Institute
and Faculty
of Actuaries

SESSIONAL MEETING DISCUSSION

"Allowing for shocks in portfolio mortality models" by Stephen Richards

[Institute and Faculty of Actuaries, Sessional Event, Wednesday 15 December 2021]

This discussion relates to the paper by Stephen J. Richards presented at the IFoA Sessional event held on 15 December 2021.

The Moderator (Prof. A. S. Macdonald, F.F.A.): Good morning and welcome to this sessional meeting. My name is Angus Macdonald, and I will be chairing the meeting. Our speaker this morning is Dr Stephen Richards who is the Managing Director of Longevitas Limited, a specialist provider of actuarial tools for longevity risk and annuities. Longevitas was founded in 2006 and has users in the UK, USA, Canada and Switzerland. Before founding Longevitas, Stephen headed the longevity analysis team at Prudential and prior to that he headed the product pricing team at Standard Life. He is an Honorary Research Fellow at Heriot-Watt University and he publishes regularly in the actuarial research literature, addressing practical longevity issues.

Dr S. J. Richards, F.F.A.: Thank you Angus, and welcome everyone to this presentation on the paper "Allowing for shocks in portfolio mortality models." The motivation for this is the ongoing COVID-19 pandemic, which has led to some sharp spikes in mortality. The question for actuaries is how does COVID-19 mortality affect portfolio mortality experience, and in particular how to avoid bias when deriving mortality bases from actual portfolio experience? This bias is a business problem for annuity portfolios and pension schemes. If you are setting a long-term basis for reserving, it would be imprudent to include recent excess shock mortality when deriving that basis. For bulk annuities and longevity swaps, it would also be inappropriate to include excess shock mortality when pricing any kind of risk transfer.

We consider some aspects of the COVID shocks in the UK and elsewhere. We will look at a semi-parametric approach to mortality shocks, then we will take what seems like a detour into the topic of reporting delays. (It will become clearer later in the presentation why reporting delays are also relevant for the methodologies presented.) We will then move to a fully parametric model of portfolio mortality, and we will look at the rather different requirements for modelling mortality by age on one hand and by calendar time on the other. Then at the end we will have a brief look at what can be done for portfolio-specific mortality improvements.

We begin with the COVID-19 shocks. Figure 1 shows the all-cause deaths in England and Wales at ages 65 and over, from January 2019, and we can see the two well-known COVID shocks in April 2020 and January 2021.

Mortality levels were fairly normal until 20 March 2020, with 9000 weekly deaths. Just four weeks later, by 17 April, the number of weekly deaths had more than doubled. So, we had a very intense short term mortality shock unfolding over a matter of weeks.

Figure 2 shows the deaths in the UK where COVID-19 was listed as one of the causes of death, and we can see the two sharp peaks clearly.

These peaks look roughly equal in height, reaching nearly 1,500 deaths per day due to COVID-19. It is worth contrasting Figure 1 with Figure 2, as the first spike in Figure 1 was
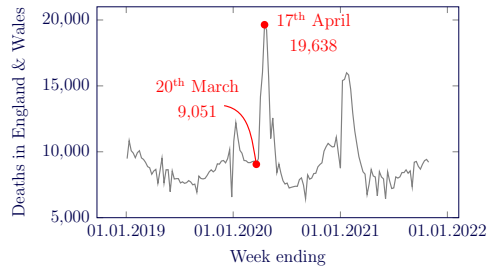
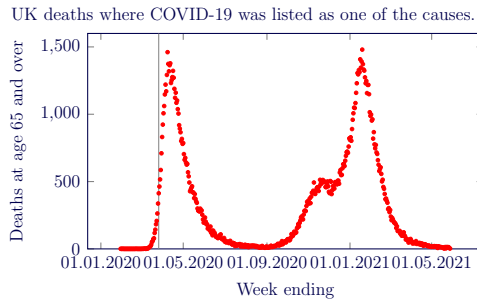**Figure 1.** All-cause deaths aged 65+ (accreditation underneath as normal).

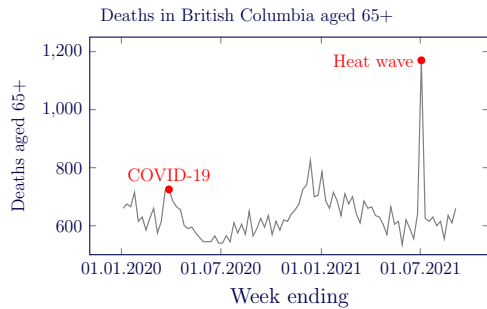**Figure 2.** COVID-19 deaths.

Figure 3. Deaths in British Columbia at ages 65 and over. Source: Statistics Canada

considerably higher than the second. One reason for this is that the coding for causes of death during the first shock was incomplete, and that not all the deaths in the first shock have been correctly coded as being due to COVID-19.

In fact, there is more to mortality shocks than just COVID-19. If we look at deaths in British Columbia in Figure 3, the COVID mortality is noticeable, but not particularly dramatic. What is quite spectacular, however, are the deaths that occurred in the week of the heat dome event in the summer of 2021 in Northwestern America.

Figure 3 shows that mortality shocks have causes other than just COVID, so there is possibly an application for the methods that we are going to look at today for climate-stress mortality. In both cases shocks occur over a very short space of time. On that basis annual mortality rates cannot hope to capture the true nature of mortality shocks, and this suggests that we need continuous-time methodologies, rather than methodologies based on annual q_x rates.

We next consider some "non-solutions" for actuaries managing portfolios. One idea might be to build some kind of cause-of-death model, not least because the COVID-19 shocks obviously entail a very specific cause of death. However, pension schemes do not record the cause of death in their administration systems, and neither do most insurers. So, cause-of-death data are usually unavailable for portfolios managed by actuaries. Also, even if cause-of-death data have been recorded, we saw in comparing Figures 1 and 2 that the coding was not wholly reliable.

Another option might be to simply ignore the periods affected by mortality shocks. However, pension schemes often only have data for the last one or two triennial valuations, so it would be wasteful to throw away some of that limited data. If a model is incapable of handling some fundamental real-world features of the data, then we need a better model. We should not be trimming data to fit the model but finding a model that fits the data. Our requirements are that we need methodologies that work with the data that is available in the administration systems. We need methodologies that are going to work with the entire exposure period and do not demand that we throw away data. And we need methodologies that can handle sharp spikes in mortality rates.

We will look at a semi-parametric approach to this, starting with some definitions:

y	Start of period of interest.
$y + t_i$	A unique date with at least one death.
$d_y + t_i$	Number of deaths at $y + t_i$.
$l_{y+t_i^-}$	Number of lives immediately before $y + t_i$.

Once a portfolio of pensioners has around 15,000 lives or more, there is usually at least one death every day. We use following version of the Nelson–Aalen estimator to estimate mortality in time:

$$\hat{\Lambda}_{y,t} = \sum_{t_i \leq t} \frac{d_{y+t_i}}{l_{y+t_i^-}}$$

The Nelson–Aalen estimator is usually defined with respect to age, this being the most important risk factor, but here we define $\hat{\Lambda}_{y,t}$ as a function of calendar time. Nelson–Aalen estimators are simply the sum of the crude daily mortality rates on the days that at least once death occurs.

$\hat{\Lambda}_{y,t}$ estimates the integrated hazard over calendar time, which we illustrate with three international annuity portfolios:

- Top-up annuities for academics in France.
- Annuities from defined-contribution pensions in the UK.
- Buy-out annuities in the United States.

If we look at the Nelson–Aalen estimate over time for the French portfolio in the left panel of Figure C, it resembles a fairly straight line and appears relatively uninteresting. To get more detail, we calculate the first central difference for the Nelson–Aalen estimator around a point in time:

$$\hat{\mu}_{y+t} = \frac{(\hat{\Lambda}_{y,t+\frac{c}{2}} - \hat{\Lambda}_{y,t-\frac{c}{2}})}{c}$$

where $c > 0$ is the bandwidth parameter. $\hat{\mu}_{y+t}$ here is just the portfolio hazard estimate over time without any risk factors. The right panel of Figure 4 shows the clear seasonal patterns in mortality. We have winter peaks, usually in January each year, and we also have summer troughs, usually around July.

Here we see varying peaks. Trough-to-peak variation is also more pronounced in some years than in others. Using a shorter period and a smaller bandwidth parameter of $c = 0.2$ Figure 6 gives the first indication of the COVID-19 spike in April 2020. We can also see some other interesting features. The winter mortality of 2016/2017 was quite severe for this portfolio, and the COVID-19 spike is not much higher than the worst recent winter for the French annuitants.

It is also worth noting for later that we have a sharp fall in mortality on the right in Figure 6. That is not due to the mortality shock taking frailer lives and leaving a more robust population of survivors. Rather, it is due to delays in the reporting of deaths.
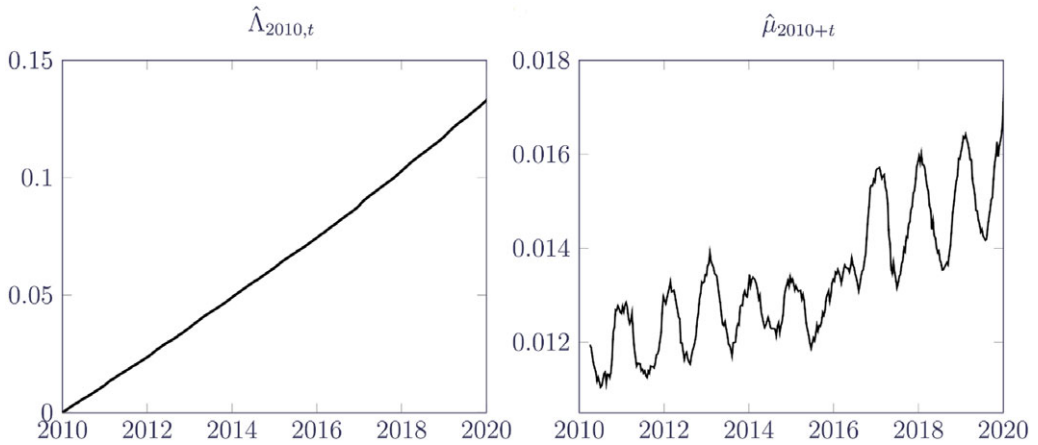
Figure 4. Semi-parametric estimates of mortality in time for French top-up annuities. Left panel: Nelson–Aalen estimate, $\hat{\Lambda}_{2010,t}$. Right panel: estimate $\hat{\mu}_{2020+t}$ with $c = 0.5$.
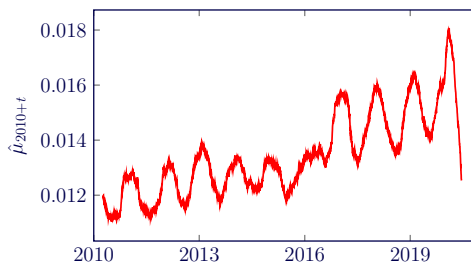Source: Richards (2022b, Figure 2)

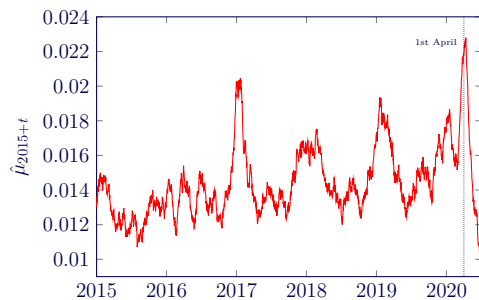**Figure 5.** FRA, $\hat{\mu}_{2015+t}$, $c = 0.5$.

Figure 6. Mortality estimate $\hat{\mu}_{2015+t}$ for French annuities $c = 0.2$.

We have seen that the Nelson–Aalen estimator over time might appear featureless at first glance, but if we take first differences then μ reveals some richly detailed patterns in time. In fact, reducing our smoothing parameter will allow us to see the extent of the COVID-19 shocks in these three portfolios.

If we reduce the smoothing parameter to 0.2 for the French portfolio, we can see the spike is the COVID mortality. Following the winter spike in January 2020 there was then the COVID spike in April. We can also see some other interesting features. The winter mortality of December and January 2017 was quite severe for this portfolio, and the COVID spike is not much higher than the worst recent winter for the French annuitants.

The picture is rather different for the UK annuitants, as shown in Figure 7:
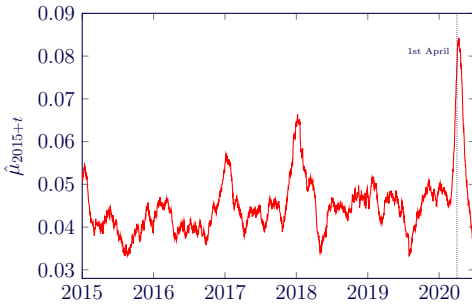
Figure 7. Mortality estimate $\hat{\mu}_{2015+t}$ for UK annuities, $c = 0.2$.
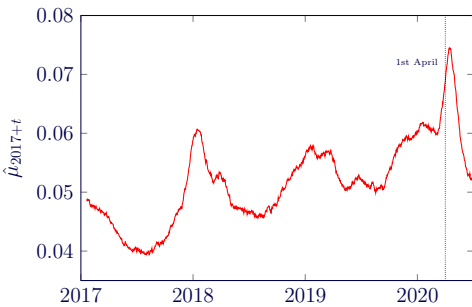

Figure 8. Mortality estimate $\hat{\mu}_{2015+t}$ for US annuities, $c = 0.2$.

The UK annuitants had a bad winter in 2017/2018. However, the COVID-19 spike for the UK annuitants is considerably higher than the most recent bad winter. We can also see that the winter of 2019–2020 was developing to be fairly mild but was then followed by a rather dramatic shock.

For the US portfolio, we only have data spanning a three-year period. It is a much larger portfolio, so the lines in Figure 8 are smoother than the French and UK portfolios for a given value of c.

As before, we have pronounced seasonal variation, and the COVID shock in April 2020 as expected. Using this semi-parametric estimator, we find that the first COVID-19 shock hit annuity portfolios in France, the UK and US at much the same time, peaking in early April 2020.

It is worth noting that $\hat{\Lambda}_{y,t}$ and $\hat{\mu}_{y+t}$ are useful from a data-privacy point of view because they only need (i) the date that the annuity commences, (ii) the date the annuity ceases, and (iii) the nature of cessation (i.e. whether it is a death, a withdrawal or some kind of censoring event). This means that the semi-parametric approach requires no personal data, so GDPR does not apply nor do PIPEDA or (say) the Californian data protection rules. These non-parametric approaches are privacy safe.

This semi-parametric approach has a number of advantages. It reveals seasonal variation, the timing and extent of mortality shocks, and it requires no personal data. In fact, it is sufficiently simple that it can be implemented in spreadsheets or in R; Richards (2022b, Appendix C) gives R code for calculating $\hat{\Lambda}_{y,t}$ and $\hat{\mu}_{y+t}$.

However, the semi-parametric approach has some important disadvantages. First and foremost, it ignores key risk factors like age, so it cannot be used for pricing (although it can be used for exploratory data work and is very effective for communicating with non-actuarial specialists). It is also not defined for the most recent $c/2$ years and is affected by reporting delays.

All pension schemes and annuity portfolios experience delays in the reporting of the death of an annuitant or pensioner to the administrator. To illustrate, we consider two extracts at calendar times, the first at u_1 and the second at u_2. We calculate the ratio, $R(s, u_1, u_2)$, of the non-parametric estimates using the two extracts:
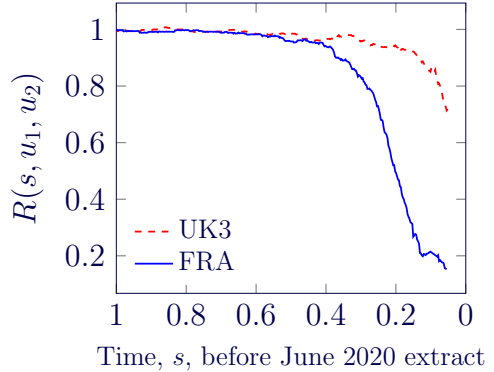
Figure 9. $R(s, u_1, u_2)$ for French and UK annuitants with reversed horizontal axis.
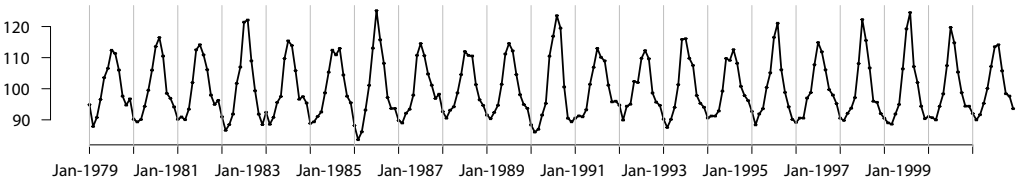Source: Richards (2022b, Figure 7)

Figure 10. Percentage of average daily number of deaths in Australia, all causes.
Source: de Looper (2002)

$$R(s, u_1, s_2) = \frac{\hat{\mu}_{u_1 - s} \text{using extract at time } u_1}{\hat{\mu}_{u_1 - s} \text{ using extract at time } u_2}$$

and this will give us an indication of the impact of delays in mortality reporting. In our example u_1 is June 2020 and u_2 is September 2020. We calculate the ratio of the non-parametric estimates using the two extracts, where the later extract will be less affected by reporting delays for the time leading up to time u_1. Figure 9 shows the ratio for the French and UK annuitants with the horizontal axis reversed to show the time before the extract at u_1.

Figure 9 shows that the further we go back in time before the extract in June 2020, the less impact there is for unreported deaths, so the ratio is close to 1. The ratio falls to zero as we approach the extract date. The French portfolio is subject to considerably more pronounced delays in death reporting. This could be partly driven by differences in insurer administration practices. Many UK insurers, for example, receive regular data feeds of death registrations from the General Registrar's Office and these are matched with annuitants to detect unreported deaths.

There can be pronounced differences in mortality reporting delays. These reporting delays are another source of potential bias, this time working in the opposite direction to mortality shocks. In actuarial work we normally ignore the exposure period leading up to the extract data, as we do not want an analysis biased by the unreported deaths (Macdonald et al, 2018, page 38). Indeed, a rule of thumb from Figure 9 might be to ignore six months prior to the date of extract. However, it is generally not a good idea to throw away data if it can be avoided, so we will revisit this topic.

We finally arrive at a solution that addresses both mortality shocks and unreported deaths with the parametric model of the main paper. There are two components to this model. We first look at the nature of the age component. Figure 5(a) in the paper (Richards, 2022a) shows the log (mortality hazard) against age for the UK portfolio with no other risk factors. There is a gradual increasing trend with age. Apart from some random variation, log(mortality) is essentially a monotonic increasing function of age, and it seems to have an underlying smoothness, especially
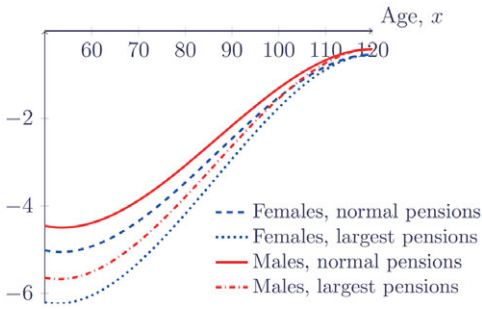
Figure 11. Hermite spline model of mortality differentials.
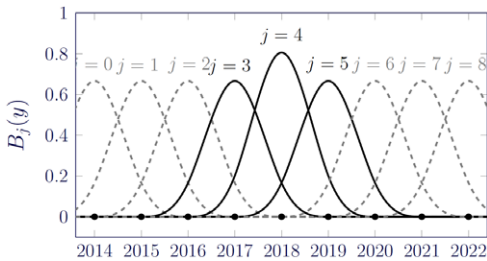Source: Richards (2020, Figure 5)

Figure 12. Basis of B-splines with fifth spline scaled independently of the others.

between the ages of 65 and 100. Indeed, seasonal mortality patterns are very reliable, as shown in Figure 10:

We have seen that there are pronounced seasonal patterns in the UK, France and the USA, with peaks in winter mortality in January of each year and troughs in the summer. In the Southern Hemisphere, this pattern is shifted by six months, which is why Figure 10 has mortality peaks in the middle of the year. This seasonal variation is universal, and Richards et al. (2020) found the same pattern in Canada, the US, Kuwait, Australia, UK and France. Unpublished results for South African pensioner data found a similar pattern to Australia.

Unlike age, mortality across calendar time is never monotonic. However, mortality rates are still smooth on a day-to-day basis, even during a pandemic – the daily deaths due to COVID-19 in Figure 3 in the paper (Richards, 2022a) still form a smooth underlying curve. This gives rise to very different modelling requirements. For mortality by age, we have relatively slow, gradual, monotonic changes for which we will need relatively little flexibility. In contrast, mortality by calendar time is subject to relatively fast, non-monotonic changes, for which we need greater flexibility.

This brings us to the continuous age-period model. The log(mortality hazard) at age x and time y will be a monotonic age component plus a locally flexible period or time component, as in equation (7) in the paper (Richards, 2022a). For the age component we use Hermite splines described in section 4 in the paper (Richards, 2022a). For mortality modelling we often just need two of these four splines; Figure 5 in the paper (Richards, 2022a) shows that a reasonable fit can be obtained with just the h_{00} and h_{01} Hermite splines. The fitted coefficient of the h_{00} spline is $\hat{\alpha} = -5.68$, while the fitted coefficient of the h_{01} spline is $\hat{\omega} = -0.625$. We can add other risk factors by varying α, meaning that modelled mortality rates will converge with increasing age without crossing over. We can add factors like pension size, gender, early retirement status etc., and we will have differentials at younger age. However, varying α only means that these differentials largely vanish by older ages. Crucially, the Hermite spline approach stops modelled mortality rates from crossing over, as shown in Figure 11.

To model the time component, we need more flexibility. We use a basis of identical, overlapping cubic B-splines (Schoenberg, 1964), as shown in Figure 12 in the paper (Richards, 2022a); the thick black dots are equally spaced knot points. Crucially for our purposes, B-splines are local functions, so we can raise or lower them without affecting neighbouring functions, as shown in Figure 12:

The summation term in equation (7) in the paper (Richards, 2022a) is a weighted sum of these B-splines, forming a flexible time-varying component to mortality. Each B-spline is multiplied by a corresponding κ_0 coefficient to vary the spline's height locally, as in Figure 12.

This continuous age-period model fitted to the UK data set using knot points at the first of January in each year is shown in Figure 14(a) in the paper (Richards, 2022a). Instead of multiplying solely by $\kappa_{0,j}$, we introduce a normalisation constant that is chosen so that the curve takes the value 0 in October 2019, as in Figure 14(b) in the paper (Richards, 2022a). October 2019 is chosen here because it is midway between the summer 2019 trough and the following winter peak. It is also the final summer-winter mid-point before the COVID-19 pandemic started. The curve takes the value 0 at that reference point and varying positive and negative values express differing relative levels of mortality in time.

Figure 14(a) in the paper (Richards, 2022a) uses one knot point per year, which results in quite a heavy level of smoothing. If we increase to two knot points per year, as in Figure 15(a) in the same paper, we immediately see the extent of seasonal variation in the portfolio. This figure shows, shows for example that the 2017/2018 winter peak is higher than the 2018/2019 winter peak, which we saw with the non-parametric estimator in Figure 5 earlier. The difference is that Figure 15(a) (Richards, 2022a) is just the time component in isolation, as other risk factors like age, gender and pension size are accounted for by the Hermite spline component. Figure 15(a) (Richards, 2022a) is the remaining time-varying mortality only.

Figure 15(a) (Richards, 2022a) has merged the 2019-2020 winter peak with the COVID-19 shock. To separate the two, we can increase to four knot points a year to see more detail, as in Figure 15(b) (Richards, 2022a). Here the winter peak in 2019/2020 is subdued, but we have more clarity on the COVID-19 shock in April 2020 and a little bit more detail on the extent of the mortality peaks in the winters of 2016/2017 and 2017/2018. However, there are limits to adding knots, as shown in Figure 15(c) (Richards, 2022a) with ten knots per year. We have a very clear illustration of the COVID spike, but we also have random variation because we have now over-parameterised the model.

However, knot points do not have to be equally spaced. For example, we could use two knot points per year for the standard seasonal variation and then place extra knots only where we already know from the population data that the shocks occur. Figure 16 in the paper (Richards, 2022a) shows how the B-spline basis changes when we add additional knot points in the first half of 2020. We have the usual equally spaced knot points at the start and middle of each year, and we have added four knot points in the first half of 2020 to give greater flexibility to handle the mortality shocks. This unequal knot spacing changes the shape of the splines. The greater density of splines in the first half of 2020 gives greater local flexibility as a result.

Figure 17(a) in the paper (Richards, 2022a) shows the result of using the B-spline basis in Figure 16 and normalising in October 2019. This model captures a lot of time-varying features of portfolio mortality. We have the expected seasonal variation with a rather dramatic COVID-19 mortality spike in April 2020. Interestingly, we also have a very deep trough in mortality following the first COVID shock. There are several possible reasons for this. It may be that there is a degree of "harvesting," i.e., some of the deaths in the mortality shock were of the frail, leaving a somewhat healthier post-shock population experiencing lower mortality. However, it may also be partly due to the strict lockdown that then applied – for example there were fewer traffic accidents because people were travelling less. There could be other reasons too.

All these results have so far been plotted on a logarithmic scale, i.e., Figure 17(a) in the paper (Richards, 2022a) shows the addition to log(mortality). We can alternatively work on the hazard

scale instead, meaning we take the exponent of our *B*-spline basis, as shown in Figure 17(b) (Richards, 2022a). Here the reference value is 1 in October 2019, and we can see that the peak-to-trough seasonal variation can be as much as 30%, which is quite dramatic. We can also see that, relative to October 2019, the mortality shock in April 2020 constituted a near-doubling of the mortality hazard.

This continuous age-period model allows us to revisit the topic of reporting delays. Normally we discard data that is affected by reporting delays to avoid bias. However, with a flexible *B*-spline basis for time-varying mortality, we no longer need to throw data away. Figure 18 in the paper (Richards, 2022a) shows the mortality multiplier for the UK portfolio using all the data from 2015 right up to the date of the extract in 2021. We can see the seasonal variation and the two COVID shocks (with the second shock not as pronounced as the first). We can also see that the reporting delays are handled right up to the extract date as the multiplier tends to zero. Using a flexible *B*-spline basis therefore allows us to include periods of mortality affected by reporting delays without biasing the parts of the model that cover age and other risk factors. Separating the time-based component from the other risk factors allows us to build a mortality model for age and other risk factors that will not be biased upwards or downwards by either seasonal variation, mortality shocks or even late-reported deaths. By including this flexible time component, we can remove the bias in the rest of the model, and we can then make an actuarial judgement as to what is suitable for a long-term basis. For example, an actuary might decide in this case that October 2019 was a normal level of pre-shock mortality and can then use this insight to derive an unbiased basis for pricing or reserving.

This brings us to the closing topic of mortality improvements. The same flexible *B*-spline basis that removes the bias in modelling can also help estimate portfolio-specific mortality improvements. We consider the time component at two calendar times, $y_1 < y_2$, and we suggest that we use midsummer points as midwinter mortality levels can be more variable. Figure 13 illustrates this:
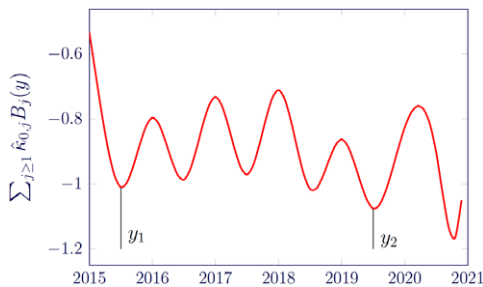

Figure 13. Addition to log(mortality) for UK annuitants using two knots per year.

We set y_1 to summer 2015 and y_2 to summer 2019 and estimate the implied improvements. We see that the summer troughs vary, so any estimate of mortality improvements depends on the choice of start and end point. We have deliberately chosen not to use the deep summer trough in 2020 because we think it is unusually deep, which would lead to an artificially high improvement rate. We calculate the annual improvement rate for the portfolio between y_1 and y_2 using equation (10) in the paper.

For the UK portfolio, the aggregate portfolio-specific annual improvement rate was about 1.2% per annum. This could be compared with the assumptions used in reserving. An additional application is for pricing bulk annuities and longevity swaps, where the reinsurer taking on the liability would be interested in knowing if a portfolio had slower or faster mortality improvement rates than the population at large.

To conclude, for such analyses we need to use all-cause mortality data because most portfolios do not have cause-of-death data. Even for the few portfolios that do, the coding of COVID-19 mortality was incomplete during the first shock. We need to use continuous-time methods that are capable of handling rapid changes in mortality levels in calendar time. For this we have two options. We have a semi-parametric estimator, which is useful for data exploration and communication with non-specialists; and we have a parametric that can be used for pricing or reserving. The parametric model is a continuous age-period model with two separate components. For modelling by age and other risk factors we need relatively little flexibility, for which we can use a basis of Hermite splines. For modelling in calendar time, we need considerably more flexibility, and for this we use a basis of B-splines. Where we know that there are pandemic shocks that require extra flexibility, we simply add extra knot points to the B-spline basis. Then we can exercise judgement as to what is the standard level of mortality for creating an unbiased best estimate for pricing or reserving. We can also use the same model to estimate portfolio-specific mortality improvement rates. We can even use the same method to allow for unreported deaths or reporting delays in the same way that we can allow for an upward bias that might come from mortality shocks.

Special thanks are due to Anil Gandhi and Caroline Roberts, who provided the data for this research. Without their kind support this research never could have happened.

The Moderator: thank you very much Stephen. That was a fascinating presentation. We now have some questions for you.

Question: The Nelson–Aalen estimate in Figure 4 showed a lowering in the trend between 2013 and 2017. Is there any intuition as to why that happened?

Dr Richards: I think the right panel of Figure 4 is what Cameron is referring to. One important drawback of the semi-parametric approach is that it does not include age. This is the reason why the semi-parametric estimator of the hazard is increasing in time. This portfolio is open to new business, but it is ageing because the new annuities being set up do not balance the deaths. The portfolio is ageing in calendar time, so the aggregate mortality hazard rate is rising over time. Figure 4 shows not only that the hazard level rising because of the increasing average age, but the sensitivity to season is also increasing over time. This was a result in Richards et al. (2020), who showed that seasonal variation becomes more pronounced with increasing age.

The reason why there is a relatively modest change in mortality hazard in 2013-2017 is probably because there was a small surge of new business in this portfolio that reduced the average age, as shown in Figure 14:
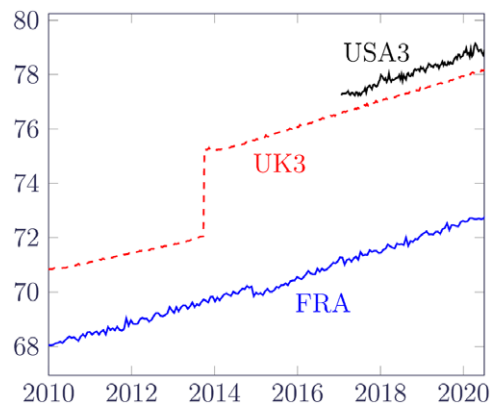

Figure 14. Average age over time of in-force annuitants.

The portfolio of French annuitants wrote a lot of new business in September 2014, which reduced the average age. Similarly, the UK portfolio reinsured many younger annuitants in late 2013, resulting in a large step increase in the average age. Source: Richards (2022b, Figure 1).

The new additions to the portfolio in the more recent years have been fewer, with less new business as the portfolio continues to age. The French portfolio is for employees at higher education establishments around France (universities and colleges). As a result, annuities to retirees start in September, this being the start of the academic year.

Question: The other risk factors such as pension size are only included in the Hermite splines, in other words the age component of the model. Does that run the risk of missing some information on how these factors may affect the COVID part of the model or the shock part of the model?

Dr Richards: In this paper I have just used a simple time-varying basis. However, you could interact this with risk factors like gender, early retirement status or socio-economic group and thus have different bases and different $\kappa_{0,j}$ values for different sub-groups. To keep things simple, I did not do this with this analysis, but it would be possible to interact the κ_0 terms with other parameters and have different time patterns for different sub-groups. Richards et al. (2020) interacted the seasonal effect with age and found that older people experienced greater seasonal variation in mortality. However, we also looked at socio-economic status and deprivation indirectly by looking at pension size quantiles and found that lower-income groups also had more extreme swings in seasonal variation. I would expect the same to apply with COVID-19. If we were to interact pension size with the coefficients of the B-splines, I think we would see a greater impact on people with lower pension values. There is nothing in the model that would stop you from interacting any risk factor with the $\kappa_{0,j}$ terms.

Question: What about including vaccination status as a downward shock? I suppose you could add to that by having the emergence of a succession of new variants as a series of rolling shocks, hence modelling vaccination versus variants.

Dr Richards: If you had information on vaccination status then it would simply be a risk factor like gender or socio-economic status. Then you would probably very much want to interact it with the $\kappa_{0,j}$ parameters, because you would expect the unvaccinated to have very different responses to the second COVID shock. Of course, at the time of the meeting (December 2021), we are on tenterhooks as to whether we are going to see a third COVID shock with the Omicron variant. We would expect to see stark differences in mortality by vaccination status. Note that vaccination status, while interesting, will not be routinely available to insurers. It would be difficult to collect and there would be additional privacy implications over and above those from using postcode and date of birth. Vaccination status would count as sensitive health data, which has specific handling rules under GDPR. It would be an interesting risk factor, but probably too hot to handle.

The Moderator: Thank you for that. That brings us to the end of this sessional meeting. I would first like to thank Stephen (Richards) for a fascinating presentation and an interesting Q&A session. I would like to thank the events management team for facilitating the meeting and finally my thanks to all the participants who have attended and taken part in the meeting.

References

de Looper, M. (2002). Seasonality of death, Bulletin No. 3, Australian Institute of Health and Welfare, available at https://www.aihw.gov.au/getmedia/e4e2219d-1745-4cd9-a221-eb7ac1a1e476/bulletin03.pdf.aspx?inline=true

Macdonald, A.S., Richards, S.J. & Currie, I.D. (2018). *Modelling Mortality with Actuarial Applications*. Cambridge, Cambridge University Press. ISBN 978-1-107-04541-5.

Richards, S.J. (2020). A Hermite-spline model of post-retirement mortality. *Scandinavian Actuarial Journal*, **2020**(2), 110–127. doi: 10.1080/03461238.2019.1642239.

Richards, S.J., Ramonat, S.J., Vesper, G. & Kleinow, T. (2020). Modelling seasonal mortality with individual data. *Scandinavian Actuarial Journal*, **2020**(10), 864–878. doi: 10.1080/03461238.2020.1777194.

Richards, S. (2022a). Allowing for shocks in portfolio mortality models. *British Actuarial Journal*, **27**, E1. doi: 10.1017/S1357321721000180.

Richards, S. (2022b). Real-time measurement of portfolio mortality levels in the presence of shocks and reporting delays. *Annals of Actuarial Science*, 1–23. doi: 10.1017/S1748499522000021.

Schoenberg, I.J. (1964). Spline functions and the problem of graduation. *Proceedings of the American Mathematical Society*, **52**, 947–950.