I'm glad to see so many actuaries in the audience tonight to talk about longevity risk, especially when we know that there are so many other pressing economic issues. For example, when we calculated the illustrative annuity factors in our paper, we used a discount rate of 5% per annum. This was the Bank of England base rate in early October 2008. Just four months later, the base rate has fallen to 1%, an historic low since the founding of the Bank of England in 1694.

Just like interest rates, mortality rates have also fallen to a historic low. However, mortality rates still have plenty more scope to fall. Figure 4 in the paper shows a clear downward trend for nearly half a century at various ages. The obvious assumption is that mortality rates will continue to fall. But how far? And how fast?

One view of mortality rates is so-called longevity risk, namely the risk that pensioners or annuitants live longer than assumed in pricing and reserving. A key question is what direction the trends in Figure 4 will take. This is not just a question of what the likeliest trend is, but of what future trends are possible and the likelihood of each. It is for this reason that we prefer statistical models for projections, rather than deterministic scenarios. A particular scenario might be interesting, for example if it involves extrapolating trends in causes of death, but a scenario is not particularly informative without a statement of its likelihood.

Both life-office annuities and defined-benefit pensions must be reserved for prudently. There is no set definition of prudence for a mortality projection for pensions and annuities. A minimal approach to prudence is that mortality rates should be more than 50% likely to be higher than assumed. In other words, a prudent reserving basis would use rates below the central projection from a statistical model. The reserving basis can itself be a deterministic scenario, possibly even one derived from a cause-of-death-based extrapolation. However, such scenarios would have to be benchmarked against the results of a statistical projection to prove themselves as prudent.

There are two main problems with this, however. The first is *model risk*, namely the fact that you cannot know which model will be appropriate for projection. This problem exists for all forms of forecasting, but was succinctly expressed by Cairns (2004) when writing about interest-rate models:

"probability statements derived from the use of a single model and parameter set should be treated with caution"

This is illustrated in Figure 5, where two different parameterisations of the Lee-Carter model lead to two very different projections. This does not necessarily mean that either — or both — of the models are wrong. Rather, it shows that even slight changes to model structure can radically change the direction and apparent uncertainty over the projections. Model risk can be demonstrated, but not completely quantified — a "known unknown" to quote Donald Rumsfeld. For this reason the prudence of a reserving basis needs to be benchmarked against more than one statistical projection. Alternatively, an actuary might include an explicit margin for model risk.

The second problem with statistical projections is less philosophical, namely which data set to use. Ideally, an actuary would have a forty-year data set for exactly the portfolio being reserved for. In practice, this is almost unattainable: few portfolios have historic mortality this far back. One solution is to use a separate data set which does have the necessary history. However, this introduces *basis risk*, namely the fact that the data used for projections is not the same as the portfolio for which the actuary is reserving. The two most obvious candidate data sets in the UK are population data from the ONS and assured-lives data from the CMI. The ONS data isn't entirely satisfactory as private pensioners and annuitants have markedly lighter mortality than the general population. Neither is the CMI data ideal: there is little data for female lives and Figure 7 suggests the data has undergone radical shifts in size and may have changed composition as a result. Actuaries have to accept these limitations and work with what data is available. However, as with model risk, an actuary should include an explicit margin for basis risk for pensioner mortality projections.

A further problem with basis risk is that all lives are not equal in financial terms. Richards (2008) showed the degree of concentration of liabilities in a small proportion of lives: the top decile of pensioners accounted for around half of all the pensions in payment, both for a large annuity portfolio (54%) and for a collection of defined-benefit pension schemes (46%). The possibility of this top decile experiencing different mortality trends from the rest is also a form of basis risk. However, it can also be categorised as part of *concentration risk*, since there are other consequences of this lack of equality. For example, different social groups have different absolute levels of mortality, which can have a much bigger impact than mortality projections.

Concentration risk also interacts with stochastic or *binomial risk*, namely the uncertainty over who dies when even if you knew exactly what the underlying mortality rates were. For example, Tables 9–11 show the volatility in run-off by lives, but also that heterogeneity in pension size increases that volatility. For very small pensions schemes (of under 50 lives, say), the binomial risk will often dominate all other sources of risk, including investment risk.

All of these risks — model risk, basis risk, concentration risk and binomial risk — can be quickly and easily explored with modern computers. Nowadays there seems no need to pick model points as entire portfolios can be simulated in run-off on a life-by-life basis. This requires deduplication to get the correct picture, however, as key sub-groups tend to have multiple policies, as shown in Figure 1. Failure to deduplicate would give a falsely comforting picture by understating concentration risk and over-stating the benefit from the law of large numbers. By parallelising the calculations across a standard eight-core system, even a portfolio of 200,000 lives can be simulated 10,000 times in an hour. The resulting data can be presented graphically in a single chart, such as Figure 10. The horizontal axis shows the reserve for the present value of benefits. The vertical axis shows the probability of a given reserve being adequate in run-off.

For a given projection model, the 10,000 run-off costs are arranged in ascending order to form an ogive, i.e. a cumulative distribution function. This can be done for as many projection models as you like on the same graph, but Figure 10 just compares two. An actuary can use any basis she likes to calculate a reserve, including a deterministic scenario, and then look up the ogives to see the corresponding probability of that reserve being adequate with respect to longevity risk (we are ignoring various other risks, such as investment returns and expenses, but these could be built in as well). For example, Figure 10 shows vertical lines for various reserving bases and the points where they cross the ogives shows the corresponding probability of adequacy in run-off under each model. For example, the line for the long-cohort reserve in Figure 10 suggests it is moderately prudent under the DDE model with around 60% probability of adequacy. However, under the CR model it is only adequate with around 10% probability. Conversely, a reserve which was minimally prudent under the CR model with a probability of adequacy of 50% would have a probability of adequacy of 95% under the DDE model.

To sum up, the prudence of mortality projections in pension and annuity reserving can only really be demonstrated with reference to the results of a statistical model. The measures of uncertainty are the key here, and these cannot be obtained from deterministic scenarios. However, both model risk and basis risk are substantial and actuaries need to consider the results of several alternative models to demonstrate prudence.