Continuous Mortality Investigation
Institute and Faculty of Actuaries
1924-2024 100 years of insight

ISSN 2044-3145

Mortality Projections Committee

WORKING PAPER 111

Methods for the CMI mortality monitor

July 2024

Version 1 of this working paper was originally issued as a consultation in October 2018.
Version 2 (issued in April 2019) described a simpler approach to updating the calculations for new calendar years.
Version 3 is briefer than previous versions and reflects the updated methods used as at July 2024

Summary

This paper describes the data and methods used for the CMI mortality monitors. The broad approach is to:

- Use provisional weekly deaths data from the ONS.
- Standardise the mortality data to allow for changes in the size and age profile of the population.
- Produce various analyses, based on standardised mortality rates (SMRs), to highlight different features of recent mortality.

Contents

1. Introduction 4
2. Dataset 5
3. Standardisation 7
4. Analysis 11

1. Introduction

The CMI has published quarterly "mortality monitors" since October 2018. These contain analysis of recent mortality in the general population in England & Wales, using provisional weekly deaths data published by the Office for National Statistics (ONS).

We currently also publish weekly mortality monitors, containing a subset of the quarterly analysis.

1.1 Overview

The broad approach is to:

- Use provisional weekly deaths data from the ONS – described in Section 2.
- Standardise the deaths data to allow for changes in the size and age profile of the population – Section 3.
- Produce various analyses, based on standardised mortality rates (SMRs), to highlight different features of recent mortality – Section 4.

1.2 Publication schedule

The ONS typically publishes weekly provisional deaths data every Wednesday, although this varies due to public holidays.

Our current approach is to publish:

- weekly analyses on the same day that the ONS publishes its data; and
- quarterly analyses in April (based on ONS data to week 13 of the year), July (to week 26), October (to week 39) and January (to the end of the previous year).

1.3 TAS compliance

This paper is intended to describe the methods that we propose to use to analyse recent mortality.

This paper complies with the principles in the Financial Reporting Council's Technical Actuarial Standard "TAS 100: General Actuarial Standards". Any person using this paper should exercise judgement over its suitability and relevance for their purpose.

1.4 Acknowledgements

The members of the Mortality Projections Committee involved in the production of this paper are Cobus Daneel (Chair), Steve Bale, Mark Cooper, Steven Rimmer, Neil Robjohns, Brian Sewell, and Chris Tavener.

We would also like to acknowledge Tim Gordon and Piero Cocevar, who were MPC members at the time of publication of the original version of Working Paper 111, and Matthew Fletcher who reviewed a draft of that paper.

2. Dataset

This section describes the ONS data that we use, and some key features of it.

2.1 Data sources

The analysis uses provisional weekly deaths data for England & Wales, published by the ONS. We are reliant on the ONS for their timely publication of this data.

- The ONS has published data for weeks ending on a Friday, numbered using the ISO 8601 definition. This means that week 1 of a year ends on the Friday that falls between January 2 and January 8 inclusive.
- Data for a week ending on a Friday is currently typically published on the Wednesday that falls twelve days later, although this can be delayed by public holidays.

In order to have as long a time series as possible we combine data from three ONS sources:

- For week 31 of 1999 to week 53 of 2004, we use a single file of total deaths on an occurrences basis[1], as registered deaths are not available for this period.
- For week 1 of 2005 to week 53 of 2009, we use a single file of total deaths on a registrations basis[2].
- From week 1 of 2010 onwards, we use annual files of deaths on a registrations basis[3]. These contain splits of deaths by gender and age band.

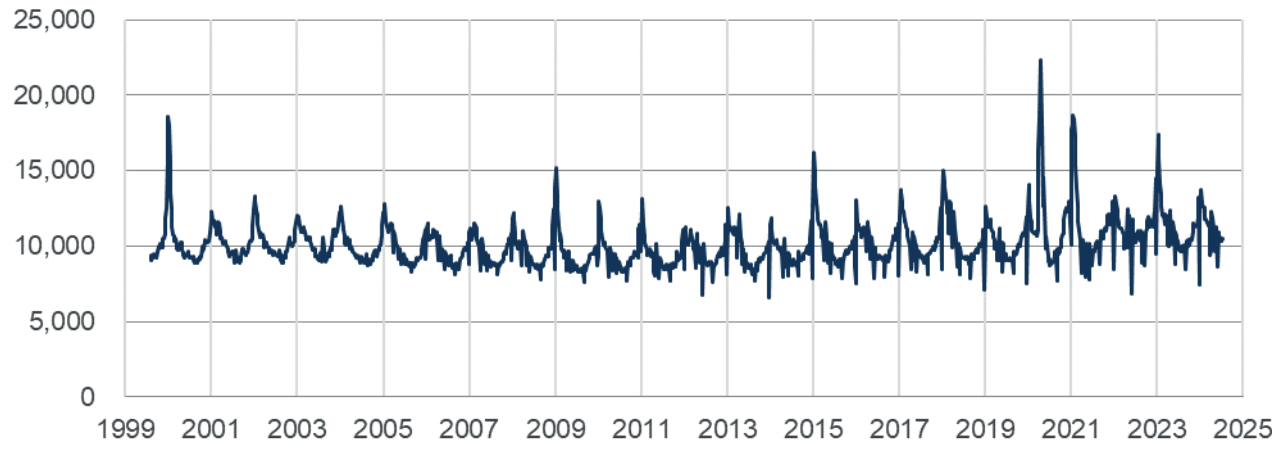
We do not show results for ages below 15 for years before 2010. We consider these to be unreliable as they are inferred from total deaths and our method for assigning deaths to single ages does not work well for these ages.

There is an inconsistency in the analysis due to the data being on an occurrences basis before 2005 and on a registrations basis afterwards. We could have adjusted the data to minimise the impact of this difference, but considered this to be unnecessary: the older data is only used after averaging over a period of at least thirteen weeks.

2.2 Features of ONS deaths data

Chart 2A shows the numbers of deaths in each week, without any adjustments, for the data source described in Section 2.1.

Chart 2A: Weekly deaths, combined data as described in Section 2.1



[1] Source: Figure 4 of http://www.ons.gov.uk/ons/rel/subnational-health2/excess-winter-mortality-in-england-and-wales/2013-14--provisional--and-2012-13--final-/stb.html
[2] Source: http://www.ons.gov.uk/ons/about-ons/business-transparency/freedom-of-information/what-can-i-request/published-ad-hoc-data/pop/may-2014/weekly-provisional-figures-on-deaths-registered-in-england-and-wales-2005-2013.xls
[3] Source: https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/deaths/datasets/weeklyprovisionalfiguresondeathsregisteredinenglandandwales

Change in population

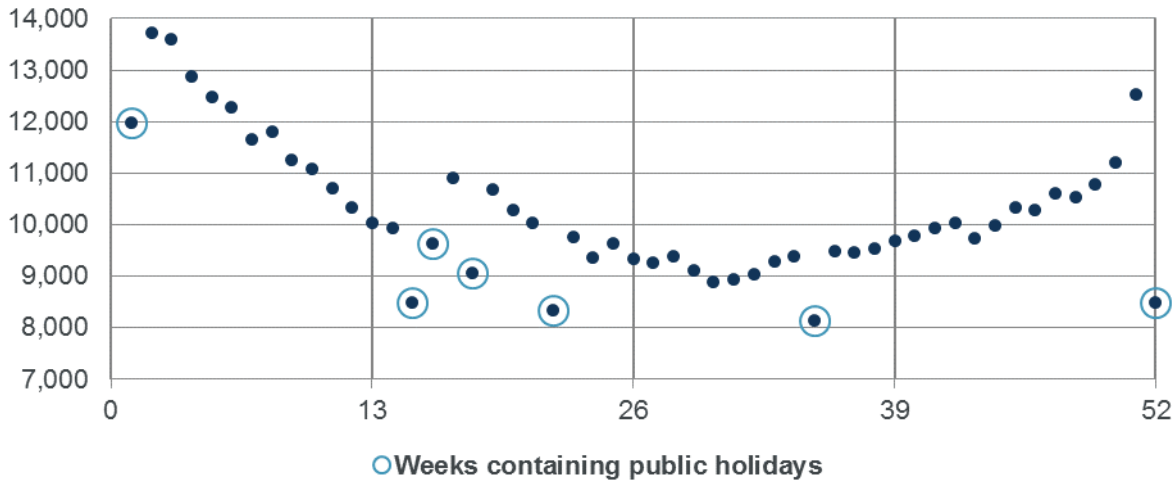
The numbers of deaths shown in Chart 2A are affected by changes in the size and average age of the population over time, as well as mortality rates. In order to focus on mortality rates, we standardise our results to remove the impact of the growing and ageing population. This is described in Section 3.

Volatility

There are several sources of variation in the numbers of deaths from week to week:

- The numbers of deaths in Chart 2A show a strong seasonal pattern, with larger numbers of deaths seen around the start and end of each year.
- The registrations data shows a number of spikes. The downward spikes typically correspond to weeks containing public holidays, when Register Offices are closed so fewer deaths are registered. This is clearly seen in Chart 2B, which shows registered deaths in 2017 and highlights weeks that contain one or more public holidays. Depending on the timing of the holidays within those weeks, there can be a corresponding upward spike when Register Offices reopen.
- There will also be some idiosyncratic variation from week to week. For example, if we expect 10,000 deaths in a particular week, with deaths following a Poisson distribution, then the actual number of deaths would have a standard deviation of 100.
- There will also be some variation linked to specific events such as the COVID-19 pandemic; severe flu-seasons; and abnormally high or low temperatures.

Chart 2B: Weekly death registrations in 2017, with public holidays highlighted



Section 4 describes how we smooth the standardised mortality rates for some charts in order to make the underlying patterns clearer. This smoothing is intended to reduce the impact of idiosyncratic variations and of public holidays.

We considered whether to adjust the data to attempt to reduce the impact of distortions caused by public holidays, but decided not to do so:

- For Charts A to C of the quarterly monitor, the impact of public holidays is greatly reduced by smoothing, so that adjustments seem unnecessary.
- For Charts D to G, we want our estimate of registered deaths in a calendar year to closely match the ONS's annual figure. For this reason, we do not want to make an adjustment for public holidays.

3. Standardisation

Over the period that we consider, the population of England & Wales has grown, and its average age has increased. This would lead to an increased number of deaths even if mortality rates remained constant. To fairly assess changes in mortality rates we need to remove the impact of changes in the population. We do this by calculating Standardised Mortality Rates (SMRs).

There is an established method for calculating an SMR using annual data, but the nature of the available weekly deaths data means that weekly SMRs cannot be calculated in exactly the same way. Consequently, we have devised a method for weekly SMRs which uses the same underlying principles as the annual SMR calculation as far as possible. Although it cannot mathematically be entirely consistent, we expect close alignment in practice.

Section 3.1 gives an overview of the method, and we give further details of its implementation in later sections.

3.1 Overview

We want a method which can calculate a Standardised Mortality Rate (SMR) consistently for various time intervals. We are particularly interested in:

- Calendar years (as in working papers that accompanied releases of the CMI Mortality Projections Model);
- The first part of calendar years (i.e. from 1 January to a later date within the same calendar year) in order to show cumulative mortality through a year; and
- Weeks (ending on a Friday, to coincide with the ONS data).

We use the formula:

$$SMR(G,i) = \sum_{(x,g)\in G}(P(x,g) \times D(x,g,i) \div E(x,g,i)) \div \sum_{(x,g)\in G} P(x,g)$$

for all SMRs, where:

- G is the group (defined by some combination of age and gender) for which we calculate the SMR.
- i is the time interval for which we calculate the SMR.
- x is age.
- g is gender.
- $P(x,g)$ is the standard population for age x and gender g.
- $D(x,g,i)$ is the number of deaths for age x and gender g in time interval i.
- $E(x,g,i)$ is a measure of the exposure for age x and gender g in time interval i. As a proxy we use the estimated population at the midpoint of interval i.

For calendar years this method is identical to our current practice, e.g. in the analysis in Working Paper 189 that accompanied the publication of CMI_2023.

We use the 2013 European Standard Population[4] as our standard population P_0. This is provided in age bands, which we have split into single years of age. We assume that the population within each five-year age band is split equally between its five ages; and that the open age band 95+ is split equally between the six ages from 95 to 100 inclusive. This is the same approach as in Working Paper 189.

We discuss $D(x,g,i)$ and $E(x,g,i)$ further below.

[4] This is available from https://webarchive.nationalarchives.gov.uk/ukgwa/20160106020035/http:/www.ons.gov.uk/ons/guide-method/user-guidance/health-and-life-events/revised-european-standard-population-2013--2013-esp-/index.html.

3.2 Determining deaths, $D(x, g, i)$

We first consider periods, such as the current year, for which we need to make an estimate based on age-grouped provisional weekly deaths data.

A. Start with provisional weekly deaths data by band

Write $D^{(0)}(b, w, y)$ for the number of deaths for band b in ISO week w of ISO year y, from the ONS provisional weekly deaths data.

Note that before 2010 there is only one band (all ages and both genders combined), but for 2010 onwards the bands incorporate gender and age (i.e. we have deaths data for males and females separately, and split by age bands <1, 1-14, 15-44, 45-64, 65-74, 75-84 and 85+). For 2009 and earlier we have a single band containing all ages and both genders.

B. Assign the deaths to individual days

We assign deaths to individual days so that we can eventually calculate the numbers of deaths in any part of any year. We do this step first, before assigning deaths to single ages, to be consistent with the method for constructing the CMI Model calibration data, which first splits weeks that span multiple years.

If a week lies entirely within one calendar year, then we assign one-seventh of the deaths for that week to each of the days within it.

If a week spans two calendar years, then we first apportion the deaths between the years according to the number of working days in each year. This is consistent with the approach used for the CMI_2022 calibration dataset (for CMI_2023, we were provided with provisional annual deaths by the ONS). Then within each part-week within each calendar year, we divide the deaths equally according to the number of days (all days, not working days) within each part-week.

If week w spans years y_1 and y_2 then write n_1 and n_2 for the numbers of days of that week in each year, and w_1 and w_2 for the numbers of working days of that week in each year. Then if there are d deaths in the week we assign $\frac{dw_1}{n_1(w_1+w_2)}$ deaths to the days in year y_1 and $\frac{dw_2}{n_2(w_1+w_2)}$ deaths to the days in year y_2.

For example:

- The week ended on 4 January 2013 has 3,164 deaths for females aged 85+.
- That week has three days in 2012, one of which is a working day; and four days in 2013, three of which are working days.
- We split the 3,164 deaths between 2012 and 2013 in the ratio of the numbers of working days; i.e. 791 in 2012 and 2,373 in 2013.
- We further the split the deaths by days, so for females aged 85+ in that week we have 263.7 deaths on each of 29-31 December 2012, and 593.3 deaths on each of 1-4 January 2013.

In the final step we could have instead assigned deaths to individual days based on working days, so that no deaths are assigned to weekends or public holidays. However, this would give a “bumpier” picture. The approach uses working days where necessary, to apportion deaths between years in order to be consistent with the CMI Model method, but we need not use working days to apportion deaths within a year.

Note that the method can mean that some days have no deaths assigned (e.g. if 31 December is a Saturday then it will have no deaths assigned as deaths for the week ending on Friday 6 January will all be assigned to the following year).

At this stage we have $D^{(1)}(b, d, y)$ for the deaths for band b in day d of calendar year y.

C. Split into single years of age

For each day, we now split the numbers of deaths in each age band into deaths by gender and single year of age. This is done using the same method as for the CMI Model when assigning deaths in the final year of the calibration dataset to single years of age.

We want to calculate:

$D^{(2)}(x,g,d,y)$ Deaths for age x and gender g in day d of calendar year y

We do this by splitting $D^{(1)}(b,d,y)$ based on a "standard" number of deaths $D^{(Std)}(x,g,y)$:

$$D^{(2)}(x,g,d,y) = D^{(1)}(b,d,y) \times \frac{D^{(Std)}(x,g,y)}{\sum_{(x,g)\in b} D^{(Std)}(x,g,y)}$$

The standard number of deaths is based on deaths in the prior year with an allowance for change in population.

We write:

$D^{(Act)}(x,g,y)$ actual deaths for year y by age and gender; and
$E(x,g,y)$ mid-year population for year y by age and gender

then calculate the standard number of deaths as:

$$D^{(Std)}(x,g,y) = D^{(Act)}(x-1,g,y-1) \times \frac{E(x-1,g,y-1) - \frac{1}{2}D^{(Act)}(x,g,y-1) - \frac{1}{2}D^{(Act)}(x-1,g,y-1)}{E(x-1,g,y-1)}$$

D. Combine deaths within the interval

The final step is to combine daily deaths $D^{(2)}(x,g,d,y)$ into deaths for the time interval, $D(x,g,i)$

To do this, we simply add up the deaths for the days that fall within the interval:

$$D(x,g,i) = \sum_{(d,y)\in i} D^{(2)}(x,g,d,y)$$

3.3 Determining $E(x,g,i)$

We use the population at the midpoint of interval i as a proxy for the exposure. The population at the midpoint is estimated by linear interpolation between mid-year populations $E^{(A)}(x,g,y)$.

For example, for the week ended on 1 June 2018:

- the mid-point of the week is 29 May 2018;
- we interpolate between mid-year populations in 2017 and 2018, deemed to be at 1 July in each year;
- the interpolated population is $9.04\% \times E^{(A)}(x,g,2017) + 90.96\% \times E^{(A)}(x,g,2018)$

In order to calculate $E(x,g,i)$ for the current year we need to estimate the population for the two most recent years. We use the same method as used for the final year of data in the CMI Model.

3.4 Leap years

For Charts D to G of the quarterly monitor, which compare cumulative mortality within several calendar years, we want to treat leap years and non-leap years consistently. For non-leap years we simply plot the cumulative value against the day of the year. For leap years we plot values for day d of the leap year against day $\frac{365}{366}d$ of non-leap years. This means that the end-year values are compared consistently.

3.5 Updating calculations

The Committee will restate rates from time to time to reflect revised population estimates.

When the CMI Model is published, typically towards the start of each year, we will revise our population estimates and make a new estimate for the following year. For example:

- CMI_2023 was published in April 2024.
- It uses ONS mid-year population estimates up to 2022 and a CMI population estimate for mid-2023 (the final year of the calibration dataset).
- We use a consistent method to produce a population estimate for mid-2024 and mid-2025, assuming mortality rates in 2024 and 2025 are the same as for 2023 in order to estimate deaths in 2024 and 2024 based on the corresponding population estimates.
  - We note there is a circularity here in estimating deaths using estimated populations, which in turn depend on the estimates of deaths. We resolve this by using iterative calculations.
- The mortality monitors for week 26 and Q2 of 2024, published in July 2024, are the first to reflect the new population estimates made as part of CMI_2023.

4. Analysis

The analyses in the quarterly mortality update are of two types:

- Centred averages of smoothed SMRs over time; described in Section 4.1.
- Analysis of cumulative (unsmoothed) SMRs for specific years, described in Section 4.2.

4.1 Centred average SMR

We calculate centred averages of SMRs in two ways:

- The "quarterly" average for a particular week is the average of the weekly SMRs for the 13 weeks centred on that week; i.e. that week, 6 weeks before and 6 weeks after.
- The "annual" average for a particular week is the average of the weekly SMRs for the 53 weeks centred on that week; i.e. that week, 26 weeks before and 26 weeks after.

The annual average removes seasonal effects and effectively shows an annual SMR, but at weekly intervals. The quarterly average removes short-term variations but still shows seasonal patterns, e.g. allowing the identification of winters with particularly heavy or light mortality.

The quarterly mortality update shows this information in three ways:

- Chart A shows quarterly and annual centred average SMRs for the whole of the period for which we have data, to provide historical context.
- Chart B "zooms in" to the lower-right corner of Chart A to more clearly show the past five years and the current year.
- Chart C shows the same information as Chart B but plots it by ISO week number so that the years are superimposed. This enables a more direct comparison between the current year and the recent past.

4.2 Cumulative mortality by year

Charts described in the previous section show smoothed mortality at a point in time. However, because the CMI Model is calibrated to annual mortality data, it is instructive to consider cumulative weekly deaths during each year.

When publishing the results of the CMI Model each year, we typically show how results from the next version of the CMI Model might vary, assuming no change in method, for various levels of mortality improvements in the final year of the new calibration dataset.

Our key aim in this section is to produce a measure that we can track over the course of the year that will agree closely at the end of the year with the actual mortality improvement. This can inform a view, during the year, of the possible level of mortality improvement, and hence results from the next version of the CMI Model.

The method for deaths in Section 3.2 means that the emerging cumulative standardised mortality for the current year should closely match the estimate eventually used in calibrating the CMI Model.

Annual mortality improvement

When considering possible outcomes from the next version of the CMI Model, we use an "m-style" annual mortality improvement; i.e. the improvement for year y is:

$$MI_y = \log(m_{y-1}) - \log(m_y)$$

This is equivalent to:

$$MI_y = \log\left(\frac{m_{y-1}}{m_y}\right) = -\log\left(\frac{m_y}{m_{y-1}}\right) = -\log\left(1 + \left(\frac{m_y}{m_{y-1}} - 1\right)\right)$$

Since the difference between m_y and m_{y-1} should be small, a Taylor expansion lets us approximate this closely by $MI_y \approx 1 - \frac{m_y}{m_{y-1}}$.

Relative cumulative mortality

In this section we consider cumulative mortality relative to the average of the previous ten years; i.e.

$$R^*(d,y) = \frac{m^*(d,y) - \overline{m^*}(d)}{\overline{m^*}(365)}$$

where

- $m^*(d,y)$ represents cumulative mortality over the first d days of year y (after scaling leap years as described in Section 3.4); and
- $\overline{m^*}(d)$ is the average of $m^*(d,y)$ for the ten most recent complete calendar years (e.g. 2014-2023 currently).

This is shown in Chart D of the quarterly mortality update.

Cumulative standardised annual mortality improvement

We also define a cumulative standardised annual mortality improvement measure:

$$CI^*(d,y) = \frac{m^*(d,y-1) - m^*(d,y)}{m^*(365,y-1)}$$

Note that $CI^*(0,y) = 0$ and

$$CI^*(365,y) = \frac{m^*(365,y-1) - m^*(365,y)}{m^*(365,y-1)} = 1 - \frac{m^*(365,y)}{m^*(365,y-1)}$$

which corresponds to the annual mortality improvement over the whole of year y. So $CI^*(d,y)$ starts at zero and by the end of the year it corresponds to the annual mortality improvement.

Chart E of the quarterly update shows the cumulative annual standardised mortality improvement for the current year and the previous ten years.

Note that:

- Chart E shows cumulative improvements, so a higher value represents a higher improvement and lower mortality; whereas in Chart D a higher value represents higher mortality.
- Lines on Chart E show a number of short-term "bumps" corresponding to public holidays, so readers are advised not to read too much into week-to-week changes, but to concentrate on longer-term patterns.
- For leap years we plot values for day d of the leap year against day 365/366 d of non-leap years.

Charts F and G of the update are similar to Charts D and E respectively, but show variations by gender and age band.

4.3 Period covered

Writing Y for the current year (e.g. $Y = 2024$ now):

- Chart A shows results for the whole of the period for which we have data.
- Charts B and C show results for years $Y - 5$ to Y.
- Charts D to G show results for years $Y - 10$ to Y.

4.4 Weekly monitor

Chart 1 of the weekly monitor shows weekly SMRs for 2022 onwards, and the 2011-2019 range.

Charts 2 and 3 of the weekly monitor are the same as Charts D and E of the quarterly monitor.

Charts 4 and 5 of the weekly monitor show counts of deaths for specific causes, without any age-standardisation.

Disclaimer: This Document has been prepared by and/or on behalf of Continuous Mortality Investigation Limited (CMI). The CMI does not accept any responsibility and/or liability whatsoever for the content or use of this Document. Whilst care has been taken during the development of the Document, CMI does not (i) warrant its accuracy; or (ii) guarantee any outcome or result from the application of this Document or of any of CMI's work (whether contained in or arising from the application of this Document or otherwise). You assume sole responsibility for your use of this Document, and for any and all conclusions drawn from its use. CMI hereby excludes all warranties, representations, conditions and all other terms of any kind whatsoever implied by statute or common law in relation to this Document, to the fullest extent permitted by applicable law. If you are in any doubt as to using anything produced by CMI, please seek independent advice.

Copyright: You may reproduce the contents of this Document free of charge in any format or medium provided it is:

1. reproduced accurately and is unaltered;
2. not used in a misleading context; and
3. correctly referenced and includes both CMI's Disclaimer notice set out above and CMI's copyright notice, as follows:

© Continuous Mortality Investigation Limited.

Continuous Mortality Investigation Limited ("CMI") is registered in England & Wales
Company number: 8373631
Registered Office: 1-3 Staple Inn Hall, High Holborn, London, WC1V 7QJ

Correspondence address: Two London Wall Place, 123 London Wall, London, EC2Y 5AU
Email: info@cmilimited.co.uk
Tel: 020 7776 3820

Website: www.cmilimited.co.uk (redirects to www.actuaries.org.uk)

Continuous Mortality Investigation Limited is wholly owned by the Institute and Faculty of Actuaries.