Institutional Members: CEPR, NBER and Università Bocconi

WORKING PAPER SERIES

Progress in Medicine, Limits to Life and Forecasting Mortality

Carlo A. Favero and Marco Giacoletti

Working Paper n. 406

This Version: July 21, 2011

IGIER – Università Bocconi, Via Guglielmo Röntgen 1, 20136 Milano –Italy
http://www.igier.unibocconi.it

Progress in Medicine, Limits to Life and Forecasting Mortality

Carlo A. Favero, Marco Giacoletti*

July 26, 2011

Abstract

In this paper we propose a model to forecast future mortality that includes information on the limits to life and on progress in medicine. We apply the model to forecasting future mortality and survival rates for the males population in England and Wales. Our proposal extends the benchmark stochastic mortality model along two dimensions. First, we try and deal explicitly with tail risk in the cross-sectional estimation. by including information about the "limit to life" in the sample used to construct factors for the cross-sectional dimension of mortality rates. Second, we propose to substitute the usual stochastic trend model adopted for the time series of risk factors with a predictive framework based on available evidence on medical progress and causes of death. The model projects very little variability for limits to life over the next ten years and predicts that in 2020 the probability that an individual age 65 will survive until 85 is 20% with an upper bound of 23% and a lower bound of 17%.

Keywords: stochastic mortality, limits to life, medical progress, longevity risk, compression of morbidity, modal age of death

*Carlo Favero, Deutsche Bank Chair in Asset Pricing and Quantitative Finance, Università Bocconi, IGIER and CEPR. Via Roentgen 1, Milan 20136, Italy, carlo.favero@unibocconi.it. Marco Giacoletti, Dept of Finance, Università Bocconi, Via Roentgen 1, Milan 20136, Italy,marco.giacoletti@unibocconi.it. We thanks, without implicating them, Michael Amori, David Blake and Michele Foresti for providing us with comments and suggestions.

Introduction

Forecasting mortality after 65 is both a challenging and a relevant problem. The evolution of population dynamics at old ages in industrialized countries in the second half of 20th century is a largely debated issue in demographics and natural sciences with important financial implications. Insurance companies and pension plan providers face longevity risk, i.e. the risk that retirees might on average live longer than expected. In many countries this risk was traditionally born by private corporations or by the government. However, regulators are pushing for a shift towards definite contribution plans; under this new regime pensions will be paid more and more in the form of annuities, provided by life insurance companies. Under the new regime longevity risk matters, since annuities are based on contributions made by the insured before retirement and are paid until the individual is alive.

The total longevity risk in a country can be decomposed in two elements. The first one is a specific risk, which concerns single individuals and very specific subgroups of the population (for example miners or employees of the chemical sector). This idiosyncratic risk can be managed by holding a large and well diversified pool of clients.

The second element is the aggregate risk, which is due to the uncertainty around the general downward trend in mortality rates. This aspect is a structural element, which cannot be hedged either with diversification or financial strategies. Thus, insurance companies will not only be forced to bear the risk of selling annuities to a population whose life expectancy is increasing, but will also be forced by Solvency II to hold significant additional capital reserves, since their risk cannot be hedged[1].

Pricing of this aggregate longevity risk requires predicting mortality rates and having a measure of the uncertainty of predictions, in fact, longevity risk *is increasing* in the conditional volatility of mortality rates. In this paper we propose a model to forecast future mortality that includes information on the limits to life and on progress in medicine and we apply it to forecasting

[1]Longevity bonds have been recently proposed as a financial instrument capable of providing an hedge against longevity risk. The idea for longevity bonds was firt published in the Journal of Risk and Insurance in 2001.A *longevity bond* pays a coupon whose size decreases along with the survival probability of individuals who have reached a certain age. *Longevity bonds* have no principal repayment and coupons are "indexed" evolution of survival probabilities of individuals who were born in the same year (so called "cohort"). For a discussion of these instruments, see Blake Boardman and Cairns (2010).

future mortality and survival rates for the males population in England and Wales. After the seminal contribution of Lee and Carter(1992), the standard approach to forecasting mortality at old ages is based on dynamic factor models for the time series of mortality rates at ages between 65 and 89.

Modeling mortality involves a cross-section and a time-series dimension. The cross-sectional dimension is generated by the fact that data on observed mortality at each age between 65 and 89. In other words, an age-structure of mortality rates is observable at any given year. The time-series dimension is generated by the fact that mortality rates at different ages are observed over a sample of time series data that, in the case of the UK, is 1971-2009. Stochastic mortality models typically capture the cross-sectional dimension of the term structure of mortality via a small numbers of stochastic factors; these stochastic factors are then projected ahead via simple time-series models to derive future survivor functions and mortality rates along with the associated uncertainty. Figure 1 reports mortality rates, survivor probabilities and frequencies of deaths in the UK and Wales for males aged from 65 onwards in 1971 and 2009[2].

Insert Figure 1 here

Several interesting facts emerge from Figure 1. First, mortality rates at all ages have witnessed a sizeable reduction over the last 40 years. Second, such a reduction is not uniform and mortality improvements at old ages have been more drastic than the ones for individuals aged between 65 and 70. As a consequence, survivor probabilities have also changed in an heterogenous way: in 1971 an individual alive at age 65 would have a probability of 81% of being alive at 69, a probability of 34% of being alive at 79, and a probability of 5% of being alive at 89; such probabilities have shifted respectively to 92%, 64% and 20% in 2009. Third, we can observe the so called "rectangularization"[3] of frequencies of death: the modal age of death (the age with the

[2]We denote with $q(x,t)$ the mortality for individuals of age x in year t, where mortality is the probability that a person aged x and alive at the beginning of the year dies within the end of the year. We then define $s(x,t)$ as the survivor probability for individuals of age x in year t, which is the probability that an individual will be alive at age x given that he has survived up to age $x-1$. Survivor probabilities are derived recursively: if $x=65$ then $s(x,t)=1-q(x,t)$, if $x>65$ then $s(x,t)=s(x-1,t)[1-q(x,t)]$. Eventually, frequencies of death for individuals of age x at time t are determined as first differnces of survival probabilities: $fod(x,t)=s(x,t)-s(x+1,t)$.

[3]This term was coined by James F. Fries; the phenomenon is also called "compression

highest frequency of death) is shifting towards the right. In 1971 the modal age of death was 74 years, while in 2009 it was 84 years. Fourth, the available data do provid evidence on a question that is becoming increasingly more important with the lenghtening of life: "What happens after the age of 89?". In this paper, we investigate the properties of an extended version of the benchmark stochastic mortality model proposed by Cairns, Blake and Dowd (2006). Our proposal extends the benchmark model along two dimensions. First, we try and deal explicitly with tail risk in the cross-sectional estimation. The benchmark model of the age-structure of mortality is fitted to the observed mortality rates, which span the ages 65-89. The out-of-sample cross sectional projections of these models deliver implausible mortality rates for ages between 90 and 120. We try and correct for this problem by including in the estimation sample information about the "limit to life" (age of the oldest person alive in the population). In other words we complement the information contained in mortality tables for individuals aged between 65 and 89 by adding for each year the information of the age at which the rate of mortality reaches the value of one. this limit to life is mesured by the age of the oldest living person in the world. Second, for what concerns the time series dimension, we propose to substitute to the usual stochastic trend model adopted for the risk factors a predictive framework based on available evidence on medical progress. In practice, use data on causes of mortality for males in England and Wales to predict factors over time; these data are made available by the Institute of English Actuaries. The Institute provides a time series of the average age standardized mortality rate for people between 60 and 89. This rate is then disaggregated into the main causes of death, namely hearth and circulatory diseases, lung cancers, other cancers and all other causes. We provide evidence on the statistical significance of this information in explaining the observed trends in the time series of the factors explaining mortality and we build this information in the forecasting models. Our framework can be used to evaluate the potential impact of breakthroughs in different fields of medicine on the projections for future

of morbidity". Morbidity consists in the quiet path towards death followed traditionally by old people. In other words, many people survived to an age higher than the model age of death, and the profile of frequencies of death above the mode was pretty smooth. However, the increase in the modal age of death has made the profile of frequencies of death for ages above the mode close to a straight vertical line. That's why morbidity has been "compressed". See Vladimir Canudas Romo (2008), Jean Marie Robine (2008) and Nadine Ouellette (2009 and 2011).

mortality rates and longevity risk.

We propose and discuss our model in Section 2, where we also compare it with the benchmark specification. In Section 3 and 4 we present the out-of-sample performance, finally we devote a section to compare model derived mortality improvements based on different assumptions for the future evolution of causes of mortality.

1 A stochastic model with causes of mortality and limits to life

Stochastic mortality models deal simultaneously with the cross-sectional and the time-series dimension of mortality data: a factor model is fitted to the cross-sectional data to capture with a limited number of factors the entire cross-sectional variation of mortality in any given period; then a time-series model is specified for the factors to obtain projections and to reconstruct the entire mortality curve in any future period. Stochastic simulation of the model allows to explicitly recognize the role of uncertainty and to associate confidence intervals to predicted mortality and survivor functions. We propose the following stochastic model for mortality at ages over 65[4]:

$$
\begin{aligned}
q(x,t) &= \frac{e^{c_{1t}+c_{2t}(x-\bar{x})+c_{3t}(x-\bar{x})^2+c_{4t}(x-\bar{x})^3}}{1+e^{c_{1t}+c_{2t}(x-\bar{x})+c_{3t}(x-\bar{x})^2+c_{4t}(x-\bar{x})^3}} + u_{x,t} \\
c_{1,t+1} &= \beta_1 + \beta_2 c_{1,t} + \beta_3 \log \bar{m}_{t+1} + v_{1,t+1} \\
c_{2,t+1} &= \beta_4 + \beta_5 c_{2,t} + v_{2,t+1} \\
c_{3,t+1} &= \beta_6 + \beta_7 c_{3,t} + v_{3,t+1} \\
c_{4,t+1} &= \beta_8 + \beta_9 c_{4,t} + v_{4,t+1} \\
\begin{bmatrix} v_{1,t+1} \\ v_{2,t+1} \\ v_{3,t+1} \\ v_{4,t+1} \end{bmatrix} &\backsim N\left(\begin{matrix} 0 \\ 0 \\ 0 \\ 0 \end{matrix}, \sum \right) \\
\bar{m}_{t+1} &= m^h_{t+1} + m^{cd}_{t+1} + m^{lc}_{t+1} + m^{oc}_{t+1} + m^{oth}_{t+1}
\end{aligned}
$$

[4]This choice is common to studies focusing on longevity. Bear in mind that longevity risk concerns only people in the oldest segments of the population. Thus, it makes sense to focus on the fit of the model to the last part of the term structure, rather than on the one to the whole curve.

Where $q(x,t)$ is the mortality at age x in year t, defined as the probability that a person aged x and alive at the beginning of the year dies within the end of the year. The cross-sectional variation of mortality is captured by four factors: $c_{i,t+1}$, $i \in \{1,2,3,4\}$ .The first factor is not age specific, while the other three factors capture the cross-sectional evolution of mortality rates at different ages. We relate improvement in mortality rates at all ages to $\log \bar{m}_{t+1}$, where $\bar{m}_{t+1}$is the average age standardized mortality rate for people between 60 and 89.

This mortality rate is generated via the aggregation of the main causes of death. Namely hearth, m^h_{t+1}, and circulatory diseases,m^{cd}_{t+1} lung cancers,m^{lc}_{t+1} other cancers, m^{oc}_{t+1}, and all other causes,m^{oth}_{t+1}. This model extends the Cairn, Blake and Dowd (2006) (CBD) model. CBD is obtained from our specification by setting $c_{3,t+1} = 0, c_{4,t+1} = 0$, and $\beta_3 = 0$ and $\beta_2 = \beta_5 = 1$.

We take from CBD the logit specification of mortality because it has the intuitive property of constraining mortality, which is effectively a conditional probability, to lie between zero and one.

We introduce the higher order factors $c_{3,t+1}$ and $c_{4,t+1}$ to obtain a better measure of mortality at old ages, as we fit our specification to the mortlaity rates between 65 and 89 augmented with the information on limits to life. The point is illustrated in Figure 2. Figure 2 reports the mortality function in year 2000 for the age range 65-120; projections based on the CBD model are plotted along with those generated by our extended model, the observed mortality rates for the age range 65-89 and the limit to life in year 2000 which sets a mortality rate at a value of 1 for the age 113, as 113 was the age of the oldest living person on earth in year 2000 .

Insert Figure 2 here

The CBD model fits very well the observable data but it generates implausible mortality for the age between 90 and 120 years. In fact, mortality grows too slowly with age (note, for example, that mortality at 100 years is only 50%). In other words, the model seems to overestimate the survival probability for very old individuals and therefore overstates tail risk.

In order to address the issue of underestimation of mortality at age over 89 in the CBD model we have augmented the dataset on mortality by collecting evidence on the oldest living person in the world over the period 1971-2009[5].

[5]For further reference see the website of the Gerontology Research Group, www.grg.org, and the following link http://en.wikipedia.org/wiki/Oldest_people# Chronologi-

This evidence, reported in Figure 3, shows that the age of the oldest living person in the world over the relevant sample period increases, but it does not change dramatically over time.

Insert Figure 3 here

We have incorporated this information in our dataset by setting the mortality rate equal to 1 in correspondence of the age of the eldest person in the world population. We have then re-estimated the logit model by including this extreme observations on the "limit to life". Our logit model for mortality rates is not any longer linear; a quadratic and a cubic trend in age are necessary to make the fit on the extended cross-section comparable to the one of the original model on the age span from 65 and 89. The quadratic and the cubic trend in age are in practice captured by the factors $c_{3,t+1}$ and $c_{4,t+1}$.

Note how the shape of the age-structure of mortality over 90 changes dramatically. The quadratic component determines a quick increase in mortality rates for people above 89 years old; the cubic component act as a linking spline between the linear in-sample model and the quadratic out-of-sample component. In year 2000 the oldest person in the population was 113; at this age the mortality rate reaches one in our model, while the original CBD specification is far from capable of replicating this structure. On the other hand, the performances of the two models for the age range 65-89 are very similar.

The second important sources of difference between our model and the traditional factor models of mortality lies in the specification of the time-series models for the factors. In fact, CBD adopts a unit root specification for all factors, and projects future values on the basis of the drift and conditional volatility estimated on past data. This procedure ensures that the uncertainty associated to the future evolution of factors is not underestimated, since the variance of projections diverges to infinite for all factors as the predictive horizon grows to infinity. However, this simple time series model appears to be based on a very restrictive set of assumptions, that for instance do not allow to link the stochastic process of the factors to exogenous evidence on medical progress and causes of death. In our specification the only factor that has a trend is $c_{1,t+1}$. We relate this trend to medical progress by making it a function of the average standardized mortality rates

cal_list_of_the_verified_oldest_living_man_since_1961

at ages between 65 and 90, which aggregates mortality rates for different causes. Figure 4 motivates intuitively the relation between c_{1t} and $\bar{m}_t$ in our specification.

Insert Figure 4 here

The Figure reports the time-series of the estimated first factor in the CBD model, c_{1t}, and logarithm of the average age-weighted mortality rate, $\bar{m}_t$. There is a close match between the behavior of the two series, which relates the downward trend in c_{1t} to the general improvements in mortality over time at all ages.

Since we build explicitly the relationship between c_{1t} and $\bar{m}_t$ into our stochastic model, we are able to evaluate the impact on future mortality projections of different scenario for progress in medicine. In fact, $\bar{m}_t$ is obtained by summing the average mortality rates for different causes: $m^h_{t+1}, m^{cd}_{t+1}, m^{lc}_{t+1}, m^{oc}_{t+1}, m^{oth}_{t+1}$. Possible breakthroughs in different fields of medicine can be incorporated into the model by providing future patterns for these (exogenous) variables. Figure 5 illustrates this point by showing that the fall in average mortality due to vascular diseases has dominated the aggregate trend, causing a general decline in the average age weighted mortality rate; starting from the seventies, the average mortality rate for hearth and circulatory diseases has declined, while the mortality rate for cancers has remained stable over time.

Insert Figure 5 here

Quentions on the effect on future mortality of a medical breakthrough that drives to zero the average death rate for lung cancer can be promptly simulated within our specification, It would be impossible to "parameterize" the same question in a model with stochastic trends for the factors determining mortality at different ages. For what concerns the other factors, we do not impose a priori a stochastic trend in the age related factors (in fact we do not impose the CBD coefficient restrictions $\beta_2 = \beta_5 = 1$). After having related the trend c_{1t} to $\bar{m}_t$, we let the data speak freely about the time series behavior of the other factors.

1.1 Placing our Model in the Literature

Our model falls in the field of stochastic mortality frameworks. Stochastic mortality approaches traditionally focus on the term structure of mortality

rates and have been mainly developed in order to answer the need of life insurers and of the financial sector for a reliable probabilistic structure to deal with population dynamics and their uncertainty. The two most popular frameworks introduced so far are the one by Lee and Carter (1992) and the one by Cairns Blake and Dowd (2006). One of the main practical implementation of this research has been the estimation and projection of longevity risk in England and Wales (but also in the United States). While in the Lee and Carter model the dynamics of mortality are captured just by one single factor, the CBD model uses two different stochastic factors, one for the intercept and one for the slope of the logit equation for mortality rates. Both models have been extended to take into account cohort effects (Renshaw and Haberman, 2006). In general, CBD shows a better fit to the data (see Cairns Blake and Dowd et Al. 2008).

In this paper we propose an extension to the original CBD framework along two dimensions: modelling mortality rates in the neighbourhood of the "limit to life" and using information on the different causes of mortality in the time series domain of the factors. The main objective of our specification is to establish a closer link between the statistical framework and the available demographics data (and their projections). Moreover, we use data on causes of mortality to explain the general improvement of population longevity and to run scenarios for future trends. The model is fitted on mortality rates, but its feature allows us to obtain model consistent projections for the maximum age of death and for frequencies of death. This makes our work closer to the debate on maximum life expectations (see Vallin and Meslé 2010), and to the one on the progressive increase of the modal age of death in low mortality countries (Robine Michel and Institut 2008, Canudas Romo 2008, Oullette 2011). In fact, our model can be useful also to simulate alternative scenarios for the "compression of morbidity" in the next decades.

Stochastic mortality models are not the only approach used to project future mortality in the UK. An alternative framework is provided by the Institute of English Actuaries and is called Constant Mortality Investigation (CMI). The CMI model is considered the benchmark for the projection of mortality rates in England and Wales. The model is based on scenarios for future Mortality Improvement Rates, which are defined as the rates of change of mortality rates. The future trajectory of improvement rates is obtained via a spline method that takes into account both cohort and age specific effects. The projections are based on the definition of current mortality improvement rates and of asymptotic long-term improvement rate. Once the

initial and the long-term improvement rate for a specific cohort are chosen, the future evolution of improvements is assumed to converge smoothly to the long-term rate with a half-life of 10 years (see CMI Working Papers 38, 39 and 49). The framework is not probabilistic and no uncertainty is associated to future trajectories. Different assumptions on the Long-term improvement rate deliver different patterns of future improvements and therefore different shapes of the term structure of mortality rates. The Institute of English Actuaries is now suggesting 1.25% ("one and a quarter") as a "base case scenario" for the Long-term improvement rate. An alternative "worst case" scenario is an improvement rate of 3.25% ("three and a quarter"). CMI is the provider of the data on average age standardized mortality rates per cause of death that we use in the estimation of our model. In addition, CMI releases projections for the future evolution of causes of death consistent with its scenarios. Our framework allows to feed such scenarios in a factor model and derived model based prediction with the associated uncertainty based on the CMI projections for the exogenous variables.

2 Estimation and Out-of-Sample Predictive Performance

Our empirical application is on England and Wales mortality rates over the sample 1971-2009 as available from Lifemetrics[6]. The data of the time series of average mortality rates for different causes are taken from CMI[7]. The data on the age of the eldest person in the world population has been reconstructed by using information available from the ebsite of the Gerontology Research Group and wikipedia[8].

The model delivers a good fit to the cross-section of mortality rates at any year in the historical sample. To provide evidence on this issue we report in Figure 6.1 the estimated factors along with their 95 per cent confidence intervals and in Figure 6.2 the cross-sectional R^2 for all years.

Insert Figures 6.1-6.2 here

[6](http://www.jpmorgan.com/pages/jpmorgan/investbk/solutions/lifemetrics/data)

[7](http://www.actuaries.org.uk/research-and-resources/pages/continuous-mortality-investigation-data)

[8]http://www.grg.org, and http://en.wikipedia.org/wiki/Oldest_people#Chronological_list_of_th living_man_since_1961

Interestingly, the only factor that features a clear non-mean reverting behavior is c_{1t}. We report in Table 1 results from the estimation of the time-series models for the factors.

Insert Table 1 here

All coefficients are significantly different from zero, in particular the coefficient on the average age standardized mortality rate is positive, very significant and estimated with high precision. R-squares are high for all equations and residuals do not show any mis-specification symptoms (see Figure 7).

Based on our estimation results, we conduct two different projection exercises. First, we conduct pseudo out -of-sample analysis by estimating the model with data up to 2005 and then projecting the relevant variables up to 2009 to assess model predictions against realized data. Second, we generate true out of sample predictions for 2020.

Out-sample predictions require to generate a scenario for the exogenous variables $m^h_{t+1}, m^{cd}_{t+1}, m^{lc}_{t+1}, m^{oc}_{t+1}, m^{oth}_{t+1}$. We have proceeded as follows:

$$\begin{aligned}
&\log HC_{t+1} = \gamma_1 + \log HC_t + u_{1,t+1} \\
&\log OC_{t+1} = \gamma_2 + \log OC_t + u_{2,t+1} \\
&\log LC_{t+1} = \gamma_3 + \log LC_t + u_{3,t+1} \\
&\log Other_{t+1} = \gamma_4 + \log Other_t + u_{4,t+1} \\
&\log \bar{m}_{t+1} = \log \left[HC_{t+1} + OC_{t+1} + LC_{t+1} + Other_{t+1}\right]
\end{aligned}$$

Projections are then obtained by setting $\gamma_1 = -0.02$, $\gamma_2 = \gamma_3 = \gamma_4 = 0$. We report the results from pseudo out-of sample simulations and for out-of sample predictions respectively on the left and the right column of Figure 8. In the simulation exercise we expliciitly allow for parameters uncertainty so there each replication is based on a draw from the ditribution of the residuals of the models and the distribution of the estimated parameters.

Insert Figure 8 here

The empirical evidence shows that the model does well in pseudo-out of sample simulation especially at modelling mortality for age over 80.

The first column of figure 9 reports the results for true out-of-sample projections by showing mortality rates, the survivor function and predicted frequencies of death for age 65 and over in 2020, with the associated 95 per cent confidence interval. The model predicts further improvements in

mortality rates and a survivor function shifted above that observed in 2009 for ages for above 80 years, interestingly, the degree of uncertainty is rather limited. The projected probability that in 2020 an individual aged 65 will survive until 85 is 20% with an upper bound of 23% and a lower bound of 17%. The evidence from the simulation of Frequencies of death is consistent with a further progressive "compression of morbidity" of mortality in the next decade.

Note that our framework allows us to project the functions at all ages above 65. In particular, we report in Figure 3, the model based projections for the age of the oldest individual in the population, that shows very little variability moving from the current value of 114 years to a projected value of 116 years in 2020.

3 Compression of Morbidity: Scenario Analysis

In this section we exploit the dependence of the dynamic of our factors on exogenous variables to evaluate how different scenarios on the improvement rates of mortality impact on the evidence of compression of morbidity described in the previous section.

In particular we considere the model based projections derived in the previous section as a baseline scenario and we evaluate them against two alternative scenarions. the first alternative scenario is based on the projections of the Institute of English Actuaries for the long-term improvement rates of mortality. The Institute of English Actuaries cosndier a rate of 1.25% ("one and a quarter") as a "base case scenario" for the Long-term improvement rate and an alternative "worst case" scenario of an improvement rate of 3.25% ("three and a quarter"). We consider for our first alternative scenario projections consistent with the hypothesis that the long term mortality improvement rate for the population is 3.25%. We then generate a second alternative scenario by changing the projections for the pattern of mortality for lung cancer, considering an hypothetical situation in which it is projected to reduce progressively to reach zero in 2016.

We use these scenarios for the exogenous variables determining c_{1t} and feed it into our model to stochasticallly simulate mortality rates from 2001

onwards for old ages. We compare in figure 9 projections based on the baseline and on the two alternative scenarios.

Insert Figure 9 here

The results show that the CMI scenario for the total average mortality rates produces lower estimates of mortality at old ages than those based on our baseline scenario and predicts a much stronger "compression of morbidity" of mortality in the next decade. This point is further illustrated in Figure 10 where we report the frequencies of death observed in 2009 along with all the different projections for the frequencies of death in 2020 based on three different scenarios

Insert Figure 10 here

The compression of morbidity delivered by the CMI projections is even stronger than the one projected by the model under the maintained hypothesis that death for lung cancer will disappear by 2016.

4 Conclusions

This paper has proposed a structural framework for the cross-section and the time series of mortality rates, survival probabilities and frequencies of death for a given population. We extend the stochastic mortality model originally proposed by CBD to include information on the limits to life and on the impact of medical progress on the average mortality rate aggregated by causes of death. In this exercise, we have fitted our model to the male population of England and Wales for ages above 65 years old and we projected mortality rates and limits to life up to 2020 by stochastic simulation. Our projections generate very little change in the "limit to life" (age of the oldest person in the population) over the next decade but sizeable shifts in mortality rates and survivor function for ages between 65 and 89 and a further compression of morbidity of mortality. The modal age of death (the age with the highest frequency of death) is predicted to shift from 84 years in 2009 to 90 years in 2020.

References

[1] Blake D. Boardman T. and A. J. Cairns, Sharing Longevity Risk: Why Governments Should Issue Longevity Bonds, Pension Institute Discussion Paper, 2010.

[2] Cairns A. J. G. Blake D. and Dowd K., A Two Factor Model for Stochastic Mortality with Parameter Uncertainty, Journal of Risk and Insurance, 2006.

[3] Cairns A. J. G. Blake D. and K. Dowd, Mortality Density Forecasts: An Analysis of Six Stochastic Mortality Models, Pension Institute Discussion Paper, 2008.

[4] Cairns A. J. G. Blake D. Dowd K., A quantitative comparison of stochastic mortality models using data from England and Wales and the United States, Pension Institute Discussion Paper, 2007.

[5] Canudas-Romo V., The Modal Age at Death and the Shifting Mortality Hypothesis, Demographic Research 19(30): 1179-1204, 2008.

[6] Canudas-Romo V. and Engelman M., Maximum life expectancies: Revisiting the best practice trends, Genus 65(1): 59-79, 2009.

[7] Continuous Mortality Investigations, Working Paper 38: A prototype mortality projections model: Part one - an outline of the proposed approach, CMI Website, 2009.

[8] Continuous Mortality Investigations, Working Paper 39: A prototype mortality projections model: Part two - a detailed analysis, CMI Website, 2009.

[9] Continuous Mortality Investigations, Working Paper 49: The CMI mortality projections model: CMI 2010, CMI Website, 2010.

[10] JP Morgan Pension Advisory Group, Lifemetrics: A toolkit for measuring and managing longevity and mortality risks, Technical Document - Version 1.1, 2007.

[11] Lee R. D. and Carter R. L., Modeling and Forecasting U.S. Mortality, Journal of the American Statistical Association, 1992.

[12] Oullette N., Smoothing mortality data and summarizing recent trends at older ages in low mortality countries, Presented at the Institut Louis Bachelier Conference: Longevity Modeling an Interdisciplinary Approach, 2011.

[13] Oullette N., Changes in age-at-death distribution in low mortality countries: a nonparametric approach, Presented at the XXVI International Population Conference of the International Union for the Scientific Study of Population, 2009.

[14] Renshaw A. E. and Haberman S., A cohort based extension to the Lee-Carter model for mortality reduction factors, Insurance: Mathematics and Economics, 2006.

[15] Robine J. Michel J.P. and Institut S., Has there been a compression of morbidity in countries with low mortality?, Prevention of Functional Dependency :139-148, 2008.

[16] Vallin J. and Mesle F., Will life expectancy increase indefinitely by three months every year?, Population and Societies, 2010.

Table 1: Estimation results, sample 1972-2005				
	$Coeffs$	$StdErrs$	$Tstats$	$p-values$
$\hat{\beta}_1$	0.155594	0.048967	3.177526	0.0019
$\hat{\beta}_2$	0.743237	0.095425	7.788667	0.0000
$\hat{\beta}_3$	0.394905	0.124283	3.177476	0.0019
$\hat{\beta}_4$	0.032356	0.007787	4.154964	0.0001
$\hat{\beta}_5$	0.808094	0.046544	17.36199	0.0000
$\hat{\beta}_6$	-0.000944	0.000268	-3.521158	0.0006
$\hat{\beta}_7$	0.828572	0.046689	17.874662	0.0000
$\hat{\beta}_8$	$2.95E-05$	$9.20E-06$	3.210272	0.0017
$\hat{\beta}_9$	0.778372	0.068254	11.40400	0.0000
	R^2	R^2_{adj}	DW	
$eq1$	0.9821	0.9810	2.4572	
$eq2$	0.6321	0.6206	2.2956	
$eq3$	0.7449	0.7370	2.1185	
$eq4$	0.7108	0.7018	1.7107	

$$c_{1,t+1} = \beta_1 + \beta_2 c_{1,t} + \beta_3 \log \bar{m}_{t+1} + v_{1,t+1}$$
$$c_{2,t+1} = \beta_4 + \beta_5 c_{2,t} + v_{2,t+1}$$
$$c_{3,t+1} = \beta_6 + \beta_7 c_{3,t} + v_{3,t+1}$$
$$c_{4,t+1} = \beta_8 + \beta_9 c_{4,t} + v_{4,t+1}$$
$$\begin{bmatrix} v_{1,t+1} \\ v_{2,t+1} \\ v_{3,t+1} \\ v_{4,t+1} \end{bmatrix} \backsim N \left(\begin{matrix} 0 \\ 0 \\ 0 \\ 0 \end{matrix}, \sum \right)$$
$$\bar{m}_{t+1} = m^h_{t+1} + m^{cd}_{t+1} + m^{lc}_{t+1} + m^{oc}_{t+1} + m^{oth}_{t+1}$$

Figure 1: Observed mortality rates, survival probabilities and frequencies of death for people aged from 65 to 89.

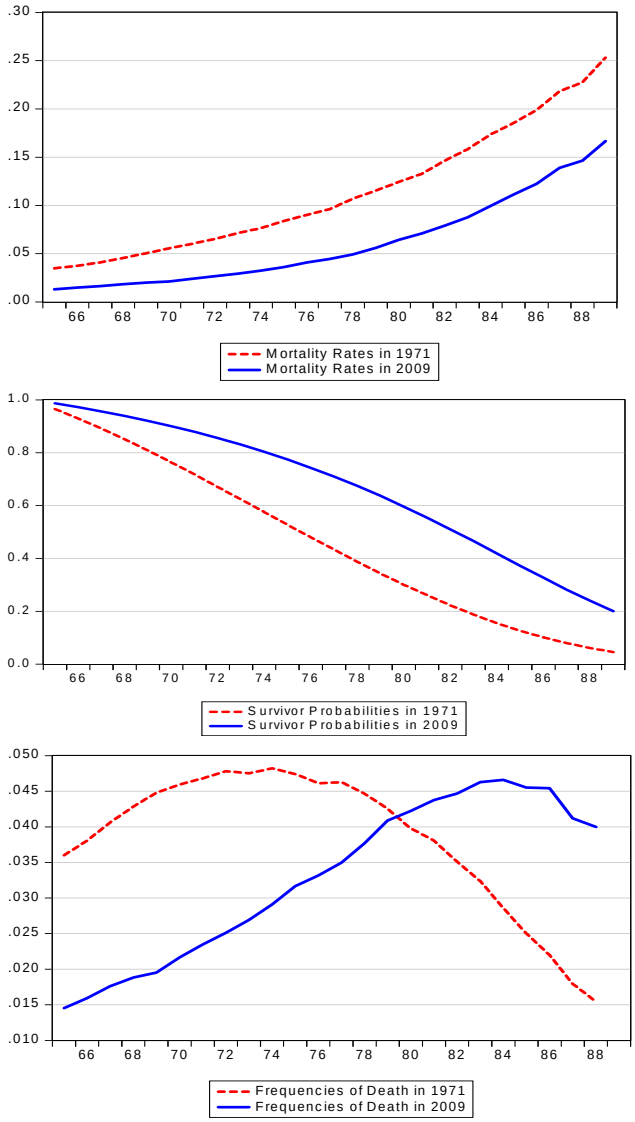

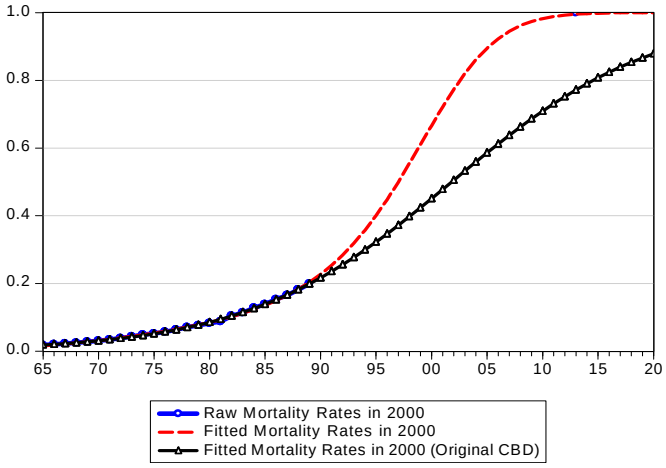

Figure 2: Comparions of the cross-sectional fit of the model vs original CBD for year 2000.

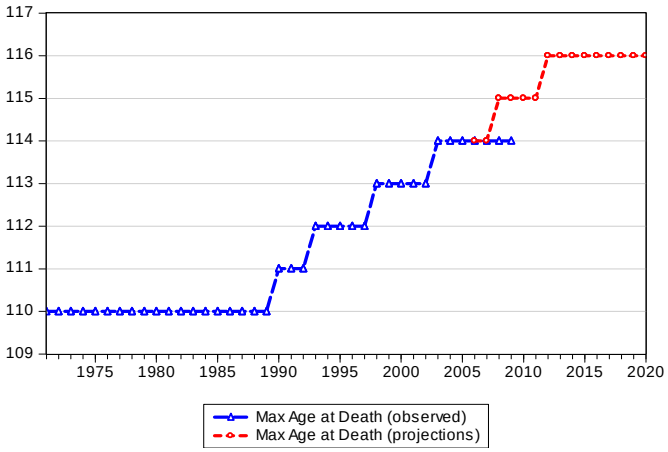

Figure 3: Observed maximum age at death (sample 1971-2009) and out-of-sample projections generated by the model (sample 2006-2020)

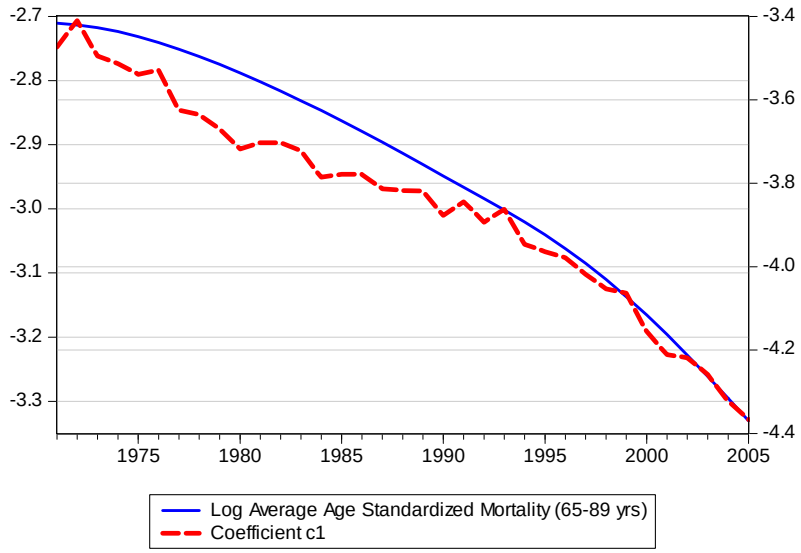

Figure 4: Average age standardized mortality for people aged from 65 to 89 vs coefficient c1 (sample 1971-2005).

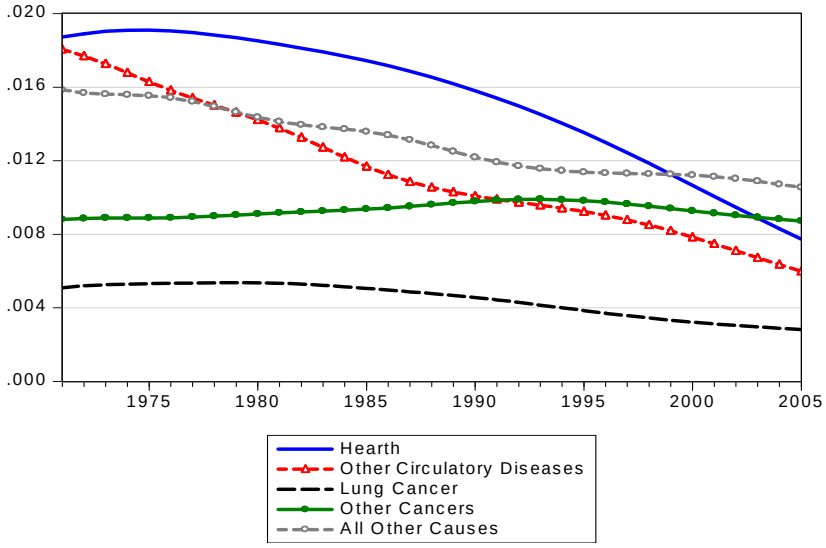

Figure 5: Time series of average age standardized mortality per cause for people aged from 65 to 89 (sample 1971-2005).

Figure 6: Time series of estimated model coeffcients and cross sectional R-squares (sample 1971-2009)

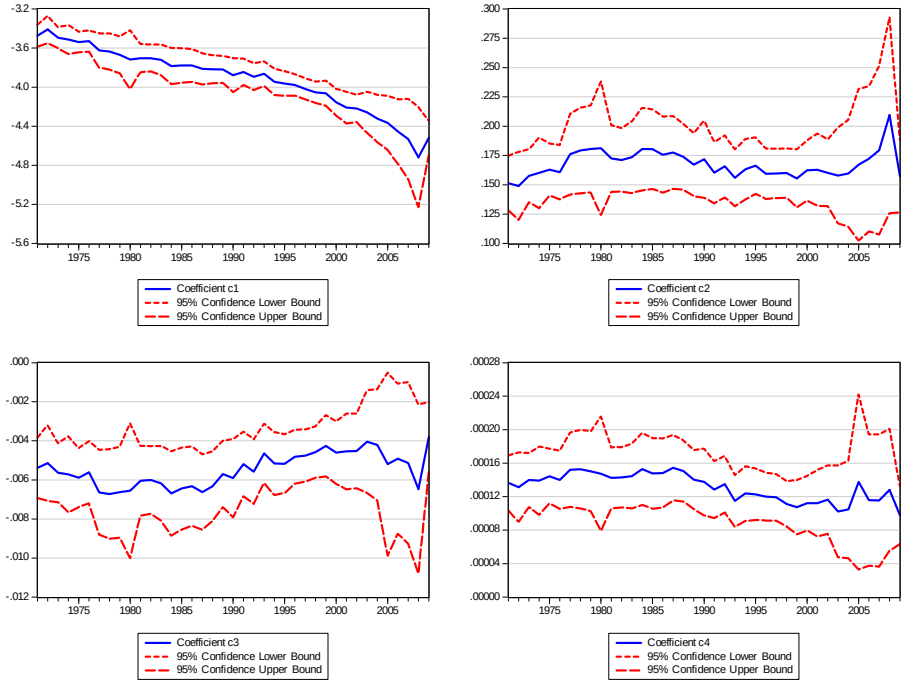

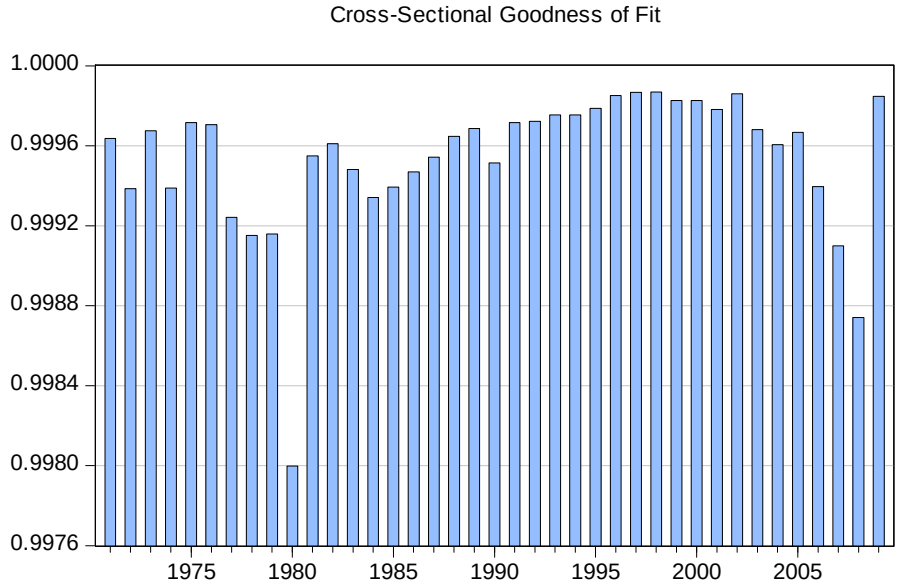

Figure 7: Time series of residuals from from the 4 equations in the model (sample 1971-2005).

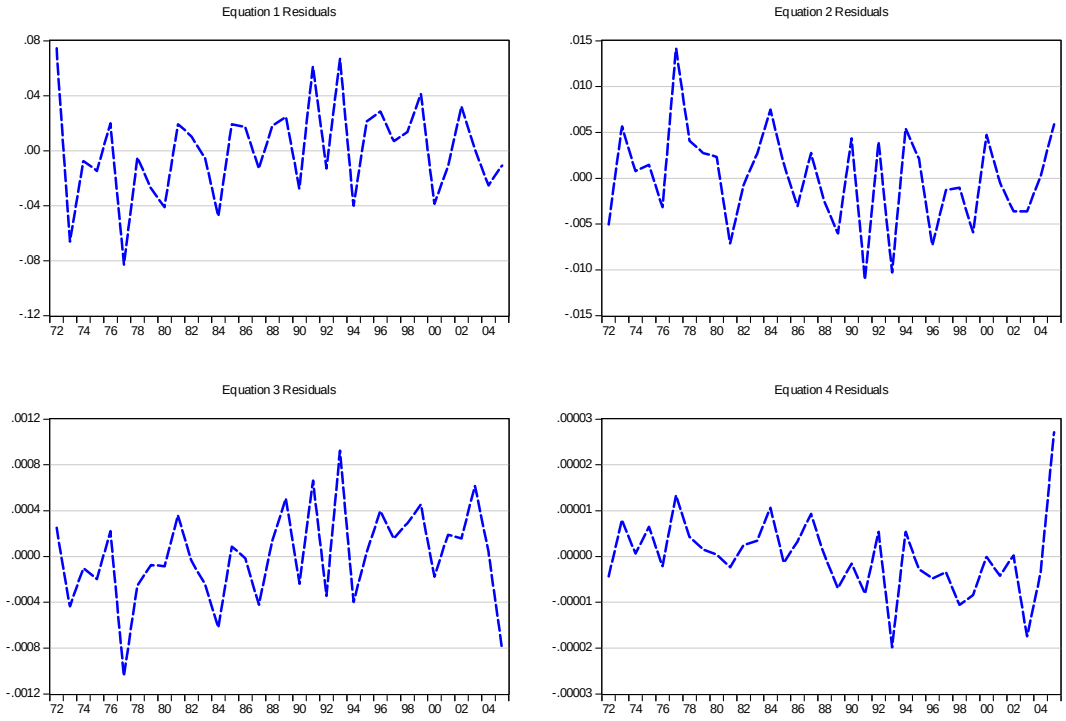

Figure 8: Pseudo out-of-sample predictions for 2009 vs realized values (left column); out-of sample forecast for 2020 (right column)

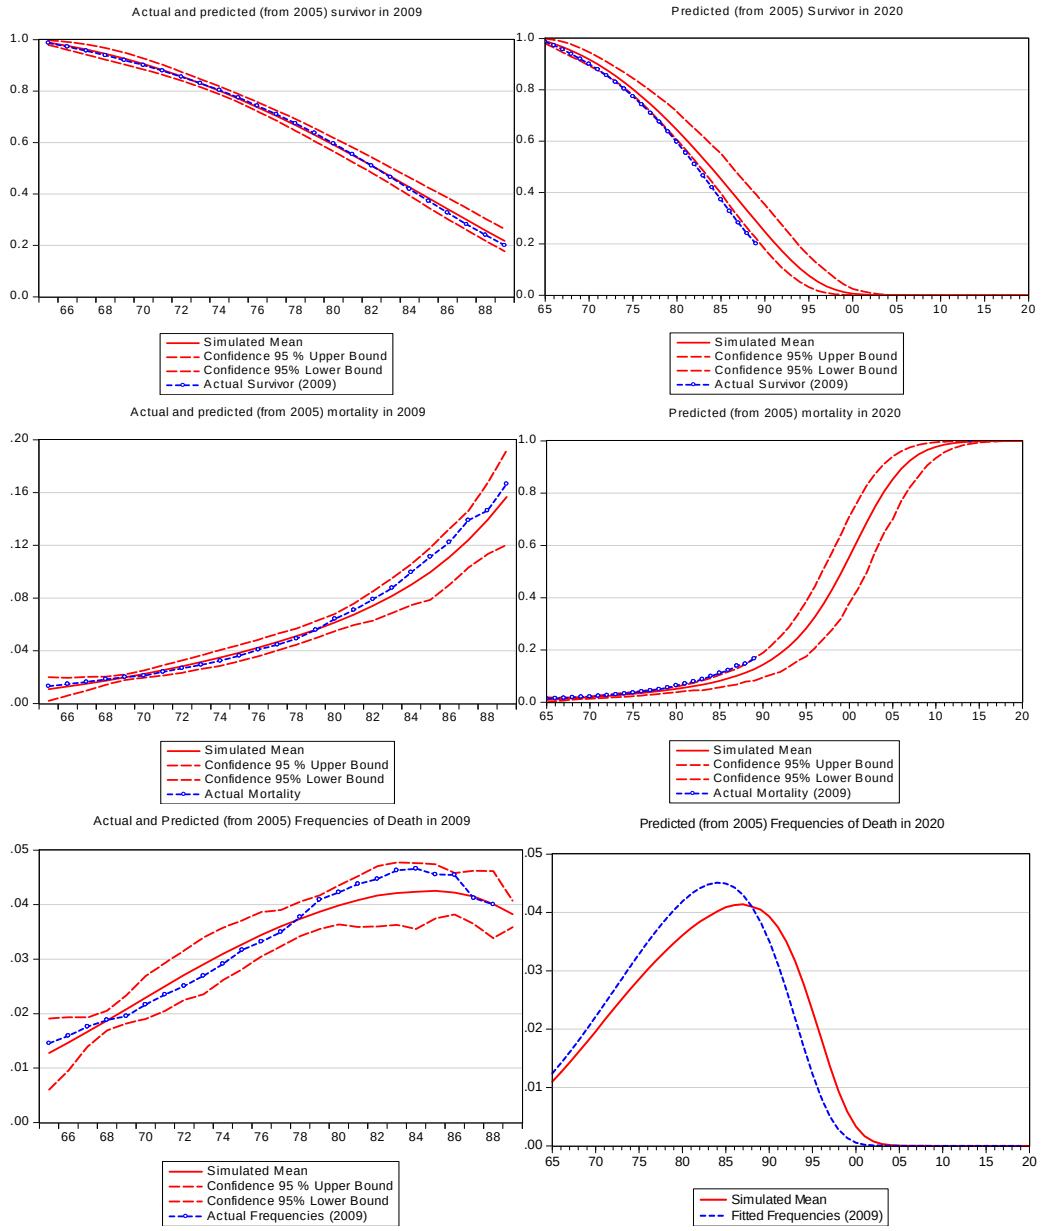

Figure 9: Out-of-sample predictions for 2020 vs realized values in 2009; on the left CMI projections for causes of mortality, on the right the our own hypothesis.

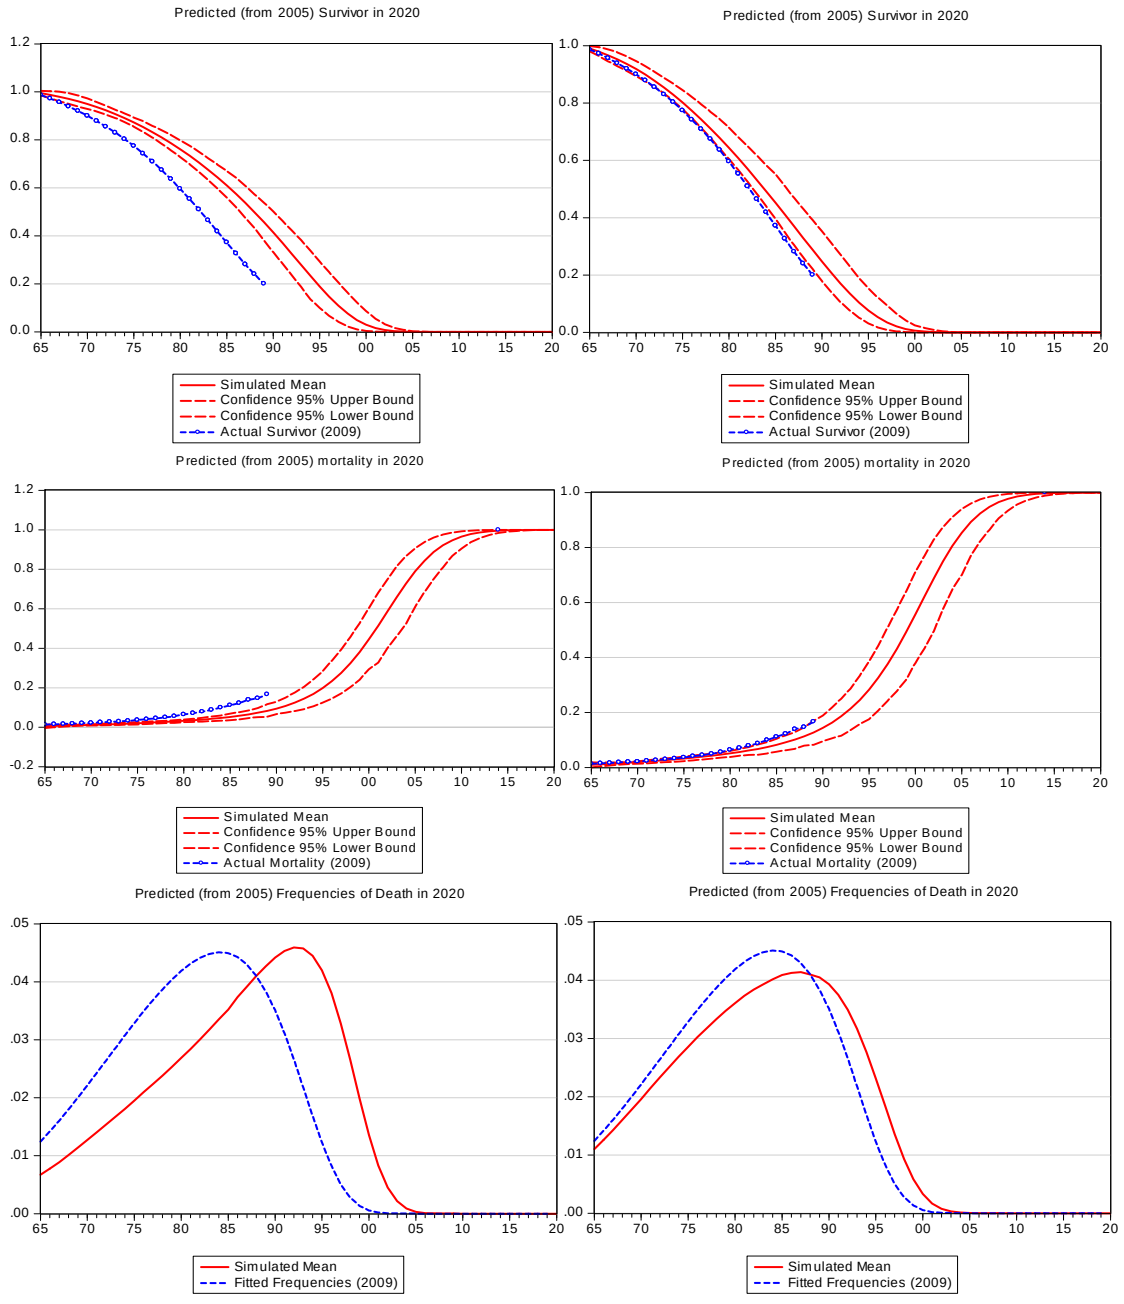

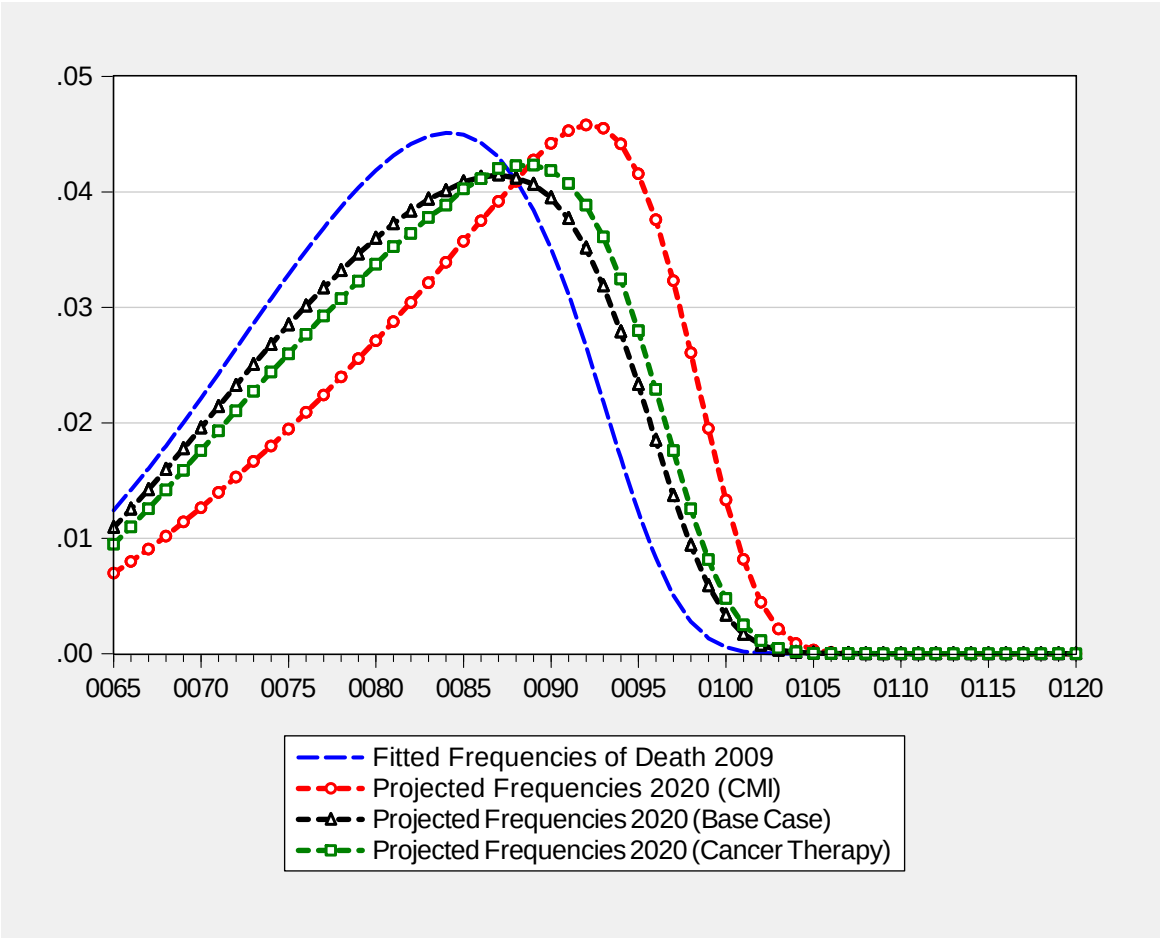

Figure 10: Scenarios for compression of morbidity