model for aggregate population data. Section 6 concludes, while an appendix provides additional proofs and details.

2 A continuous-time mortality model

We consider the time interval $[0, T^*]$ and fix a filtered[2] probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\in[0,T^*]}, \mathbb{P})$. We set $\mathcal{F} = \mathcal{F}_{T^*}$ and assume $\mathcal{F}_0$ to be trivial. In the following, t, s and T will denote times in $[0, T^*]$; equalities and inequalities between random variables are in the almost sure sense; superscripts usually refer to coordinates of vector-valued processes rather than powers. Consider m insureds aged $x_1, \dots, x_m$ at time 0, and introduce the vector of death indicator processes[3] $N_t = (N_t^1, \dots, N_t^m)'$, meaning that each component is defined by $N_t^i := 1_{\tau^i \le t}$, with τ^i the random time of death of the i-th insured.

For ease of exposition, we consider first the case of a single insured ($m = 1$). It is convenient to model the death time of the insured as a stopping time with respect to $\mathbb{F} = (\mathcal{F}_t)_{t\ge0}$. This means that $\mathbb{F}$ carries enough information to tell whether τ has occurred or not by each time t. If we wish to distinguish between the world of factors influencing the occurrence of τ and the occurrence of τ itself, it is natural to assume that $\mathbb{F}$ is given by $\mathbb{G} \vee \mathbb{H}$, where $\mathbb{G} = (\mathcal{G}_t)_{t\ge0}$ is a strict subfiltration of $\mathbb{F}$ providing information on relevant risk factors and $\mathbb{H} = (\mathcal{H}_t)_{t\ge0}$ makes $\mathbb{F}$ the smallest enlargement of $\mathbb{G}$ with respect to which τ is a stopping time.[4] We may think of $\mathbb{G}$ as carrying information based on medical/demographical data collected at population/industry level, and of $\mathbb{H}$ as capturing the actual occurrence of death in an insurance portfolio.

We characterize the conditional law of τ in several steps. By the Doob-Meyer decomposition (e.g., Protter, 2004, p. 111), there is a unique $\mathbb{F}$-predictable increasing process A, called the *compensator* of N, such that $N - A$ is a martingale. Since N is constant after jumping to one at death, its compensator must be constant as well, and thus $A_t = A_{t\wedge\tau}$ for all t. In view of the assumed information structure, we can always find a $\mathbb{G}$-predictable process Λ coinciding with A up to τ (Protter, 2004, p. 370), i.e., such that $A_t = A_{t\wedge\tau} = \Lambda_{t\wedge\tau}$ for all t. It turns out that Λ is continuous (Protter, 2004, Thm. 7, p. 106). If it is absolutely continuous with respect to the Lebesgue measure, we can write $\Lambda_t = \int_0^t \mu_s \mathrm{d}s$ for some nonnegative $\mathbb{G}$-predictable process μ. We call Λ the *martingale hazard process* and μ the *intensity* of τ. To model the time evolution

[2] All filtrations are assumed to satisfy the usual conditions, i.e. right-continuity and $\mathbb{P}$-completeness.
[3] In the following, we use the notation $y \cdot z = y'z = \sum_{i=1}^m y_i z_i$, for $y, z \in \mathbb{R}^m$.
[4] That is, $\mathcal{F}_t = \cap_{s>t}\mathcal{G}_s \vee \sigma(\tau \wedge s)$ for all t. We use the notation $t \wedge s := \min(t, s)$ throughout the paper.

of mortality, it is natural to introduce the process[5] $F_t := \mathbb{P}(\tau \leq t|\mathcal{G}_t)$, with $F_0 = 0$. Since τ is not a $\mathbb{G}$-stopping time, we must have $F_t < 1$ for all t, and can then apply Bayes rule to obtain (see Proposition A.1 in the appendix)

$$\mathbb{P}(\tau > T|\mathcal{F}_t) = 1_{\tau>t}\frac{\mathbb{P}(\tau > T|\mathcal{G}_t)}{\mathbb{P}(\tau > t|\mathcal{G}_t)}. \tag{2.1}$$

We can then set $F_t = 1 - \exp(-\Gamma_t)$ for some $\mathbb{G}$-adapted nonnegative process Γ satisfying $\Gamma_0 = 0$. We call Γ the *hazard process of* τ and use it to rewrite expression (2.1) as

$$\mathbb{P}(\tau > T|\mathcal{F}_t) = 1_{\tau>t}E\left[\exp(\Gamma_t - \Gamma_T)|\mathcal{G}_t\right]. \tag{2.2}$$

A natural question to ask is, what is the relationship between the hazard process Γ and the martingale hazard process Λ? The following proposition gives conditions under which the two notions coincide:

Proposition 2.1. *Assume that the stochastic basis* $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ *satisfies the structural conditions described above, and that the following conditions hold:*

- **(A1)** *F is continuous.*
- **(A2)** *For all t, the σ-fields $\mathcal{H}_t$ and $\mathcal{G}_{T^*}$ are conditionally independent, given $\mathcal{G}_t$.*

Then, the martingale hazard process, Λ, is unique and coincides with the hazard process, Γ. The compensator of τ is thus given by $A_t = \Lambda_{\tau\wedge t} = \Gamma_{\tau\wedge t}$.

Proof. In the appendix. □

Given Λ, we can construct τ by introducing a unit-mean exponential random variable, Θ, independent of $\mathcal{G}_{T^*}$, and setting

$$\tau = \inf\left\{t : \Lambda_t > \Theta\right\}. \tag{2.3}$$

See Blanchet-Scalliet and Jeanblanc (2004) for a proof. We use the convention $\inf\varnothing = \infty$, in which case $\tau > T^*$. When $\Lambda_t = \int_0^t \mu_s \mathrm{d}s$, construction (2.3) shows that τ coincides with the

[5] We consider its right-continuous-with-left-limits modification.

first jump time of a doubly stochastic (equivalently, Cox or conditionally Poisson) process with intensity μ driven by $\mathbb{G}$ (Brémaud, 1981; Duffie, 2001).

Mortality is often modeled by specifiying the hazard process Γ and working with the convenient formula (2.2). (The situation when $\mathbb{G}$ and $\mathbb{H}$ are independent captures the case of deterministic hazard functions.) While we know that disregarding the notion of martingale hazard process is immaterial under the assumptions of Proposition 2.1, this perspective may not be inconsequential in more general situations, in particular when considering changes of measure, as we will argue in the following sections.

The previous analysis can be immediately extended to the case of $m > 1$ insureds. We disregard the case of simultaneous deaths, i.e., we assume that $\tau^i \neq \tau^j$ almost surely on $[0, T^*]$ for $i \neq j$. We set $\mathbb{H} = \vee_{i=1}^m \mathbb{H}^i$, where each $\mathbb{H}^i$ is the augmented filtration generated the death indicator process N_t^i. For each fixed i, we can use the result described above by assuming that conditions (A1)-(A2) hold with respect to $\mathbb{F}^i := \mathbb{G} \vee \mathbb{H}^i$ and $F_t^i := \mathbb{P}(\tau^i \leq t | \mathcal{G}_t)$. To ensure tractability when the entire information $\mathbb{F} := \mathbb{G} \vee \mathbb{H}$ is available, however, we need to impose the additional condition

(A3) for all t, the σ-fields $\mathcal{H}_t^1, \ldots, \mathcal{H}_t^m$ are conditionally independent, given $\mathcal{G}_{T^*}$.

When the above holds, conditional survival probabilities can be conveniently computed as

$$\mathbb{P}(\tau^i > T | \mathcal{F}_t) = 1_{\tau^i > t} E\left[\exp\left(\Lambda_t^i - \Lambda_T^i\right) | \mathcal{G}_t\right] = \mathbb{P}(\tau^i > T | \mathcal{F}_t^i).$$

Assumption (A3) means that once the information carried by $\mathbb{G}$ is taken into account, the residual randomness in death times is independent across individuals. A vector of random death times $(\tau^1, \ldots, \tau^m)$ can be constructed by generating m independent draws from a unit-mean exponential distribution, and using construction (2.3) for each component τ^i. If each death time admits the intensity μ^i, then each τ^i can be shown to coincide with the first jump in the i-th component of an m-variate doubly stochastic process.

Remark 2.2. In the insurance literature it is common to model the deaths in an insurance portfolio as the jump times of a point process $N^* \wedge m$, where m represents the portfolio size (e.g., Dahl and Møller, 2006). $N^* \wedge m$ is typically assumed to have intensity $(m - (N^*_{t-} \wedge m))\mu_t$, for some $\mathbb{G}$-predictable process μ. Consider now an m-variate doubly stochastic process $\widehat{N} = (\widehat{N}^1, \dots, \widehat{N}^m)$, whose components have all intensity μ. The process $\sum_{i=1}^m \widehat{N}^i$ is then a univariate doubly stochastic process with intensity $m\mu$. If we look at the first jump times of the components of $\widehat{N}$, we see equivalence of our setup with using $N^* \wedge m$, at least in the case of conditionally i.i.d. death times admitting an intensity μ. For the more general case of death times with intensities $\mu^1, \dots, \mu^m$, we see that the process $\sum_{i=1}^m \widehat{N}^i$ is doubly stochastic with intensity $\sum_{i=1}^m \mu^i$, and the instantaneous conditional jump probability of each $\widehat{N}^j$ is $\mu^j (\sum_{i=1}^m \mu^i)^{-1}$. If we focus again on the first jump times of the components of $\widehat{N}$, the intensity of $\sum_{i=1}^m \widehat{N}^i$ is now $\sum_{i=1}^m \mu^i$ only before occurrence of the first death. As soon as the first death is recorded, say $\widehat{N}^j = 1$, the intensity of $\sum_i \widehat{N}^i$ jumps to $\sum_{i \neq j} \mu^i$. The setup outlined in the previous section easily allows for heterogeneous lives, while avoiding the cumbersome procedure of updating the intensity of the remaining survivors every time a death is recorded. Moreover, the setup can be used beyond the conditionally Poisson setting, for example by working directly with the hazard process as in (2.2).

Example 2.3 (Generalized LC model). For an example of stochastic intensity, we consider the LC model (Lee and Carter, 1992). Let us introduce the set of ordered dates $\mathcal{T} = \{t_0, \dots, t_h\} \subset [0, T^*]$ and m representative individuals with ages in $\mathcal{I} = \{x_1, \dots, x_m\}$ at the reference time $t_0 = 0$. We further let $m_x(t)$ denote the central death rate relative to age $x \in \mathcal{I}$ and date $t \in \mathcal{T}$. This means that, for each i and j, $m_{x_i}(t_j)$ gives the average number of deaths between ages x_i and x_{i+1} to the exposed to risk between t_j and t_{j+1}. Under the assumption of a piecewise constant intensity between each age pair $(x_i + t, x_i + t + 1)$, we can approximate μ^i_t with $m_{x_i+t}(t)$. The LC approach is based on a relational model of the form

$$\ln \widehat{m}_x(t) = \alpha_x + \beta_x \kappa_t + \epsilon_x(t) \tag{2.4}$$

where $\widehat{m}_x(t)$ is the unconstrained maximum likelihood estimator of $m_x(t)$, the $\epsilon_x(t)$'s are homoskedastic centered error terms and the parameters $\{\alpha_x\}, \{\beta_x\}$ are subject to the identification constraints

$$\sum_{t\in\mathcal{T}}\kappa_t = 0 \text{ and } \sum_{x\in\mathcal{I}}\beta_x = 1. \tag{2.5}$$

The model (2.4) is fitted to a matrix of age-specific observed death rates using singular value decomposition (SVD), and the resulting estimate of the time-varying parameter κ is then forecasted as a stochastic time series using standard Box-Jenkins methods. For example, Lee and Carter (1992) model κ as a random walk with drift. The interpretation of the parameters is as follows: the $\exp(\alpha_x)$'s describe the general shape across age of the mortality schedule, while the mortality intensities change according to an index κ modulated by the age responses (β_x). In light of what was observed in the previous section, we may view the LC model as a discretized version of the stochastic intensity

$$\mu_t^i = \exp(\alpha(x_i + t) + \beta(x_i + t)\cdot\kappa_t), \tag{2.6}$$

where α, β are continuous functions and κ is an $\mathbb{R}^d$-valued $\mathbb{G}$-predictable process. (Care must be taken in reconciling (2.4) with (2.6): α and β depend on time because each μ_t^i represents the intensity of an individual aged $x_i + t$ at time t. Thus, the parameters α_x's and β_x's in (2.4) correspond to the values of the functions $\alpha(\cdot), \beta(\cdot)$ computed at each age in $\mathcal{I}$.) We refer to the intensity (2.6) as to a *generalized LC intensity*. If κ is a Brownian motion with drift (i.e., $\mathrm{d}\kappa_t = \delta \mathrm{d}t + \sigma \mathrm{d}W_t$ for some real coefficients δ and σ), after discretization along the time dimension we are back to the model originally proposed by Lee and Carter (1992). If k is a one-dimensional Ornstein-Uhlenbeck process, we recover the model used for example by Milevsky and Promislow (2001). We note that some authors (e.g., Brouhns, Denuit and Vermunt, 2002) have paired the exponential affine form (2.4) with a Poisson regression model for death counts: since the estimates of κ are treated as realizations of a random walk with drift, and (2.6) is seen as the intensity of a Poisson inhomogeneous process, we see agreement with the doubly stochastic setup. ◦

3 Stochastic Mortality under Measure Changes

For simplicity from now on we assume that $\mathbb{G}$ is the augmented filtration generated by a d-dimensional Brownian motion W. (In general, $\mathbb{G}$ could be generated by discontinuous risk factors, for example evolving as jump diffusions. The main features outlined in Section 2 would be unaltered as long as jump times are not death times; see Example 4.2.) We are now interested in changes from the reference measure $\mathbb{P}$ (e.g., physical/objective probability) to an arbitrary equivalent measure $\widetilde{\mathbb{P}}$ defined on $(\Omega, \mathcal{F})$. Two probability measures are said to be equivalent when they assign null probability to the same events. Depending on the context, additional constraints may be imposed: see Section 4 for no-arbitrage restrictions. We introduce the Radon-Nikodym density process ξ by setting

$$\xi_t := \frac{\mathrm{d}\widetilde{\mathbb{P}}}{\mathrm{d}\mathbb{P}}\Big|_{\mathcal{F}_t} = E\left[\xi_{T^*}|\mathcal{F}_t\right], \tag{3.1}$$

where ξ_{T^*} is a strictly positive $\mathcal{F}$-measurable random variable such that $E[\xi_{T^*}] = 1$. In the following we write $E[\cdot]$ and $\widetilde{E}[\cdot]$ for expectations under $\mathbb{P}$ and $\widetilde{\mathbb{P}}$. A martingale representation theorem given in Appendix A.2 ensures that ξ can be expressed as[6]

$$\xi_t = 1 - \int_0^t \xi_{s-}\eta_s \cdot \mathrm{d}W_s + \sum_{i=1}^m \int_0^t \xi_{s-}\phi_s^i \mathrm{d}M_s^i, \tag{3.2}$$

with η and $(\phi^1, \ldots, \phi^m)$ $\mathbb{F}$-predictable processes satisfying suitable integrability conditions, and with each ϕ^i valued in $(-1, \infty)$ to ensure the strict positivity of ξ. The processes $M^1, \ldots, M^m$ are purely discontinuous $\mathbb{F}$-martingales defined by $M^i := N^i - \Lambda^i_{\cdot\wedge\tau^i}$. Appendix A.2 shows that the explicit factorization $\xi = \xi'\xi''$ holds, with ξ' and ξ'' given by

$$\xi_t' = \exp\left(-\int_0^t \eta_s \cdot \mathrm{d}W_s - \frac{1}{2}\int_0^t ||\eta_s||^2 \mathrm{d}s\right) \tag{3.3}$$

$$\xi_t'' = \prod_{i=1}^m \xi_t''^{(i)} = \prod_{i=1}^m \left(1 + \phi_{\tau^i}^i 1_{\tau^i \le t}\right) \exp\left(-\int_0^{t\wedge\tau^i} \phi_s^i \mathrm{d}\Lambda_s^i\right). \tag{3.4}$$

[6] We use the convention that $\int_0^t$ stands for integration over $(0, t]$.

Define the processes $\widetilde{W}_t := W_t + \int_0^t \eta_s \mathrm{d}s$, $\widetilde{M}_t^i := M_t^i - \int_0^{t\wedge\tau^i} \phi_s^i \mathrm{d}\Lambda_s^i$ and $\widetilde{\Lambda}_t^i = \int_0^t (1+\phi_s^i)\mathrm{d}\Lambda_s^i$. An application of the Girsanov-Meyer theorem (e.g. Protter, 2004, p. 132) yields the following result:

Proposition 3.1. *Consider the probability measure* $\widetilde{\mathbb{P}}$ *on* $(\Omega, \mathcal{F})$ *equivalent to* $\mathbb{P}$ *and with density* $\xi = \xi'\xi''$ *given by* (3.3)-(3.4)*. Then*

(i) $\widetilde{W}$ *is an* $\mathbb{F}$*-Brownian motion under* $\widetilde{\mathbb{P}}$.

(ii) For each i, $\widetilde{M}^i$ *is an* $\mathbb{F}$*-martingale under* $\widetilde{\mathbb{P}}$*, orthogonal to* $\widetilde{W}$ *and to* M^j *for each* $j \neq i$*. Moreover, the martingale hazard process of* τ^i *is given by* $\widetilde{\Lambda}^i$.

An immediate consequence of Proposition 3.1 is that the martingale hazard process of each death time τ^i may not remain the same process when moving to the new measure $\widetilde{\mathbb{P}}$. Moreover, each $\widetilde{\Lambda}^i$ may not coincide with the hazard process anymore, as assumptions (A1)-(A2)-(A3) are in general not preserved under equivalent changes of measure.[7] The following proposition gives conditions ensuring that this is the case:

Proposition 3.2. *Consider the density* $\xi = \xi'\xi''$ *given by* (3.3)-(3.4) *and assume that the following condition is satisfied:*

(A4) η *and* $\phi^1, \ldots, \phi^m$ *are* $\mathbb{G}$*-predictable.*

Then assumptions (A1)-(A2)-(A3) *hold under* $\widetilde{\mathbb{P}}$.

Proof. See Appendix A.2. □

Condition (A4) ensures that ξ'_{T^*} is $\mathcal{G}_{T^*}$-measurable and each $\xi''^{(i)}_{T^*}$ is $\mathcal{G}_{\tau^i}$-measurable. (Since for each $\mathbb{F}^i$-predictable process ρ we can always find a $\mathbb{G}$-predictable process $\bar{\rho}$ coinciding with ρ up to time τ^i, we could state condition (A4) in terms of $\mathbb{F}^i$-predictable processes up to each death time.) If for all i we have $\widetilde{\Lambda}^i = \int_0^{\cdot} \widetilde{\mu}_s^i \mathrm{d}s$, then each death time τ^i coincides with the first jump of a conditionally Poisson process with intensity $\widetilde{\mu}^i = (1+\phi^i)\mu^i$ under $\widetilde{\mathbb{P}}$. To understand to what extent assumption (A4) may be restrictive for practical applications, we examine two

[7] With regard to (A1), we mean that $\widetilde{F}_t := \widetilde{\mathbb{P}}(\tau \leq t | \mathcal{G}_t)$ may not be continuous.

examples. In the first one, we consider an insurer faced with intensity dynamics that are only partially observable. In the second example, we look at LC intensities preserving their structure under equivalent changes of measure.

Example 3.3 (Incomplete information). We slightly modify the setup of Section 2 by considering $\mathbb{F} = \mathbb{G} \vee \mathbb{H} \vee \sigma(\psi)$, with ψ a random variable defined on $(\Omega, \mathcal{F})$ and independent of the Brownian motion W, which for simplicity is assumed to be one-dimensional. It can be verified that the properties described in Section 2 still hold in the present setting, although $\mathcal{F}_0$ is clearly nontrivial. We think of ψ as an unobservable source of randomness affecting the intensities $\mu^1, \ldots, \mu^m$. As a simple example, assume that the random times of death $\tau^1, \ldots, \tau^m$ have common intensity μ such that $Y := \log \mu$ evolves according to

$$\mathrm{d}Y_t = (a(t) + b(t)Y_t + c(t)\psi)\,\mathrm{d}t + \sigma(t, Y_t)\mathrm{d}W_t, \tag{3.5}$$

where the functions a, b, c, σ satisfy the conditions given in Appendix A.3. We assume that the insurer observes Y, recovers σ from the quadratic variation of Y, and draws inferences about the drift in a Bayesian fashion, backing out from observations of Y the true value of ψ. For tractability, we let the insurer have a Gaussian prior on ψ (i.e., under $\mathbb{P}$ the conditional distribution of ψ, given $\mathcal{F}_0$, is Normal with mean zero and standard deviation σ^ψ) and update her conditional estimate of ψ by computing $\Psi_t := E\left[\psi | \mathcal{F}_t^Y\right]$, where $\mathbb{F}^Y$ is the filtration generated by Y. We define the innovation process $\widetilde{W}$ induced by the updating procedure in the classical filtering sense (e.g., Lipster and Shiryaev, 2001),

$$\mathrm{d}\widetilde{W}_t = \frac{\mathrm{d}Y_t - (a(t) + b(t)Y_t + c(t)\Psi_t)}{\sigma(t, Y_t)}, \tag{3.6}$$

so that Y is perceived to evolve as

$$\mathrm{d}Y_t = (a(t) + b(t)Y_t + c(t)\Psi_t)\,\mathrm{d}t + \sigma(t, Y_t)\mathrm{d}\widetilde{W}_t. \tag{3.7}$$

The above is 'observationally equivalent' to (3.5), since $\mathbb{F}^Y = \sigma(Y_0) \vee \mathbb{F}^{\widetilde{W}}$.

Assume now that τ is doubly stochastic with intensity μ under $\mathbb{P}$, driven by $\mathbb{G}^{\psi} := \mathbb{G} \vee \sigma(\psi)$ (our reference subfiltration). Under the filtering procedure described above, we can think of the insurer as being endowed with the filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}^Y, \widetilde{\mathbb{P}})$. The process $\widetilde{W} = W + \int_0^{\cdot} \eta_s \mathrm{d}W_s$ is a Brownian motion under a subjective probability measure $\widetilde{\mathbb{P}}$ reflecting the insurer's prior beliefs on ψ and belonging to the class introduced in Section 3, with

$$\eta_t = \frac{c(t)\,(\psi - \Psi_t)}{\sigma(t, Y_t)}, \quad \phi_t^i = 0.$$

Since η is $\mathbb{G}^{\psi}$ predictable, by Proposition 3.2 we see that τ is still doubly stochastic with intensity μ under $\widetilde{\mathbb{P}}$, although the insurer believes the intensity to have log-dynamics (3.7). Things change, however, as soon as we extend the information available to the insurer. For example, it would be natural for the insurer to back out the true value of ψ from observation of the death times in the insurance portfolio, i.e., by using $\mathbb{F}^Y \vee \mathbb{H}$ rather than just $\mathbb{F}^Y$. If this were the case, however, the perceived dynamics of Y would have discontinuities induced by death occurrences. The subjective probability measure $\widetilde{\mathbb{P}}$ would then be characterized by a density dependent on the death times, thus violating assumtpion (A4) (e.g., Snyder, 1972; Segall, Davis and Kailath, 1975). $\circ$

Example 3.4 (LC measure changes). Assume that under $\mathbb{P}$ the death times $\tau^1, \ldots, \tau^m$ have stochastic intensities $\mu^1, \ldots, \mu^m$ of the generalized LC type (2.6). Define a class of measure changes by setting

$$\phi_t^i = \exp\left(a^i(x_i + t) + b^i(x_i + t) \cdot \kappa_t\right) - 1 \tag{3.8}$$

in (3.4), for some continuous functions a^i, b^i such that the density (3.1) is well defined. Assume further that η in (3.3) is $\mathbb{G}$-predictable. Then, under $\widetilde{\mathbb{P}}$ each stopping time τ^i has stochastic intensity $\widetilde{\mu}^i$ still of the LC type, in the sense that

$$\widetilde{\mu}_t^i = \mu_t^i(1 + \phi_t^i) = \exp\left(\widetilde{\alpha}^i(x_i + t) + \widetilde{\beta}^i(x_i + t) \cdot \kappa_t\right), \tag{3.9}$$

with $\widetilde{\alpha}^i := \alpha + a^i$, $\widetilde{\beta}^i := \beta + b^i$. This class of measure changes is quite flexible: $\widetilde{\mathbb{P}}$ can only affect the dynamics of the time-trend κ (each ϕ^i is the null process), only the age responses to κ (η is the null process), or both of them. $\circ$

To appreciate the extra flexibility delivered by the term ξ'' in (3.1), it is instructive to look at the dynamics of the new intensity under an equivalent probability measure. Consider a single death time τ and for simplicity set $\phi^i = \phi$ for all i in (3.4). Assume that the $\mathbb{P}$-dynamics of μ is

$$\mathrm{d}\mu_t = \mu_t(\delta^\mu_t \mathrm{d}t + \sigma^\mu_t \cdot \mathrm{d}W_t),$$

where the drift δ^μ and the volatility σ^μ can be recovered by applying Itô's formula to (2.6). The dynamics of μ under $\widetilde{\mathbb{P}}$ is then

$$\mathrm{d}\mu_t = \mu_t(\widetilde{\delta}^\mu \mathrm{d}t + \sigma^\mu_t \cdot \mathrm{d}\widetilde{W}_t),$$

with $\widetilde{\delta}^\mu = \delta^\mu - \eta \cdot \sigma^\mu$. When the adjustment η is positive, $\widetilde{\mathbb{P}}$ gives greater weight to mortality improvements. This is only part of the story, since we know that μ needs not be the intensity process of τ under $\widetilde{\mathbb{P}}$, unless ϕ is the null process. Setting $\varphi := 1 + \phi > 0$, under (A4) we can write the dynamics of φ under $\widetilde{\mathbb{P}}$ as

$$\mathrm{d}\varphi_t = \varphi_t(\widetilde{\delta}^\varphi_t \mathrm{d}t + \sigma^\varphi_t \cdot \mathrm{d}\widetilde{W}_t),$$

with obvious meaning of the notation. Integration by parts then yields:

$$\begin{aligned} \mathrm{d}\widetilde{\mu}_t &= \varphi_{t-}\mathrm{d}\mu_t + \mu_{t-}\mathrm{d}\varphi_t + \mathrm{d}[\varphi, \mu]_t \qquad (3.10) \\ &= \widetilde{\mu}_t \left((\widetilde{\delta}^\mu_t + \widetilde{\delta}^\varphi_t + \sigma^\mu_t \cdot \sigma^\varphi_t)\mathrm{d}t + (\sigma^\mu_t + \sigma^\varphi_t) \cdot \mathrm{d}\widetilde{W}_t \right). \end{aligned}$$

Expression (3.11) shows that the drift adjustment $\widetilde{\delta}^\mu - \delta^\mu = -\eta_t \cdot \sigma^\mu_t$ is replaced by the more structured spread

$$\widetilde{\delta}^{\widetilde{\mu}}_t - \delta^\mu_t = -\eta_t \cdot \sigma^\mu_t + \widetilde{\delta}^\varphi_t + \sigma^\mu_t \cdot \sigma^\varphi_t,$$

which allows for a richer parameterization of the new intensity drift. Moreover, the volatility of the new intensity is in general different from that of the process μ under $\widetilde{\mathbb{P}}$. In particular, there are non-trivial specifications of $\widetilde{\delta}^{\varphi}$ and σ^{φ} that allow some Brownian sources of risk to disappear from the dynamics of the intensity when moving from $\mathbb{P}$ to $\widetilde{\mathbb{P}}$. (An extreme example is given by a death time with intensity that is deterministic under one measure and stochastic under the other.) This is useful when $\mathbb{G}$ is generated by both financial and mortality risk sources and one wishes to calibrate separately the corresponding risk premia to different security prices. Expression (3.11) shows that there are cases when we can assume independence of financial and demographic risk factors under $\widetilde{\mathbb{P}}$, without requiring it to hold under $\mathbb{P}$.

4 Pricing applications

In financial markets the absence of arbitrage is essentially equivalent to the existence of an equivalent martingale measure under which the discounted gain from holding a security is a martingale. The setup described in Section 2 allows one to easily extend a financial market to include securities providing mortality-contingent payouts (e.g., Artzner and Delbaen, 1995; Duffie, Schroder and Skiadas, 1996; Lando, 1998; Jamshidian, 2004). Take T^* to be a final trading date, and consider a financial market with security prices adapted to $\mathbb{G}$. In particular, assume that there is a money market account with market value $B_t = \exp(\int_0^t r_s \mathrm{d}s)$, where the $\mathbb{G}$-predictable process r represents the risk-free rate. In addition to financial securities, consider insurance contracts providing payoffs contingent on death. In this context, it is useful to rewrite assumption (A2) in the equivalent form (e.g, Brémaud and Jacod, 1978)

(A2∗) every $\mathbb{G}$-martingale is an $\mathbb{F}$-martingale.

The above makes clear that if (A2) still holds when moving to a pricing measure, then the dynamics of financial security prices are unaffected by the introduction of information on the death times $\tau^1, \ldots, \tau^m$. (When the financial market is incomplete, this may actually be more than what we need: under a pricing measure we could simply require (A2∗) to hold for the discounted gains from holding the traded securities.)

For simplicity, consider the case of a single death time. Let V denote the market value of an insurance contract issued to an individual aged x at time 0 and with random residual lifetime τ admitting an intensity μ under $\mathbb{P}$. Denoting by C the cumulated dividends paid by the security, we refer to V as to an ex-dividend price, i.e., excluding the payment of $\Delta C_t = C_t - C_{t-}$ at each time t. We assume that C is of bounded variation and follows

$$\mathrm{d}C_t = D_{t-}\mathrm{d}N_t + (1 - N_t)\mathrm{d}S_t, \tag{4.1}$$

where D and S are a $\mathbb{G}$-adapted processes representing the death benefit and the cumulated survival benefits respectively. In the absence of arbitrage, there exists a probability measure $\widetilde{\mathbb{P}}$ equivalent to $\mathbb{P}$ under which the discounted gain from holding V, $(B^{-1}V + \int_0^{\cdot} B_s^{-1}\mathrm{d}C_s)$, is an $\mathbb{F}$-martingale. We can thus write

$$V_t = B_t\ \widetilde{E}\left[\int_t^T B_s^{-1}\mathrm{d}C_s + B_T^{-1}1_{\tau>T}\overline{V}_T\middle|\mathcal{F}_t\right], \tag{4.2}$$

for all $T \geq t$, where we have assumed that the price of any security is zero at any given time if no dividends are paid thereafter. In the above, $\overline{V}$ denotes the pre-death price of the security, in the sense that $V_t = 1_{\tau>t}\overline{V}_t$. An application of Proposition A.1 in the appendix yields the pricing formula

$$\begin{aligned} V_t &= 1_{\tau>t}B_t\ e^{\widetilde{\Lambda}_t}\widetilde{E}\left[\int_t^T B_s^{-1}D_s\mathrm{d}\widetilde{\Lambda}_s + \int_t^T B_s^{-1}e^{-\widetilde{\Lambda}_s}\mathrm{d}S_s + B_T^{-1}e^{-\widetilde{\Lambda}_T}\ \overline{V}_T\middle|\mathcal{G}_t\right] \\ &= 1_{\tau>t}\widehat{B}_t\ \widetilde{E}\left[\int_t^T \widehat{B}_s^{-1}\ D_s\widetilde{\mu}_s\mathrm{d}s + \int_t^T \widehat{B}_s^{-1}\ \mathrm{d}S_s + \widehat{B}_T^{-1}\overline{V}_T\middle|\mathcal{G}_t\right], \end{aligned} \tag{4.3}$$

where $\widehat{B}_{\cdot} = \exp(\int_0^{\cdot}(r_s + \widetilde{\mu}_s)\mathrm{d}s)$ represents a 'mortality risk-adjusted money market account'. Expression (4.3) shows that the standard risk-neutral machinery passes over quite simply to the mortality-contingent setting, provided we consider fictitious securities paying a dividend $D_s\widetilde{\mu}_s\mathrm{d}s + \mathrm{d}S_s$ under a fictitious short rate $r + \widetilde{\mu}$. By (4.3) the discounted pre-death gain from holding the security, $L_t = \widehat{B}_t^{-1}\overline{V}_t + \int_0^t \widehat{B}_s^{-1}(D_s\widetilde{\mu}_s\mathrm{d}s + \mathrm{d}S_s)$, is a $\mathbb{G}$-martingale under $\widetilde{\mathbb{P}}$. As an example, assume that $D_t = d_t \overline{V}_t$ and $\mathrm{d}S_t = s_t \overline{V}_t \mathrm{d}t$ for suitable continuous processes d and s. We can then apply integration by parts to L and recover the dynamics

$$\mathrm{d}\overline{V}_t = \overline{V}_t (r_t + \widetilde{\mu}_t - d_t \widetilde{\mu}_t - s_t)\mathrm{d}t + \mathrm{d}Z_t, \tag{4.4}$$

where $Z_t = \int_0^t \widehat{B}_s \mathrm{d}L_s$ is a $\mathbb{G}$-martingale. The drift term in (4.4) shows how mortality uncertainty affects the dividend yield of pre-death prices under the pricing measure $\widetilde{\mathbb{P}}$. We elaborate on this in Example 4.1 below.

4.1 Market Price of Mortality risk

So far we have not required $\widetilde{\mathbb{P}}$ to be unique, which amounts to saying that both the financial and insurance markets may be incomplete. If the assets available for trade do not span the sources of randomness in the economy, there are infinitely many equivalent martingale measures making (4.2)-(4.3) hold. The latter determine an interval of no-arbitrage prices associated with different pairs (η, ϕ) accounting for investors' risk preferences. A way to narrow down the price range is to use market information to calibrate (4.3) and recover the risk-neutral dynamics of (η, ϕ). In the case of insurance contracts, useful information is provided by financial and reinsurance markets, or by suitably adjusted Embedded Value computations (see Biffis, 2005). The calibration approach is adopted, for example, by Cairns, Blake and Dowd (2006a), who restrict their attention to the case $\phi = 0$. An example based on the LC model is provided in Section 5. An alternative approach relies on making sensible assumptions on the preferences of market participants to pin down the equivalent martingale measure embedding their attitude toward risk. We elaborate on this approach in Example 4.2 below, where we discuss some of the examples given in Dahl and Møller (2006), who assume ϕ to be deterministic. We note that both approaches are consistent with the guidance on fair valuation provided by the International Accounting Standards Board (IASB, 2004).

The valuation framework described above allows us to identify different types of mortality risk affecting insurance contracts. First, there is *systematic* risk affecting the conditional probability of death of each individual in the insurance portfolio. This risk is channeled by randomness in the intensities of mortality. We call it longevity risk when we are concerned with systematic declines in the intensity. As an extreme example, consider the case of individuals with the same intensity μ: the evolution of μ affects the instantaneous death probability of all the insureds in the portfolio, at the same time and in the same direction, preventing the insurer from benefiting from pooling effects, as the impact of systematic risk is magnified by the size of the portfolio (see Schervish, 1995, for a version of the Strong Law of Large Numbers applying to this case). In our framework, the process η may include compensation for taking on mortality risk arising from systematic factors changing the intensity.

Second, there is the *unsystematic* risk of fluctuations in death occurrences, once the realizations of conditional death probabilities are given. If we use construction (2.3) in Section 2, we see that this risk is associated with the draws from a unit-mean exponential distribution. The extent to which fluctuations can be reduced depends on the size of the insurance portfolio at each point in time, i.e., on the trajectory of the process $\sum_{i=1}^{m} N^i$. In our framework, compensation for this risk is jointly captured by the processes $\phi^1(1-N^1), \ldots, \phi^m(1-N^m)$. From (3.2) we see that the pricing measure adjusts the intensity as soon as new deaths are recorded, to reflect an increase in unsystematic risk.

Finally, the third source of risk resides in the timing of each individual death, irrespective of portfolio size. This risk is relevant when the pricing framework captures the point of view of the policyholder, who is concerned with the death time of the insured life only, no matter how large the insurer's portfolio. As Example 4.1 shows, this risk may be present even when the product passes back to the policyholder any mortality profit generated by the insurance portfolio. Compensation for this risk is captured by each individual process $\phi^i(1-N^i)$ in isolation.

Summing up, while the first type of mortality risk premium materializes in a change of intensity dynamics, the latter two risk premia require a change in the intensity process.[8] Hence in the second case agents behave as if they were valuing cashflows contingent on some death

[8] A situation when a non-zero mortality risk premium must rely on a change of intensity is when $\mathbb{G}$ is independent of $\mathbb{H}$, and hence the hazard process is deterministic.

times $(\widetilde{\tau}^1, \ldots, \widetilde{\tau}^m)$ different from $(\tau^1, \ldots, \tau^m)$, and with intensities $(\widetilde{\mu}^1, \ldots, \widetilde{\mu}^m)$ under both $\mathbb{P}$ and $\widetilde{\mathbb{P}}$. This suggests a practical use of changes of measure as a way to capture basis risk in actuarial valuations. To illustrate, suppose that an insurer has limited experience data on an insurance portfolio, for example because it has entered a relatively new market. The mortality dynamics of an insured $\widetilde{\tau}$ could then be specified by adjusting the intensity of a representative τ in some reference population with more reliable data, for example by defining $\widetilde{\mu} = (1 + \phi)\mu$. In the context of mortality derivatives written on a public index, we could use a similar approach to capture the risk that the price of the derivative instrument does not move in line with the exposure the insurer wishes to hedge. The change in intensity corresponding to the use of $\widetilde{\tau}$ (instead of τ) under a pricing measure $\widetilde{\mathbb{P}}$ may be interpreted as a premium for basis risk.

Example 4.1 (Mortality-indexed insurance contract). Consider an insurance contract paying a continuous dividend yield, contingent on survival, equal to the risk-free rate plus the physical intensity of mortality of a reference population. We assume that the insured's time of death has the same law as the individuals in the reference population, and think of a variable annuity paying out the risk-free rate and passing back to the policyholder the demographic profit realized by the insurer on the pool of annuitants. Formally, set $D_t = 0$ and $\mathrm{d}S_t = (r_t + \mu_t)\overline{V}_t \mathrm{d}t$ in (4.1). Then, for any fixed $T > 0$, by (4.3) the initial price of the security is given by

$$V_0 = \widetilde{E}\left[\int_0^T \exp\left(-\int_0^s (r_u + \widetilde{\mu}_u)\mathrm{d}u\right)\overline{V}_s(r_s + \mu_s)\mathrm{d}s + \exp\left(-\int_0^T (r_s + \widetilde{\mu}_s)\mathrm{d}s\right)\overline{V}_T\right].$$

If $\phi = 0$ (i.e., $\widetilde{\mu}$ and μ are the same process) V_0 is equal to its pre-death price $\overline{V}_T$ at each fixed future date T. Since T is chosen arbitrarily, we have that V_0 is constant and the buyer of the contract has no way to express her aversion to mortality risk. In other words, the buyer must find the payment of the intensity μ enough to compensate her for her own risk of dying and ceasing to receive benefits, irrespective of how averse she is to changes in conditional death probabilities. Things are different if we allow for a change of intensity through a positive ϕ, because the pre-death price $\overline{V}$ has then higher drift. For example, if ϕ is deterministic, V_0 is decreasing in ϕ, capturing the fact that the insured is averse to her own risk of death and is willing to pay a lower price to enter the contract. ◦

Example 4.2 (Quadratic hedging). Throughout this example we assume that $\mathbb{G}$ is generated by a two-dimensional Brownian motion ($d = 2$). We consider a financial market with a money market account, B, and a risky security, P^1, with dynamics

$$\begin{cases} \mathrm{d}P_t^1 = P_t^1 \left(\delta_t^1 \mathrm{d}t + \sigma_t^1 \mathrm{d}W_t^1\right) \\ P_0^1 = p^1 > 0, \end{cases}$$

where r, δ^1 and σ^2 are strictly positive, $\mathbb{G}$-predictable, and uniformly bounded processes. Consider an insurer facing a time-T liability arising from a portfolio of equity-linked pure endowments issued to m individuals with common $\mathbb{G}$-predictable intensity of mortality μ. Denote the liability by $H_T \in \mathcal{F}_T$, and set $H_T = \sum_{i=1}^m 1_{\tau^i > T} f(P_T^1)$, with $f(\cdot)$ a bounded positive function. Suppose now that a reinsurer or investment bank offers a security providing payments linked to the number of survivors in the insurer's portfolio at time T. We may think of a reinsurance arrangement, or of a structured product embedding a longevity swap. Let the price process of this security, P^2, admit the representation

$$\begin{cases} \mathrm{d}P_t^2 = P_t^2 \left(\delta_t^2 \mathrm{d}t + \sigma_t^2 \mathrm{d}W_t^2 + \vartheta_t \cdot \mathrm{d}N_t\right) \\ P_0^2 = p^2 > 0, \end{cases}$$

where the coefficients δ^2, σ^2 and ϑ are $\mathbb{G}$-predictable and uniformly bounded, with $\delta^2, \sigma^2 > 0$ and ϑ valued in $(-1, \infty)^m$. Suppose that the insurer wishes to set up a dynamic hedging strategy by investing in (B, P^1, P^2), with the objective of minimizing the mean-square error of the mismatch between the terminal value of the strategy and the liabilities at maturity. The minimal cost to implement such strategy identifies a no-arbitrage price consistent with the quadratic hedging criterion (see Cerny and Kallsen, 2007, and references therein). In particular, we look for a self-financing strategy $\pi = (\pi^1, \pi^2)'$ (where π_t^j denotes the amount of wealth invested in asset P^j at time t) in the class of square-integrable, $\mathbb{F}$-predictable processes solving the optimization problem

$$\min_{\pi} E\left[\left(H_T - X_T^{x,\pi}\right)^2\right], \tag{4.5}$$

where $X^{x,\pi}$ is the wealth process associated with the investment strategy π starting from a given initial wealth level $x > 0$ (see Appendix A.4 for details). For simplicity, assume $\vartheta_t^i = \theta_t$ for all i, where θ is a process valued in $(-1, \infty)$. It can be shown that in our setup the solution to problem (4.5) leads to an initial price given by

$$V_0 = \widetilde{E}\left[\exp\left(-\int_0^T r_s \mathrm{d}s\right) H_T\right], \tag{4.6}$$

where $\widetilde{\mathbb{P}}$ is a measure equivalent to $\mathbb{P}$, with Radon-Nykodim density given by (3.3)-(3.4) (e.g., Lim, 2005). Denoting by $n_t = m - \sum_{i=1}^m N_{t-}^i$ the number of survivors in the portfolio at each point in time, we can express the pair (η, ϕ) as

$$\eta_t = \begin{bmatrix} \frac{\delta_t^1 - r_t}{\sigma_t^1} \\ \frac{\sigma_t^2\left(K_t\left(\delta_t^2 - r_t + n_t\theta_t\mu_t\right) + \sigma_t^2 R_t^2\right)}{K_t\left((\sigma_t^2)^2 + n_t(\theta_t)^2\mu_t\right)} - \frac{R_t^2}{K_t} \end{bmatrix}; \; \phi_t^i = -\frac{K_t\theta_t\left(\delta_t^2 - r_t + n_t\theta_t\mu_t\right) + \theta_t\sigma_t^2 R_t^2}{K_t\left((\sigma_t^2)^2 + n_t(\theta_t)^2\mu_t\right)}, \tag{4.7}$$

where the $\mathbb{F}$-predictable processes K and $R = (R^1, R^2)'$ are defined in Appendix A.3. Expression (4.7) shows that the quadratic hedging criterion commands a mortality risk premium involving a change of intensities, in the sense that ϕ^i is non zero in general. However, the doubly stochastic setup is not preserved under $\widetilde{\mathbb{P}}$, as (A4) is not satisfied because η and ϕ^i depend on the number of survivors at each point in time.

Things change if we assume that the liability does not depend directly on the death times affecting the traded security P^2. For example, P^2 may represent a security indexed on a mortality index broadly matching the demographic characteristics of the insureds. One way of formalizing this situation is to consider a liability $H_T^{(k)} = \sum_{i=1}^k 1_{\tau^i > T} f(P_T^1)$ and assume $\theta_i = 0$ for $i = 1, \ldots, k$ (with $k < m$). This means that the liability depends on the first k components of N (i.e., on $N^{(k)} := (N^1, \ldots, N^k)'$), while the dynamics of P^2 depends only on the remaining $m - k$ death indicators, (i.e., on $N^{(m-k)} := (N^{k+1}, \ldots, N^m)'$). In this case the pair (η, ϕ) takes the form

$$\eta_t = \begin{bmatrix} \frac{\delta_t^1 - r_t}{\sigma_t^1} \\ \frac{\sigma_t^2\left(K_t\left(\delta_t^2 - r_t + n_t^{(m-k)}\theta_t\mu_t\right) + \sigma_t^2 R_t^2\right)}{K_t\left((\sigma_t^2)^2 + n_t^{(m-k)}(\theta_t)^2\mu_t\right)} - \frac{R_t^2}{K_t} \end{bmatrix}, \qquad \phi_t^i = 0 \quad \text{for all } i, \tag{4.8}$$

where $n_t^{(m-k)} = (m-k) - \sum_{j=1}^{m-k} N_{t-}^j$. Taking $\mathbb{G}^{(m-k)} := \vee_{j=1}^{m-k}(\mathbb{G} \vee \mathbb{H}^{k+j})$ as our reference subfiltration, we see that assumption (A4) clearly holds with respect to $\mathbb{G}^{(m-k)}$: under $\widetilde{\mathbb{P}}$, the death times $\tau^1, \ldots, \tau^k$ are doubly stochastic with the same intensity process μ driven by $\mathbb{G}^{(m-k)}$; on the other hand, the $\widetilde{\mathbb{P}}$ dynamics of μ are clearly affected by η, and hence by $N^{(m-k)}$ through availability of security P^2. Expression (4.9) now takes the convenient form

$$\begin{aligned} V_0 &= \widetilde{E}\left[\exp\left(-\int_0^T r_s \mathrm{d}s\right) H_T^{(k)}\right] = \sum_{i=1}^{k} \widetilde{E}\left[\exp\left(-\int_0^T (r_s + \widetilde{\mu}_s)\mathrm{d}s\right) f(P_T^1)\right] \\ &= k\widetilde{E}\left[\exp\left(-\int_0^T (r_s + \mu_s)\mathrm{d}s\right) f(P_T^1)\right]. \end{aligned} \tag{4.9}$$

The above examples can be compared with the analysis offered by Dahl and Møller (2006), who restrict their attention to the case of deterministic ϕ and affine intensities of mortality. They first examine risk-minimizing strategies under a given $\widetilde{\mathbb{P}}$ (i.e., strategies that are self-financing only on average), and then determine the equivalent martingale measure supporting mean-variance indifference prices and quadratic hedging prices, always obtaining a null adjustment ϕ. The above examples show that the intensity μ needs not be the same process under $\widetilde{\mathbb{P}}$ when a more general setting is examined.

5 Calibration example

In this section, we provide an example of calibration of a risk-neutral LC model to the survival probabilities implied by a mortality table employed in the French annuity market (referred to as the TPRV life table). We estimate the margins to be added to a classical LC model fitted to French population data and find evidence that the underlying change of measure involves a change in the intensity process and not just a drift adjustment in the dynamics of the time trend.

As a simple example, consider an annuity issued to an individual aged x at time 0 and paying unitary amounts at dates in $\mathcal{T} = \{t_1, \ldots, t_m\}$. From the results of Section 4, we can exploit the independence assumption (under $\widetilde{\mathbb{P}}$) to write the time-0 price of such security as

$$\widetilde{E}\left[\sum_{t\in\mathcal{T}} B_t^{-1} 1_{\tau>t}\right] = \sum_{t\in\mathcal{T}} \widetilde{E}\left[B_t^{-1}\right] \widetilde{E}\left[\exp\left(-\int_0^t \widetilde{\mu}_s \mathrm{d}s\right)\right] = \sum_{t\in\mathcal{T}} \widetilde{E}\left[B_t^{-1}\right] \widetilde{\mathbb{P}}(\tau > t).$$

The last expression shows that we can separately calibrate, the financial component to zero-coupon bond prices, the mortality component to risk-neutral survival probabilities. We provide an example of the latter below.

5.1 Lee-Carter modeling for the population mortality rates

We use data concerning French males, general population. They have been downloaded from the Human Mortality Database.[9] The period considered is 1960-2001 and the age range is 50-100. Figure B.1 depicts the shape of the mortality surface. More specifically, the $\ln \widehat{m_x}(t)$ are displayed in function of age x and time t, where $\widehat{m_x}(t)$ is the death rate at age x in calendar year t (obtained by dividing the observed number of deaths by the corresponding central exposure-to-risk). The fit of the LC model (2.4) to the French population data gives the results displayed in Figure B.3. The estimated parameters of the discretized Brownian motion with drift $\mathrm{d}\kappa_t = \delta \mathrm{d}t + \sigma \mathrm{d}W_t$ are given by $\widehat{\delta} = -0.666778$ and $\widehat{\sigma}^2 = 1.660415$.

< Figure B.1 about here >

5.2 Presentation of the TPRV life table

The life tables currently in use by the insurers operating in France were established by the French National Institute for Statistics and Economic Studies (Institut national de la statistique et des études économiques - INSEE). Since 1993, the insurers are allowed to use tables based on their own experience to establish premium rates for annuity contracts. These experience life tables have to be certified by an independent actuary and submitted to a commission for approval. In

[9] See http://www.mortality.org.

any case, the result cannnot go below a minimum given by the INSEE projected life table (that incorporates mortality trends for generations from 1887 to 1993). In practice, the projected life table for annuities (known as the TPRV 93 life table) is used as a single threshold. It includes a comprehensive table of the 1950 cohort and, applying an age shift in order to establish the policyholder's technical age, leads to premium amounts very similar to those obtained through a direct application of generation tables. Figure B.2 depicts the shape of the mortality surface given by the TPRV life table. More specifically, the $\ln m_x^{TPRV}(t)$ are displayed in function of age x and time t, where $m_x^{TPRV}(t)$ is the intensity of mortality prevailing at age x in year t implied by the TPRV. Note that the $m_x^{TPRV}(t)$'s do not fill a rectangular array of data.

The vast majority of annuity providers operating in France have adopted the TPRV life table for pricing and reserving. This table is conservative because it is derived from female mortality experience, at the general population level, but applied to both male and female annuitants. The wide adoption of this life table by the French insurance companies suggests that the industry considers that the basis risk is appropriately mitigated in this way. Expenses are not taken into account in the present study.

< Figure B.2 about here >

5.3 Estimation of the κ_t's keeping the population α_x's and β_x's

We first carry out the LC estimation on the basis of the TPRV life table by keeping the α_x's and β_x's fixed at their population values. The estimated κ_t's are then obtained by the linear regressing of the $\ln \widehat{m}_x(t) - \widehat{\alpha}_x$ on the $\widehat{\beta}_x$'s, without intercept and separately for each value of t. The resulting κ_t's are displayed in the third plot of Figure B.3. We see that the resulting time index is driven by a totally different process, with an estimated variance of 0.1192059. It is thus impossible to get a standard LC model consistent with the TPRV lifetable if the underlying change of measure only involves the Brownian motion driving the time index. The finding is consistent with what has been observed elsewhere (see Biffis and Denuit, 2006, for the case of the Italian annuity market) and is not surprising. Indeed, what we have actually done was testing whether the use of the TPRV table implies that annuitants' random death times are conditionally independent, given the population time-trend dynamics. The answer is likely to be negative, because there are well-documented discrepancies between population and annuitants risk characteristics and because insurers' annuity portfolios may not be large enough to diversify away the idiosyncratic risk associated with the individual death times (see Section 4.1). In the next section, we show how to identify the different risk premiums emerging in the standard LC model when moving to the risk-neutral world.

< Figure B.3 about here >

5.4 Estimation of the α_x's and κ_t's by keeping the population β_x's

We consider the change of measure (3.3)-(3.4) with $\eta \in \mathbb{R}$ and $\phi_t^x = a(x+t) + b(x+t)\kappa_t$ for all x, where the dynamics of κ under $\widetilde{\mathbb{P}}$ is assumed to follow

$$\mathrm{d}\kappa_t = (\delta - \eta\sigma)\mathrm{d}t + \sigma \mathrm{d}\widetilde{W}_t, \tag{5.1}$$

with $\widetilde{W}$ a one-dimensional Brownian motion and where the coefficients δ, σ have been estimated in Section 5.1. Our aim is to estimate the functions a and b, and the parameter η entering the drift of κ under $\widetilde{\mathbb{P}}$. We start by keeping the population β_x's and by estimating the α_x's and the time-index implied by the TPRV life table. We denote the implied time-index by $\widetilde{\kappa}$. The new estimates for the drift and volatility of $\widetilde{\kappa}$ will then enable us to recover the adjustment b underlying the change of measure, as we now show.

The fit of the LC model to the TPRV implied intensities gives the results displayed in Figure B.4. Note that the α_x's have been modified in order to satisfy the constraints (2.5). The parameters of the random walk with drift modeling the time trend are given next: $\widetilde{\delta}^{\widetilde{\kappa}} = -0.3075042$ and $(\sigma^{\widetilde{\kappa}})^2 = 0.02767568$. Now, using a superscript 'FM' to indicate that the quantity relates to French males, and 'TPRV' to the TPRV table, the intensity of mortality can be written as

$$\widetilde{\mu}_t^x = \exp\left(\alpha^{\mathrm{TPRV}}(x+t) + \beta^{\mathrm{FM}}(x+t)\widetilde{\kappa}_t\right)$$

where $\alpha^{\text{TPRV}} = \alpha^{\text{FM}} + a$ and $\widetilde{\kappa}$ has $\widetilde{\mathbb{P}}$-dynamics

$$\mathrm{d}\widetilde{\kappa}_t = \widetilde{\delta}^{\widetilde{\kappa}}\mathrm{d}t + \sigma^{\widetilde{\kappa}}\mathrm{d}\widetilde{W}_t,$$

where $\sigma^{\widetilde{\kappa}}$ is different from the volatility coefficient σ of κ (as seen from the estimates), so that $\widetilde{\kappa}$ and κ are two different processes under $\widetilde{\mathbb{P}}$. Now, we can also write

$$\widetilde{\mu}_t^x = \exp\left(\alpha^{\text{TPRV}}(x+t) + \left(\beta^{\text{FM}}(x+t) + b(x+t)\right)\kappa_t\right)$$

with κ having $\widetilde{\mathbb{P}}$-dynamics (5.1) and with

$$b(x+t) = \beta^{\text{FM}}(x+t)\left(\frac{\sigma^{\widetilde{\kappa}}}{\sigma} - 1\right).$$

The interpolated point estimates of the functions $a = \alpha^{\text{TPRV}} - \alpha^{\text{FM}}$ and b are displayed in Figure B.5. Finally, we can compute η as

$$\eta = \frac{\delta}{\sigma} - \frac{\widetilde{\delta}^{\widetilde{\kappa}}}{\sigma^{\widetilde{\kappa}}} = 1.33097.$$

The resulting estimates for the adjustment functions a, b and for the coefficient η can be employed for the fair valuation of annuity business in the framework of Section 4. They also quantify the adjustments implied by table TPRV for the different types of mortality risk described in Section 4.1. Specifically, let us see how the simple procedure described above has allowed us to disentangle an adjustment for longevity risk from the other risk margins. Given the estimate of η, the risk-neutral drift $\widetilde{\delta}^{\kappa}$ of the time-trend process κ is seen to be equal to -2.1838272. From this, we can carry out a standard LC projection based on the drift implied by the TPRV table and compare it to a projection based on real-world estimates. The results depicted in Figure B.6 tell us that a stronger decrease in the mortality-index κ is incorporated in the risk-neutral version of the LC model implied by the TPRV table. We note that the width of the risk-neutral bands (at ± 2 standard deviations) is equal to the width of the real-world bands, consistently with the fact that the same time-trend process is considered in the real and risk-neutral world.

< Figure B.4 about here >

< Figure B.5 about here >

< Figure B.6 about here >

6 Conclusion

In this work we have presented a self-contained analysis of stochastic mortality models in continuous time, emphasizing the meaning and implications of some of the key assumptions commonly used in the literature, and analyzing their stability under equivalent measure changes. As an example, we have introduced a continuous-time version of the classical model proposed by Lee and Carter (1992) and examined how the model can be consistently used under equivalent probability measures. We have provided applications related to the fair valuation of life insurance liabilities and mortality-indexed securities, including an example on quadratic hedging. Finally, we have offered an example of model calibration based on the French annuity market.

Acknowledgements

The financial support of the Belgian Government under the contract "Projet d'Actions de Recherche Concertées" ARC 04/09-320 is gratefully acknowledged by Michel Denuit and Pierre Devolder. This work was partly done when Enrico Biffis visited the Institute of Actuarial Science at the University of Louvain, supported by ARC 04/09-320. Michel Denuit also thanks the Banque Nationale de Belgique for financial support.

References

Artzner, P. and F. Delbaen (1995). Default risk insurance and incomplete markets. *Mathematical Finance*, vol. 5(3):187–195.

Bauer, D., F. Benth and R. Kiesel (2009). Modeling the forward surface of mortality. Working Paper, University of Ulm.

Bielecki, T. and M. Rutkowski (2001). *Credit Risk: Modeling, Valuation and Hedging*. Springer, Berlin.

Biffis, E. (2005). Affine processes for dynamic mortality and actuarial valuations. *Insurance: Mathematics & Economics*, vol. 37(3):443–468.

Biffis, E. and M. Denuit (2006). Lee-Carter goes risk-neutral: An application to the Italian annuity market. *Giornale dell'Istituto Italiano degli Attuari*, vol. LXIX:1–21.

Blanchet-Scalliet, C. and M. Jeanblanc (2004). Hazard rate for credit risk and hedging defaultable contingent claims. *Finance and Stochastics*, vol. 8:145–159.

Brémaud, P. (1981). *Point Processes and Queues. Martingale Dynamics*. Springer-Verlag, Berlin.

Brémaud, P. and J. Jacod (1978). Changes of filtrations and of probability measures. *Zeitschrift für Wahrscheinlichkeitstheorie*, vol. 45:269–295.

Brouhns, N., M. Denuit and J. Vermunt (2002). A Poisson log-bilinear approach to the construction of projected lifetables. *Insurance: Mathematics & Economics*, vol. 31:373–393.

Cairns, A., D. Blake and K. Dowd (2006a). A two-factor model for stochastic mortality with parameter uncertainty: Theory and calibration. *Journal of Risk & Insurance*, vol. 73:687–718.

Cairns, A., D. Blake and K. Dowd (2006b). Pricing death: Frameworks for the valuation and securitization of mortality risk. *ASTIN Bulletin*, vol. 36:79–120.

Cerny, A. and J. Kallsen (2007). On the structure of general mean-variance hedging strategies. *The Annals of Probability*, vol. 35:1479–1531.

Dahl, M. (2004). Stochastic mortality in life insurance: Market reserves and mortality-linked insurance contracts. *Insurance: Mathematics & Economics*, vol. 35(1):113–136.

Dahl, M. and T. Møller (2006). Valuation and hedging of life insurance liabilities with systematic mortality risk. *Insurance: Mathematics & Economics*, vol. 39:193–217.

Duffie, D. (2001). *Dynamic Asset Pricing Theory*. Princeton University Press, Princeton, third edn.

Duffie, D., M. Schroder and C. Skiadas (1996). Recursive valuation of defaultable securities and the timing of resolution of uncertainty. *The Annals of Applied Probability*, vol. 6(4):1075–1090.

Girosi, F. and G. King (2008). *Demographic Forecasting*. Princeton University Press, Princeton, New Jersey.

IASB (2004). *International Financial Reporting Standard N. 4*. International Accounting Standards Board, London.

Jamshidian, F. (2004). Valuation of credit default swaps and swaptions. *Finance & Stochastics*, vol. 8(3):343–371.

Jarrow, R., D. Lando and F. Yu (2005). Default risk and diversification: Theory and empirical implications. *Mathematical Finance*, vol. 15(1):1–26.

Kusuoka, S. (1999). A remark on default risk models. In *Advances in Mathematical Economics*, vol. 1, pp. 69–82. Springer-Verlag.

Lando, D. (1998). On Cox processes and credit risky securities. *Review of Derivatives Research*, vol. 2(2/3):99–120.

Lee, R. (2000). The Lee-Carter method for forecasting mortality, with various extensions and applications. *North American Actuarial Journal*, vol. 4(1):80–93.

Lee, R. and L. Carter (1992). Modeling and forecasting the time series of US mortality. *Journal of The American Statistical Association*, vol. 87(419):659–671.

Lim, A. (2005). Mean-variance hedging when there are jumps. *SIAM Journal of Control and Optimization*, vol. 44:1893–1922.

Lipster, R. S. and A. N. Shiryaev (2001). *Statistics of Random Processes*, vol. I-II. Springer, Heidelberg, second edn.

Milevsky, M. and S. Promislow (2001). Mortality derivatives and the option to annuitize. *Insurance: Mathematics & Economics*, vol. 29(3):299–318.

Milevsky, M., S. Promislow and V. Young (2006). Killing the Law of Large Numbers: Mortality risk premiums and the Sharpe ratio. *Journal of Risk and Insurance*, vol. 73(4):673–686.

Møller, T. and M. Steffensen (2007). *Market-valuation Methods in Life and Pension Insurance*. Cambridge University Press, Cambridge, UK.

Norberg, R. (1987). A note on experience rating of large group life contracts. *Mitteil. Verein. Schweiz. Vers. Math.*, vol. 87:17–34.

Norberg, R. (1989). Experience rating in group life insurance. *Scandinavian Actuarial Journal*, vol. 1989(4):194–224.

Protter, P. (2004). *Stochastic Integration and Differential Equations*. Springer-Verlag, Heidelberg, second edn.

Samuelson, P. (1963). Risk and uncertainty: A fallacy of large numbers. *Scientia*, vol. 98:1–6.

Schervish, M. (1995). *Theory of Statistics*. Springer, New York.

Segall, A., M. Davis and T. Kailath (1975). Nonlinear filtering with counting observations. *IEEE Transactions on Information Theory*, vol. 21(2):143–149.

Snyder, D. (1972). Filtering and detection for doubly stochastic poisson processes. *IEEE Transactions on Information Theory*, vol. 18(1):91–102.

Steffensen, M. (2000). A no arbitrage approach to Thiele's differential equation. *Insurance: Mathematics and Economics*, vol. 27(2):201–214.

Yong, J. and X. Zhou (1999). *Stochastic Controls: Hamiltonian Systems and HJB equations*. Springer, New York.

A Additional details and proofs

A.1 Section 2

We provide two useful results that are used throughout the paper. We refer the reader to Bielecki and Rutkowski (2001) for some classical references.

Proposition A.1. *In the setup of Section 2, the following results hold:*

(i) For every integrable random variable $H \in \mathcal{F}$, we have

$$E\left[1_{\tau>T}H\middle|\mathcal{F}_t\right] = \frac{1_{\tau>t}}{1-F_t}E\left[1_{\tau>T}H\middle|\mathcal{G}_t\right]. \tag{A.1}$$

(ii) For any bounded $\mathbb{G}$-predictable process Z and for each $0 \le t \le T$, we have:

$$E\left[1_{t<\tau\le T}Z_\tau|\mathcal{F}_t\right] = \frac{1_{\tau>t}}{1-F_t}E\left[\int_t^T Z_s \mathrm{d}F_s\middle|\mathcal{G}_t\right].$$

Proof. See, for example, Bielecki and Rutkowski (2001). □

A useful observation that will be employed below is that condition (A2) is not just equivalent to (A2$*$), but also to

(A2$\diamond$) $\mathbb{P}(\tau \le t|\mathcal{G}_{T^*}) = F_t$ for all t.

See Brémaud and Jacod (1978) for a proof. The above emphasizes that the conditional probability of death happening before time t only depends on the evolution of mortality (and other) risk factors up to that time. Nothing would change if we were given their entire paths.

Proof of Proposition 2.1. To find the $\mathbb{F}$-compensator of τ, i.e. the unique predictable process A such that $N-A$ is an $\mathbb{F}$-martingale, we fix $t>0$ and let π_n be a sequence of refining partitions of $[0,t]$ such that $\lim_{n\to\infty}\mathrm{mesh}(\pi^n)=0$. We then set $A^n_t := \sum_{\pi^n} E[N_{t_{i+1}} - N_{t_i}|\mathcal{F}_{t_i}]$ and observe that, since $\{\tau > t_i\}$ is an atom for $\mathcal{H}_{t_i}$, we can write

$$\begin{aligned} E[N_{t_{i+1}}|\mathcal{F}_{t_i}] =& E[N_{t_{i+1}}(1_{\tau\le t_i}+1_{\tau>t_i})|\mathcal{F}_{t_i}] = 1_{\tau\le t_i}N_{t_{i+1}} + 1_{\tau>t_i}E[N_{t_{i+1}}|\mathcal{F}_{t_i}] \\ =& N_{t_i} + 1_{\tau>t_i}\frac{E[N_{t_{i+1}}1_{\tau>t_i}|\mathcal{G}_{t_i}]}{\mathbb{P}(\tau>t_i|\mathcal{G}_{t_i})} = N_{t_i} + 1_{\tau>t_i}\frac{\mathbb{P}(t_i<\tau\le t_{i+1}|\mathcal{G}_{t_i})}{\mathbb{P}(\tau>t_i|\mathcal{G}_{t_i})}, \end{aligned}$$

where we have used Proposition A.1(i) and the fact that $N_{t_{i+1}}1_{\tau\le t_i} = 1_{\tau\le t_i}$ and $1_{\tau>t_i}N_{t_{i+1}} = 1_{t_i<\tau\le t_{i+1}}$. Now, we note that by (A2$\diamond$) the process $F_t = \mathbb{P}(\tau \le t|\mathcal{G}_t)$ is increasing. (Indeed, for all $s \le t$ we have $\{\tau \le s\} \subseteq \{\tau \le t\}$ and hence $\mathbb{P}(\tau \le s|\mathcal{G}_{T^*}) \le \mathbb{P}(\tau \le t|\mathcal{G}_{T^*})$. (A2$\diamond$) finally yields $F_s \le F_t$.) We can then write

$$E[N_{t_{i+1}} - N_{t_i}|\mathcal{F}_{t_i}] = 1_{\tau > t_i} \frac{\mathbb{P}(t_i < \tau \le t_{i+1}|\mathcal{G}_{t_i})}{\mathbb{P}(\tau > t_i|\mathcal{G}_{t_i})} = 1_{\tau > t_i} \frac{E\left[F_{t_{i+1}}\middle|\mathcal{G}_{t_i}\right] - F_{t_i}}{1 - F_{t_i}},$$

to obtain

$$A_t^n = \sum_{\pi^n} 1_{\tau > t_i} \frac{E\left[F_{t_{i+1}}\middle|\mathcal{G}_{t_i}\right] - F_{t_i}}{1 - F_{t_i}} \underset{n\to\infty}{\longrightarrow} \int_0^{t\wedge\tau} \frac{\mathrm{d}F_u}{1 - F_{u-}}.$$

Now, since F is continuous by (A1), we have $A_t = -\ln(1 - F_{t\wedge\tau})$. The $\mathbb{G}$-predictable increasing process $-\ln(1 - F) = \Gamma$ then coincides with the martingale hazard process Λ. By the Doob-Meyer decomposition, Λ is unique, in the sense that if two increasing processes Λ^1 and Λ^2 are martingale hazard processes, they are modifications up to τ. □

A.2 Section 3

To simplify exposition, we prove the factorization (3.3)-(3.4) in the single policyholder case (i.e., we set $m = 1$ in Section 3). A suitable martingale decomposition theorem (see Kusuoka, 1999, Thm. 3.2) yields that, under the assumptions of Section 2, and (A2) in particular, any $\mathbb{F}$-square-integrable martingale ξ can be expressed as

$$\xi_t = 1 - \int_0^t \varsigma_s \cdot \mathrm{d}W_s + \int_0^t \zeta_s \mathrm{d}M_s,$$

where $M_t = N_t - \Lambda_{t\wedge\tau}$ and ς, ζ are $\mathbb{F}$-predictable processes satisfying $E[\int_0^{T^*} ||\varsigma_t||^2 \mathrm{d}t] < \infty$ and $E[\int_0^{T^*} |\zeta_t|^2 \mu_t \mathrm{d}t] < \infty$. In view of the strict positivity of ξ, we have that its left continuous version ξ_- is strictly positive as well. We can then set $\varsigma = \xi_-^{-1}\eta$ and $\zeta = \xi_-^{-1}\phi$ for some $\mathbb{F}$-predictable processes η, ϕ and get the alternative representation:

$$\xi_t = 1 - \int_0^t \xi_{s-}\eta_s \cdot \mathrm{d}W_s + \int_0^t \xi_{s-}\phi_s \mathrm{d}M_s, \tag{A.2}$$

with $\phi > -1$ to ensure that ξ is strictly positive. It is a standard result that (A.2) admits a unique solution $\xi = \mathcal{E}(Z)$ given by the stochastic exponential of the semimartingale $Z_t = -\int_0^t \eta_s \cdot \mathrm{d}W_s + \int_0^t \phi_s \mathrm{d}M_s$ started at 0 (e.g. Protter, 2004, p. 84):

$$\mathcal{E}(Z)_t = \exp\left(Z_t - \frac{1}{2}[Z]_t^c\right) \prod_{0<s\leq t} (1+\Delta Z_s)\exp(-\Delta Z_s). \tag{A.3}$$

Here $[Z]_t^c$ is the continuous part of the quadratic variation process $[Z] \doteq [Z,Z]$ (i.e., the process making $Z^2 - [Z]$ a local martingale) and $\Delta Z_s = Z_s - Z_{s-}$. By employing (A.2) and the properties of the quadratic variation process (e.g. Protter, 2004, p. 75), we see that:

$$[Z]_t = [-\int_0^t \eta_s \cdot \mathrm{d}W_s + \int_0^t \phi_s \mathrm{d}M_s] = \int_0^t ||\eta_s||^2 \mathrm{d}[W]_s + \int_0^t \phi_s^2 \mathrm{d}[M]_s = \int_0^t ||\eta_s||^2 \mathrm{d}s + \int_0^t \phi_s^2 \mathrm{d}N_s,$$

since $[W]_t = t$ by Lévy's theorem (Protter, 2004, Thm. 39, Ch. II), and $\Delta M = \Delta N$ because M is continuous except for a jump of size 1 at τ (recall that in our setup the compensator $\Lambda_{\cdot\wedge\tau}$ is continuous; equivalently, τ is a totally inaccessible $\mathbb{F}$-stopping time: see Protter, 2004, p. 106). Hence, $[Z]_t^c = \int_0^t ||\eta_s||^2 \mathrm{d}s$. Moving on to the discontinuous part of (A.3), we have $\Delta Z_t = \phi_t \Delta M_t = \int_0^t \phi_s \mathrm{d}N_s$ and the factorization $\xi = \xi'\xi''$ holds, with ξ' given by (3.3) and ξ'' by (3.4).

Proof of Proposition 3.2. We recall that by (A4) we have $\xi_t' \in \mathcal{G}_t$ all t. By (A3) we also know that under $\mathbb{P}$ the process ξ'' admits the factorization $\xi_t'' = \prod_{i=1}^m \xi_t''^{(i)}$, where for each t the r.v.'s $\xi_t''^{(1)}, \ldots, \xi_t''^{(m)}$ are conditionally independent, given $\mathcal{G}_t$. Furthermore, by (A4) we have $E[\xi_t''^{(i)}|\mathcal{G}_t] = 1$. We can then show that (A1) holds by computing, for each i,

$$\begin{aligned} 1 - \widetilde{F}_t^i =& \widetilde{E}\left[1_{\tau^i>t}|\mathcal{G}_t\right] = \frac{E\left[\xi_t 1_{\tau^i>t}|\mathcal{G}_t\right]}{E\left[\xi_t\middle|\mathcal{G}_t\right]} = \frac{\xi_t' E\left[\prod_{j\neq i}\xi_t''^{(j)}\middle|\mathcal{G}_t\right] E\left[\xi_t''^{(i)} 1_{\tau^i>t}\middle|\mathcal{G}_t\right]}{\xi_t' E\left[\prod_{j\neq i}\xi_t''^{(j)}\middle|\mathcal{G}_t\right] E\left[\xi_t''^{(i)}\middle|\mathcal{G}_t\right]} \\ =& E\left[\xi_t''^{(i)} 1_{\tau^i>t}\middle|\mathcal{G}_t\right] = \exp\left(-\int_0^t \phi_s^i \mathrm{d}\Lambda_s^i\right)(1 - F_t^i) = \exp\left(-\int_0^t (1+\phi_s^i)\mathrm{d}\Lambda_s^i\right), \end{aligned}$$

which is clearly continuous. Moving on to (A2), for each i we can compute

$$\begin{aligned}\widetilde{\mathbb{P}}(\tau^i \leq t|\mathcal{G}_{T^*}) =& \frac{E\left[\xi'_{T^*}\xi''^{(i)}_{T^*}1_{\tau^i\leq t}\middle|\mathcal{G}_{T^*}\right]}{E\left[\xi'_{T^*}\xi''^{(i)}_{T^*}\middle|\mathcal{G}_{T^*}\right]} = \frac{E\left[\xi''^{(i)}_{T^*}1_{\tau^i\leq t}\middle|\mathcal{G}_{T^*}\right]}{E\left[\xi''^{(i)}_{T^*}\middle|\mathcal{G}_{T^*}\right]}\\ =& E\left[\xi''^{(i)}_{t}1_{\tau^i\leq t}\middle|\mathcal{G}_{T^*}\right] = E\left[\xi''^{(i)}_{t}1_{\tau^i\leq t}\middle|\mathcal{G}_{t}\right] = \widetilde{\mathbb{P}}(\tau^i \leq t|\mathcal{G}_t),\end{aligned}$$

showing that the equivalent property (A2$\diamond$) holds under $\widetilde{\mathbb{P}}$. Finally, we can show that for each t, s and $i \neq j$ we have

$$\begin{aligned}\widetilde{E}\left[1_{\tau^i>s}1_{\tau^j>s}\middle|\mathcal{G}_t\right] &= E\left[1_{\tau^i>s}1_{\tau^j>s}\xi''^{(i)}_s\xi''^{(j)}_s\middle|\mathcal{G}_t\right]\\ &= E\left[1_{\tau^i>s}\xi''^{(i)}_s\middle|\mathcal{G}_t\right]E\left[1_{\tau^j>s}\xi''^{(j)}_s\middle|\mathcal{G}_t\right] = \widetilde{\mathbb{P}}(\tau^i > s|\mathcal{G}_t)\widetilde{\mathbb{P}}(\tau^j > s|\mathcal{G}_t),\end{aligned}$$

so that (A3) holds under $\widetilde{\mathbb{P}}$. □

A.3 Example 3.3

The dynamics given in (3.7) can be obtained under the following assumptions (see Lipster and Shiryaev, 2001): $E\left[\psi^4\right] < \infty$; $\int_0^{T^*}\left(|a(t)| + |b(t)Y_t| + \sigma^2(t, Y_t)\right)\mathrm{d}t < \infty$, $\int_0^{T^*}\left(a^2(t) + b^2(t)Y_t^2\right)\mathrm{d}t < \infty$ a.s.; for each t, $\sigma^2(t, Y_t) \leq z_1\int_0^t(1 + Y_s^2)\mathrm{d}l(s) + z_2(1 + Y_t^2)$ a.s., where $z_1, z_2 > 0$ and $l(\cdot)$ is a nondecreasing right-continuous function taking values in $[0, 1]$; for each t, $\sigma(t, \cdot)$ is bounded away from zero; $c(\cdot)$ is bounded. Then, the *a posteriori* mean process $\Psi_t := E\left[\psi|\mathcal{F}_t^Y\right]$ is given by

$$\Psi_t = \frac{\Psi_0 + \Psi_0^{(2)} + \int_0^t \frac{c(s)}{\sigma^2(s,Y_s)}\left(\mathrm{d}Y_s - (a(s) + b(s)Y_s)\,\mathrm{d}s\right)}{1 + \Psi_0^{(2)}\int_0^t\left(\frac{c(s)}{\sigma(s,Y_s)}\right)^2\mathrm{d}s},$$

where the *a posteriori* variance $\Psi_t^{(2)} := E\left[(\psi - \Psi_t)^2|\mathcal{F}_t^Y\right]$ takes the form

$$\Psi_t^{(2)} = \Psi_0^{(2)}\left(1 + \Psi_0^{(2)}\int_0^t\left(\frac{c(s)}{\sigma(s, Y_s)}\right)^2\mathrm{d}s\right)^{-1}.$$

A.4 Example 4.2

The dynamics of the wealth process $X = X^{x,\pi}$ associated with the strategy $\pi = (\pi^1, \pi^2)'$ is

$$\begin{cases} \mathrm{d}X_t = r_t X_t \mathrm{d}t + \left(\pi_t^1(\delta_t^1 - r_t)) + \pi_t^2(\delta_t^2 - r_t)\right)\mathrm{d}t + \pi_t^1\sigma_t^1 \mathrm{d}W_t^1 + \pi_t^2\sigma_t^2\mathrm{d}W_t^2 + \pi_t^2\sum_{i=1}^m \vartheta_t^i \mathrm{d}N_t^i \\ X_0 = x > 0. \end{cases}$$

The $\mathbb{F}$-predictable process K appearing in (4.7) has dynamics given by the backward stochastic differential equation (BSDE)

$$\begin{cases} \mathrm{d}K_t = K_t \left(-2r_t + \left(\frac{\delta_t^1 - r_t}{\sigma_t^1} + \frac{R_t^1}{K_t}\right)^2 + \frac{\left(K_t(\delta_t^2 + n_t\theta_t\mu_t - r_t) + \sigma_t^2 R_t^2\right)^2}{(K_t\sigma_t^2)^2 + n_t(K_t\theta_t)^2\mu_t}\right)\mathrm{d}t + R_t^1\mathrm{d}W_t^1 + R_t^2\mathrm{d}W_t^2 \\ K_T = 1, \end{cases} \tag{A.4}$$

where R^1 and R^2 are square integrable $\mathbb{F}$-predictable processes. (For the example with liability $H_T^{(k)}$, n_t is replaced by $n_t^{(m-k)}$.) We note that no differentials $\mathrm{d}N$ or $\mathrm{d}N^{(m-k)}$ appear in (A.4) because all the coefficients in the dynamics of P^1, P^2 are assumed to be $\mathbb{F}$-predictable. As opposed to usual stochastic differential equations, in (A.4) the terminal, rather than initial, condition is given. However, no time reversal can be used, since a solution (K, R^1, R^2) to (A.4) needs to satisfy measurability conditions with respect to the (forward) filtration $\mathbb{F}$. For existence and uniqueness of solutions to (A.4), we refer to Lim (2005) and references therein. We note that when all coefficients are deterministic, R^1 and R^2 are the null process and BSDE (A.4) admits an explicit solution (see Yong and Zhou, 1999). The quadratic hedging strategy for $H_T \in \mathcal{F}_T$ makes wealth evolve under $\widetilde{\mathbb{P}}$ according to the following BSDE:

$$\begin{cases} \mathrm{d}X_t = r_t X_t \mathrm{d}t + \sigma_t^X \cdot \mathrm{d}\widetilde{W}_t + \sum_{i=1}^m \vartheta_t^{X,i}\mathrm{d}\widetilde{M}_t^i \\ X_T = H_T, \end{cases} \tag{A.5}$$

where the random coefficients $\sigma^X, \vartheta^{X,i}$ are square-integrable and $\mathbb{F}$-predictable, and $\widetilde{W} = W + \int_0^{\cdot} \eta_s \mathrm{d}s$, $\widetilde{M}^i = M^i - \int_0^{\cdot\wedge\tau^i} \phi_s^i \mathrm{d}s$ are respectively a two-dimensional Brownian motion and a compensated doubly stochastic process (up to its first jump) with intensity $\widetilde{\mu}^i = (1 + \phi^i)\mu^i$ under $\widetilde{\mathbb{P}}$. Existence and uniqueness of solutions to (A.5) are proved in Lim (2005) for the case of bounded H_T.

B Figures

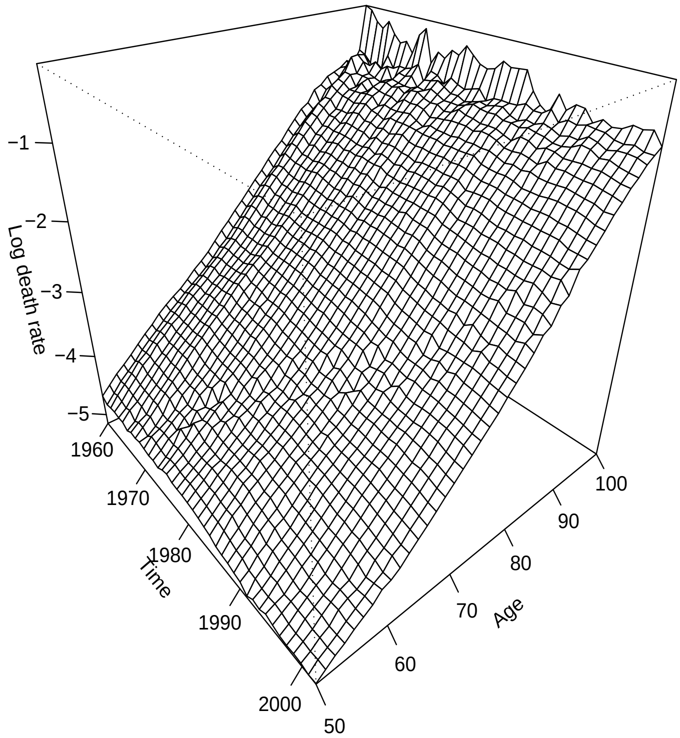

Figure B.1: Mortality surface for the French males, general population: plot of the empirical death rates $\widehat{m_x}(t)$ on the logarithmic scale, as a function of age x and calendar year t.

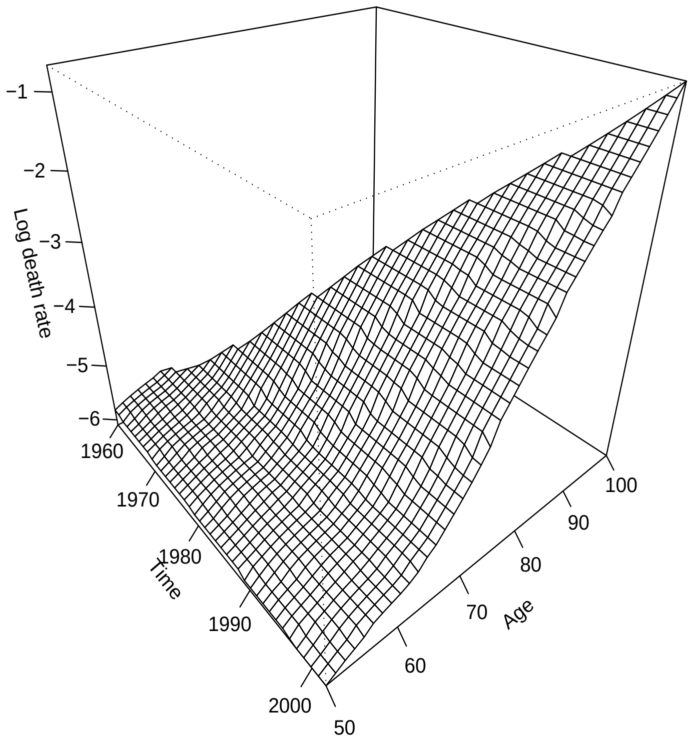

Figure B.2: Mortality surface corresponding to the TPRV life table: plot of the death rates derived from the regulatory life table on the logarithmic scale, as a function of age x and calendar year t.

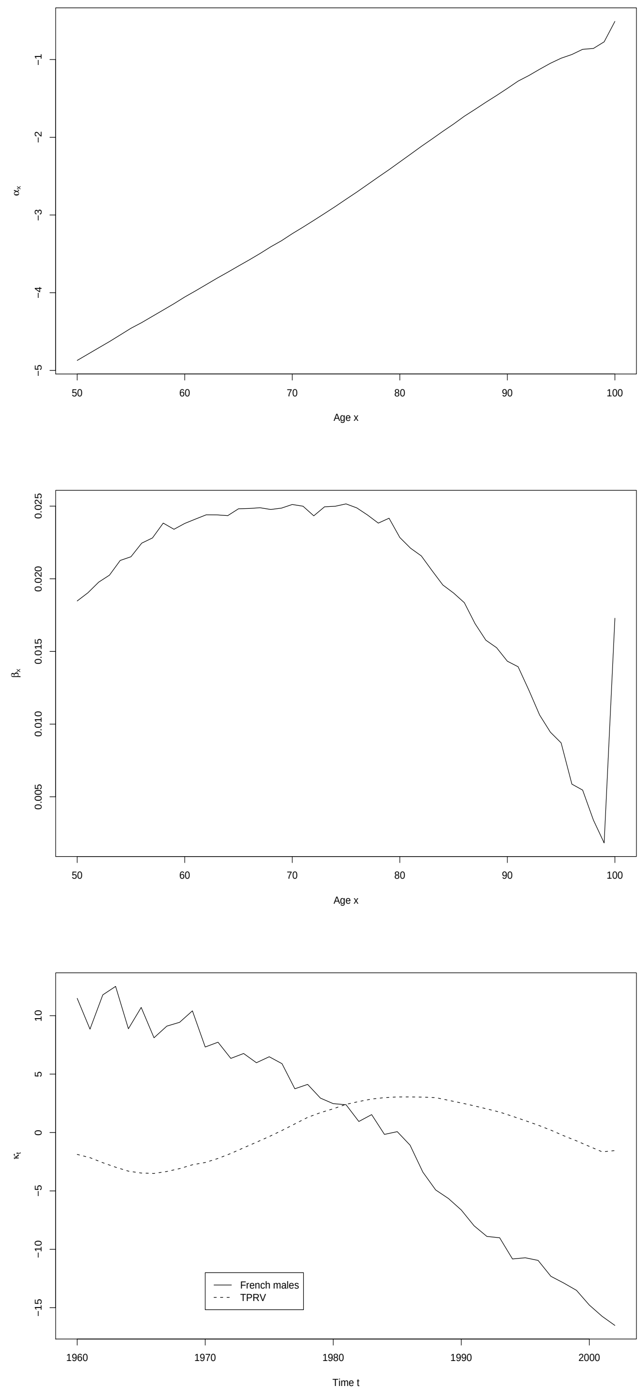

Figure B.3: Estimated Lee-Carter parameters α_x (top panel) and β_x (middle panel) for French males, general population, and comparison between the estimated Lee-Carter parameters κ_t for French males, general population, and estimated Lee-Carter parameters κ_t derived from the TPRV life table keeping the general population α_x and β_x (bottom panel).

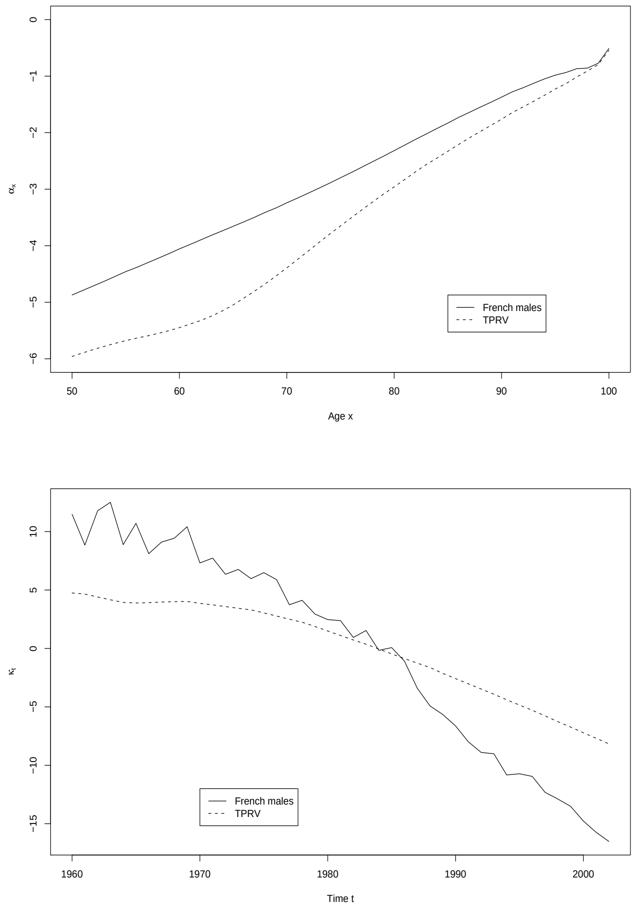

Figure B.4: Comparison between the estimated Lee-Carter parameters α_x (top panel) and κ_t (bottom panel) for French males, general population, and those dervived from the TPRV life table keeping the population β_x's.

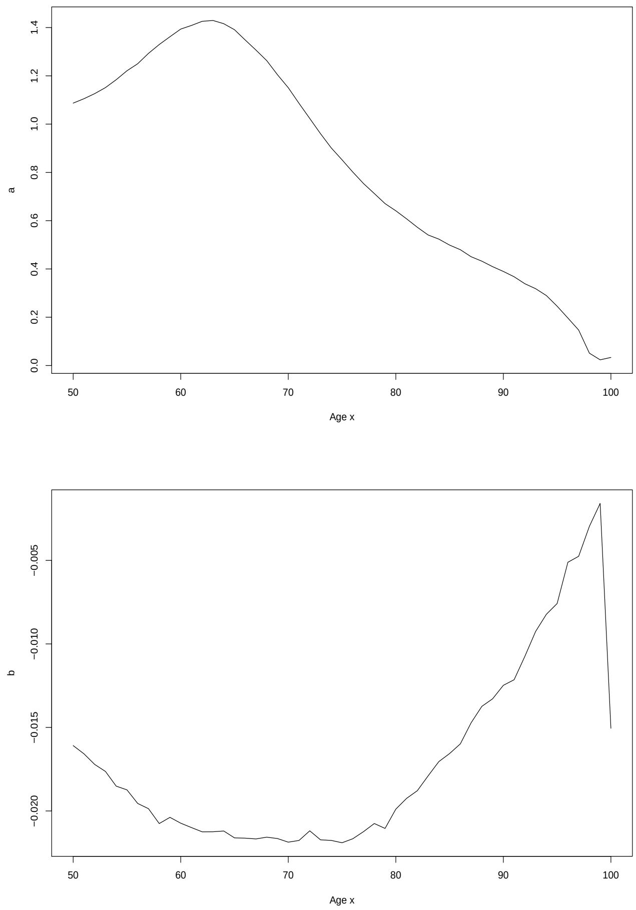

Figure B.5: Estimated functions a (top panel) and b (bottom panel).

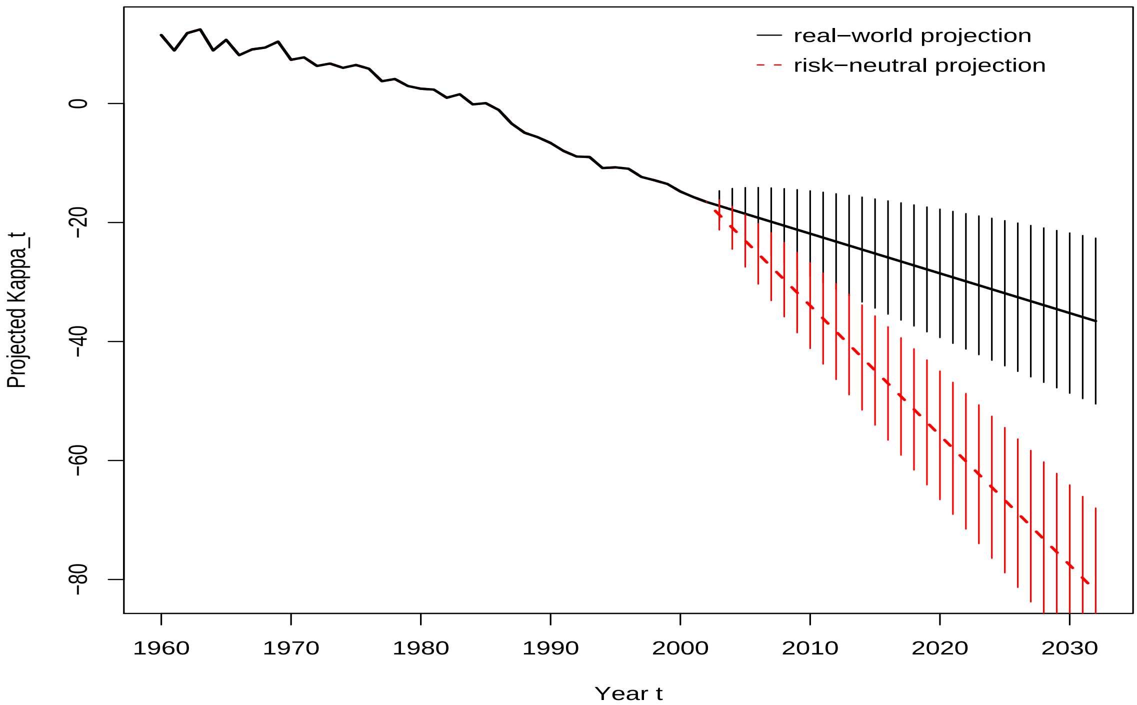

Figure B.6: Real-world and risk-neutral (TPRV-implied) mortality projection.