![Bank of England logo] BANK OF ENGLAND

Speech

What Fates Impose: Facing Up To Uncertainty

Speech given by

Mervyn King, Governor of the Bank of England

At the Eight British Academy Annual Lecture, London

1 December 2004

What Fates Impose: Facing Up To Uncertainty
The Eighth British Academy Annual Lecture 2004

Mervyn King

"*What fates impose, that men must needs abide;*
It boots not to resist both wind and tide".
(William Shakespeare, Henry VI part 3, Act IV Scene 3)

1.   Introduction

My subject tonight is risk: how we think about, how we talk about, and how we manage risk. So let me start by thanking all of you for coming and taking the risk of spending an hour with a central banker. Few writers have put central bankers in their place, but one of them was Nancy Mitford, whose centenary we celebrated last Sunday. In *The Pursuit of Love* and *Love in a Cold Climate* she describes the marriage of Linda, the heroine, to the son of the Governor of the Bank of England. The marriage between Linda and the incredibly boring Tony is doomed from the outset. As Lady Montdore explained, 'Bankers don't seem to be much to look at – so extraordinarily unsuitable having to know them at all, poor things, let alone marry them'.

I don't suppose Lady Montdore would have come to this lecture, but if she had I would point out to her that life is full of surprises. Risk, luck, fate, uncertainty, probability theory – we all have names for the game of chance. Most decisions in life involve risk. Sometimes we embrace it, as when we enjoy a bet on the Grand National, and sometimes we avoid it, as when we insure our houses against fire. The playing of the hand we are dealt may be a pleasure in bridge but a burden in life. We accept that Lady Luck has her part to play in our personal lives. But in our collective life – public policy – the role of probabilities rarely takes centre stage. An informed discussion of public policy issues, however, requires an analysis of the risks and uncertainties involved. Whether in policies for health or transport, matters monetary or meteorological, in times of war and peace,

decisions should reflect a balance of risks. Yet policy debates continue to be permeated by the 'illusion of certainty'.[1]

The reluctance to give adequate prominence to risks may reflect the fact that many of us feel uncomfortable with formal statements of probabilities. Probability theory is relatively recent in our intellectual history, dating back to a flowering of ideas around 1660 from Pascal, Leibniz, Huygens and others.[2] Despite advances since then, statistical thinking remains prone to confusion and is often avoided. Television weather forecasts in Britain rarely employ the language of probabilities used by the meteorologists themselves. Professor Gigerenzer of the Max Planck Institute in Berlin has demonstrated in a series of studies how poorly doctors, lawyers, and other professionals understand probabilities. And despite Seneca's maxim that 'luck never made a man wise', airport bookshops stock titles on how to become rich by successful investors and entrepreneurs who are confident that their success is the result of outstanding business acumen rather than good fortune.

Many of these misunderstandings stem from a failure to grasp basic statistical concepts. Juries are not informed that, in a country of our size, multiple cot deaths are likely to occur several times a year, that several people will have DNA that matches the incriminating sample, and that in themselves these coincidences are not evidence of guilt. Bookshops do not stock such titles as 'I would have been a billionaire if only Lady Luck had been faithful'.[3]

In my lecture tonight I want to illustrate two propositions. First, in a wide range of collective decisions it is vital to think in terms of probabilities. We cannot avoid taking decisions, so we must accept the need to analyse the uncertainty that inevitably surrounds them. H. G. Wells is often reported as saying that, 'Statistical thinking will one day be as necessary for efficient citizenship as the ability to read and write'. And that is even more

---

[1] To use the phrase of Professor Gerd Gigerenzer who has done more than anyone to draw attention to the lack of public understanding of risk.
[2] For an intellectual history of probability theory see Hacking (1975) and Bernstein (1996).
[3] Among recent discussions of misperceptions of probability are Gigerenzer (2002) and Taleb (2001).

true today when we are inundated with statistics from every quarter. Second, in order that public discussion can be framed in terms of risks, the public needs to receive accurate and objective information about the risks. Transparency and honesty about risks should be an essential part of both the decision-making process and the explanation of decisions.

I want to illustrate those two propositions by considering as an example public policy about pensions – an issue, you might think, of particular interest to many of us in the Academy. When the Pensions Commission reported in October, it highlighted the financing gap in our present system. But we must not lose sight of the equally important question of what are the risks incurred in pension provision and how should they be shared among us? It is not my intention to make any recommendations. That is for the Pensions Commission next year, and the Government in its turn. But I do want to show that risk is at the heart of the issue.

2.   The Nature of Risk

'In this world nothing can be said to be certain, except death and taxes' wrote Benjamin Franklin in 1789. Great man though Franklin was, it is clear that there is indeed substantial uncertainty both about death, or at least its timing, and the contributions or taxes required to pay for our pensions.

When thinking of how much to save for retirement, there are three main risks to consider. First, how much will we be earning in the future? Second, what will be the rate of return on our savings? Third, how long are we likely to live? Tonight I shall focus exclusively on longevity risk – the uncertainty about how long we shall live – in order to focus on my two propositions about probabilistic thinking and public policy.[4]

---

[4] There is a large literature on uncertainty about earnings and productivity growth; see, for example, MaCurdy (1982), Abowd and Card (1989) and Meghir and Pistaferri (2004). Uncertainty about future returns on financial assets is the essence of the modelling of financial markets. An excellent introduction to this area is contained in the two 2003 Nobel lectures by Robert Engle (2004) and Clive Granger (2004).

None of us know how long we shall live. But we can assess our chances by looking at the experience of others. A key distinction is between individual and collective risk. Chart 1 plots the observed mortality rates at different ages for women in England and Wales in 2002, the most recent year for which detailed data are available. (A similar story can be told for men.) The most common age of death for women was 87, but many died at very different ages. Mortality among individuals varies greatly – individual risk.

But there is also significant uncertainty about the average length of life of a generation as a whole – collective risk. Longevity has changed markedly over time. Chart 2 shows female mortality in England and Wales in both 1902 and 2002. It shows the remarkable reduction in infant mortality during the twentieth century as well as the fall in mortality in middle age.

Of particular relevance to decisions on pensions is the expectation of life conditional on reaching retirement age. That is directly relevant to the purchase of an annuity on retirement. Over the past century the normal retirement age in public service in this country has been constant at age 60. Chart 3 shows the distribution of deaths by age conditional on having reached the age of 60 in England and Wales in both 1902 and 2002. A woman who was 60 in 1902, and subject to the mortality rates of that year, would have expected to live for another 14½ years. By 2002 that expectation had increased to 23½ years.

Individual risk – that we die either earlier or later than the average for our cohort – can be insured by taking out life insurance against early death or by purchasing an annuity to insure against a later death.[5] Collective risk – that <u>average</u> expectation of life will rise or fall in the future – cannot, by definition, be shared among members of the same generation. The principle of insuring individual longevity risk is that those who live longer than average can be paid using the contributions of those who die earlier than

---

[5] The first estimates of mortality rates using bills of mortality from the City of London were published by John Graunt in 1662. But the widespread use of such data to price annuities began only with the calculations of Richard Price in 1780 and the Institute of Actuaries in 1869. In Holland more accurate calculations of the prices of annuities date back to de Witt in 1671. See Poterba (2004).

average. But if the average itself rises all contributions must increase – there is no possibility for insuring changes in the average from within the cohort. Can we quantify this risk? The expectation of life has risen significantly over the past century. But how will it change in the future? To answer that question we need to model and quantify longevity risk.

3.   Modelling and quantifying risks

Quantitative modelling has been one of the great advances of the social sciences over the past fifty years. But forecasts, whether of longevity or any other economic or social variable, inevitably require a judgment about how far the past can inform us about the future. Projecting longevity is fraught with difficulty. In his *First Essay on Population*, Malthus wrote in 1798 that 'with regard to the duration of human life, there does not appear to have existed, from the earliest ages of the world, to the present moment, the smallest permanent symptom, or indication, of increasing prolongation'. We know now that in 1798 life expectancy in Britain was around 40. Today it is nearly 80.

There is no agreed method of forecasting life expectancy. The most widespread approach is to assume that life expectancy will gradually approach some limit determined by medical knowledge and our choice of lifestyle. A key uncertainty is what the limit to life is, and demographers have pursued many ways of estimating it.[6] But successive forecasts of the limit to life have been disproved as actual life expectancy data surpassed the

---

[6] Manton *et al.* (1991) and Caselli and Vallin (2001) describe alternative approaches to inferring the limit to human life. Some estimates assume that life expectancy will converge to that of the longest-lived population; for example, that life expectancy in the West will rise to that in Japan, the country with the highest life expectancy. But the level in Japan has itself increased significantly over the past century. So what determines Japan's life expectancy? Another approach is to examine how life expectancy might change if everyone adopted lifestyle patterns that increase longevity. Mormons and Seventh Day Adventists seem to live longer than the rest of us, but exactly which aspects of their lifestyle we should adopt is not entirely clear. See, for example, Enstrom (1989). Yet another approach is to work out the contributions that different human diseases make to premature death, and calculate the hypothetical life expectancy that would result if medical science were to be able to eliminate these diseases. This 'cause elimination' approach was adopted in the 1950s by Pichat. The number derived was 75. This is the number that the UN used until the early 1980s in their population projections. The number used from then on has been about 85, based on a similar approach by Benjamin (1982) and Olshansky *et al.* (1990). And the evidence on a limit to increases in the human lifespan is mixed; see for example, Kannisto *et al.* (1994) and Wilmoth and Horiuchi (1999).

supposed limit. Chart 4 plots the forecast limits to female life against the life expectancy achieved by the longest lived female population. As life spans have increased, forecasts of the limit to life have been revised up.

The fact is that we simply do not know how life expectancy will change in the future. Unforeseen, indeed unimagined, developments in medical science could change the prospects for life expectancy radically. In the end we have to face up to our ignorance. Providing insurance against longevity, whether as private annuities or public pensions, rests on an ability to calculate the risks involved. But our present degree of knowledge means that, in the terminology of Frank Knight (1921), there is both risk, which can be quantified, and uncertainty, which cannot.

The extent of past uncertainty about life expectancy can be seen in the revisions to its forecasts made over the past 20 years by the Government Actuary's Department. Nor has the financial sector itself proved better at forecasting. Chart 5 shows successive forecasts of expected length of life for men aged 60 published by the Actuarial Profession, a body of life assurance companies and annuity providers. It shows how reductions in mortality were expected by the industry to lead to rises in length of life. For example, the yellow line shows the Actuarial Profession forecast made in 1980 that a man who was 60 in that year could expect to live another 20 years. At that time, it was thought that someone who reached 60 in 1999 would live a further 21 years. But by 1999 the forecast was that a man of 60 would live another 26 years. Over a twenty year period, expected length of life was revised up by 5 years. Length of life has risen faster than the industry anticipated, as can be seen by the upward shifts over time both in the forecasts of expected life and in the rate at which it was expected to increase.[7]

That there are large and unavoidable uncertainties in forecasting longevity is clear. The key unknowns are evolutions in lifestyle, diseases and the ability of medical science to combat them. No amount of complex demographic modelling can substitute for good

---

[7] The data in Chart 5 are for occupational pension holders of Actuarial Profession member institutions. These individuals are typically wealthier than the rest of the population, and for that reason tend to live longer.

judgements about those unknowns. The danger of relying on computer modelling is well illustrated by the World Health Organization's early estimates of the future spread of AIDS, which, at the time, was a significant unknown in forecasting life expectancy even in developed countries. AIDS was first recognised in 1981. The virus responsible for AIDS, which we now call HIV, was identified in 1984. To date, around 20 million people are thought to have died from the disease worldwide. The proportion of those who are infected with HIV who go on to develop AIDS is now thought to be close to 100%. But the spread of HIV, and hence AIDS, has been very uneven. At the end of last year the percentage of adults infected with HIV was estimated at 0.2% in the United Kingdom, 21.5% in South Africa, 37.3% in Botswana, yet only 3.9% in Angola.

Why do infection rates differ so much between countries? In the 1980s the World Health Organization and associated researchers made rather optimistic projections of the impact of AIDS on populations in Africa using extremely detailed computer models of population structure. Complex demographic models were combined with strong simplifying assumptions about sexual behaviour. Unfortunately, this led the modellers to ignore those factors most critical to the spread of AIDS. They assumed that the probability of transmitting the HIV infection was a function of the number of sexual acts, not the number of partners of infected persons. But it is crucial to distinguish between the two.[8] Ten acts with one partner have a smaller impact on transmission of the infection than one act with each of ten different partners. As Professor (now Lord) May, President of the Royal Society, pointed out, ignoring this simple and crucial distinction led to a serious underestimate of the rate at which AIDS was likely to spread.[9] By building elaborate models of demographic variables for which a great deal of information was available and making simple assumptions for behaviour about which

---

[8] Peterman *et al*. (1988) found that the probability of infecting a partner was unrelated to the number of sexual acts with that partner.

[9] The large models made the further simplifying assumption that the rate took one of only two values. In fact, it varies widely among the population. Robert May, and his colleague Roy Anderson, showed that a key driver of the proportion of the population eventually infected was the 'variance of the rate of acquisition of new sexual partners'. Simulations of the proportion of the population ultimately infected with AIDS are sensitive to that variance, and plausible differences in the number of their partners among different people changed the estimate of the proportion likely to be infected from 25% to 90%.

little was known, the large models used by the World Health Organization produced misleading results. The fact that few data on sexual behaviour were available was not a good reason for focusing on demographic variables – the researchers were looking for the key under the lamp-post rather than where it fell.

Quantitative models are obviously essential to explain empirical evidence. But the idea of a single comprehensive model, whether of a population or an economy, which can capture all the relevant features to explain all phenomena is an illusion.[10] There is no substitute for thinking about the underlying process which is generating the data. Robert May's article 'Uses and Abuses of Mathematics in Biology', published earlier this year, should be compulsory reading for anyone involved in public policy. In pointing to the dangers of models so complex that intuition is impossible, he writes, 'It makes no sense to convey a beguiling sense of "reality" with irrelevant detail, when other equally important factors can only be guessed at. … Remember Einstein's dictum: "models should be as simple as possible, but not more so"' (May, 2004).

In many areas, successful public policy depends on a quantitative analysis of risks by experts and researchers. It is important that the general public understand that such assessments are only judgements. To enhance confidence in those judgments experts need to engage in a debate among themselves about what they know and what they do not know. On the Bank of England's Monetary Policy Committee the nine members debate among themselves our current state of knowledge of the behaviour of the economy. The results of that discussion are published once a month, and the Committee, as it did yesterday, appears regularly in front of Parliamentary committees. So I turn now to the question of how information and judgements of risks can be communicated to the general public.

---

[10] Many non-economists believe that economists use a single 'model' to explain all economic phenomena. Hence the interest of many commentators in 'the model' that is alleged to generate the forecasts of inflation and economic growth published by the Bank of England. But, as we never cease pointing out, the Bank uses not one but many models to help us understand particular aspects of economic behaviour. The relevant models change from one forecast to another.

4.   Communicating about risks with the public

Although it is often impossible to be precise about the risks surrounding a decision, it is vital that policymakers resist the temptation to communicate a false sense of certainty. Two key principles should govern the communication of risks to the public. First, information must be provided objectively and placed in context so that risks can be assessed and understood. Second, experts and policy-makers must be open about the extent of our knowledge and our ignorance. Transparency about what we know and what we don't know, far from undermining their credibility, helps to build trust and confidence in policy-makers.

(a) (Mis)understanding probabilities

The first principle states that information about probabilities should be placed in context. A simple example illustrates the difficulty we all have in interpreting probabilities when the context is unclear. The stock market is volatile and difficult, if not impossible, to predict over short periods. At the beginning of the week the chance of the market rising over the following week is roughly the same as the chance of its falling. So if I were to predict the direction of the market movement correctly for five successive weeks you might think that I knew something you didn't. Indeed, you might be willing to subscribe to an investment service with that sort of track record. So let me explain to you how one might create the illusion of clairvoyance. Select around 6000 or so names and addresses from the London telephone directory. Divide the names into two groups. To the first write predicting that the market will rise over the coming week. To the second write predicting a fall in the market. At the end of the week keep the 3000 or so names who were given the correct prediction and discard the others. Divide those names in turn into two groups. To the first predict a rise in the market and to the second a fall. Repeat this process for five weeks, at which point there will be around 200 people to whom we could write the following letter. "You may well have been sceptical when you received our first letter, but by now you will know that we have worked out the secret of predicting

successfully the direction of movement of the stock market. You know that our method really works. To subscribe to our investment service please send £5000 by return".

Quite what our President, as the former Deputy Chairman of the Financial Services Authority, would make of this scheme I do not know. But it illustrates vividly that the interpretation of *ex post* outcomes depends critically on understanding the *ex ante* process which generated those outturns.

Another example shows that it is not just lay people who find statistical inference difficult. Experts of all kinds do too. The advent of DNA profiling in the 1980s led to the use of match probabilities in criminal cases. The jury is told that the probability of finding a match between the sample taken from the scene of the crime and the DNA of the defendant is, for example, only 1 in 500,000. That is sometimes taken as evidence that the probability of the defendant being innocent is also only 1 in 500,000. Such an inference is incorrect. In a city such as London where there might be about 5 million people who could have committed the crime, around ten people would have DNA that matched the relevant sample. Hence, in the absence of any other evidence, the probability that the defendant is guilty, far from being overwhelming, is only one in 10. Of course, in practice other evidence is usually available. But this incorrect statistical reasoning has swayed enough cases to be given its own name – the prosecutor's fallacy.[11]

The cases of Sally Clark and Angela Cannings, both convicted and imprisoned for the alleged murder of their own children, show the desperate consequences of the false use of statistics to create a presumption of guilt to resolve the cause of apparently inexplicable events. Both women were convicted by juries which had in all likelihood been influenced by the assertion that the probability of two cot deaths in the same family was extraordinarily low. That assertion was based on the assumption that cot deaths in the same family were independent events, a view for which there was no scientific evidence

---

[11] See, for example, Thompson and Schuman (1987).

and which is *a priori* implausible.[12] The assumption of independence was crucial to the prosecution evidence against Sally Clark that the odds of two cot deaths in a family such as hers were 1 in 73 million and so would be expected only once a century. Given the complex and conflicting medical evidence, it would not be surprising if that striking and simple statistic had played a role in helping the jury come to its conclusion. In fact, as pointed out by Professor Ray Hill, the statistical evidence on cot deaths suggests that in England and Wales we should expect several cases a year.[13] Thankfully, both women were eventually released after winning their appeals, albeit largely on other grounds.

Why then do probabilities cause us so much difficulty? Our understanding of probabilities evolved through counting the frequencies of various events, whether the number of days in summer on which it rains or the outcomes of tossing a coin. All this was part of our evolutionary experience long before the concept of probability was invented in the seventeenth century. In his book *Reckoning with Risk,* Gigerenzer argued forcefully that risks should be presented not as probabilities but in terms of natural frequencies, which correspond to our experience in counting events. Although people do understand probabilities in familiar contexts, such as betting on horse races or tossing a coin, it is not easy to calculate odds when the context is unfamiliar.

(b) Explaining probabilities

How then can we best present information about risks to the public? Both the Monetary Policy Committee and the Met Office face this problem almost every day: the former in explaining the outlook for inflation and the latter in forecasting the weather. Both have independently developed similar visual representations of forecasts that avoid the use of the word 'probability'.

---

[12] To understand the concept of independence consider the example of a church spire that is struck by lightning several times. Is this just coincidence (implied by the assumption that the strikes are independent of each other) or does the spire have certain characteristics that lead it to be struck regularly? Given our knowledge of spires and lightening the latter is more plausible. Similarly, there may be characteristics, genetic or other, and still largely unknown, that make particular families more likely to experience cot deaths than others.

[13] Hill (2002) and (2004). See also Dawid (2004) and Eggleston (1978) for a discussion of the difficulties courts find with the concepts of conditional probability and Bayes's theorem.

Chart 6 shows the forecast for inflation over the next three years published by the Bank of England in its latest *Inflation Report* on 10 November 2004. It is in the form of a 'fan chart'.   We do not say that in our view inflation will be 2%, or any other number. Such a statement is incoherent because a forecast is inherently probabilistic. So we represent the uncertainties by the coloured bands in Chart 6.[14] Their interpretation is as follows. If we found ourselves in economic circumstances identical to today's on 100 occasions, we would expect that inflation would lie within the darkest central band on only 10 of those occasions. As we increase the width of the coloured section by moving from the dark central band to the lighter shades of red further out, the number of outturns that we think will lie within the section increases. The bands are drawn so that the number of outturns increases by 10 for each change of colour intensity. The inflation outturn would be expected to lie somewhere within the entire coloured area on 90 out of 100 occasions.

On the same day that the Bank of England produced its forecast for inflation, the Met Office produced a forecast for the temperature in various towns over the following ten days. Chart 7 shows this 'plume chart', as the Met Office refers to its fan charts, for the town of Reading. I will leave a detailed explanation to meteorologists, but the interpretation of the temperature plume chart parallels that of the inflation fan chart closely. The width of the plume chart indicates the degree of uncertainty about future temperatures, and different colours suggest different probability bands. And the probability of the central temperature projection – in this case the solid blue line – is extremely small. I am not suggesting that the BBC use such charts in its forecasts, but surely it might give a little more freedom to the meteorologists to use probabilistic thinking in describing the outlook.

Such charts can also be used to illustrate uncertainty about longevity. Chart 8 constructs a fan chart for female life expectancy at birth in the United Kingdom derived from the latest forecast published by the Government Actuary's Department. The width of the fan

---

[14] The forecast for inflation is based on the assumption that interest rates follow the path expected by financial markets over an average of the fifteen days prior to the finalisation of the *Inflation report.*

chart shows how forecasts of life expectancy are likely to change in the future, as new information and ideas about mortality evolve, causing forecasts to be revised – assuming that news about life expectancy arrives in the future at about the same rate as in the past. The colour shading depicts the same probability bands as used in the fan chart for inflation that I showed earlier. Chart 8 illustrates the uncertainty about longevity.

Of course, statistics and their visual representation can be both used and abused. The design of charts has gone backwards since that (shown as Chart 9) of Napoleon's march to Moscow and subsequent retreat constructed in 1861 by Charles Joseph Minard, Inspecteur-Général des Ponts et Chaussées and, appropriately given tonight's theme, en retraite. Brilliant in conception and the quantity of information which it conveys, the chart shows the devastating losses suffered in Napoleon's fateful Russian campaign of 1812. So I do not underestimate the difficulties of promoting the better understanding of statistics in discussions of public policy.[15] Too often figures are bandied about without a careful description of where they came from and what they might tell us. As we learn from the Rime of the Ancient Statistician (with apologies to Samuel Taylor Coleridge):

> Figures, figures, everywhere
> Enough to make us sink
> Figures, figures, everywhere
> But none to help us think.

The second key principle governing the explanation of risk to the public is the open recognition of the limits of our present knowledge. The failure to do that is damaging to public confidence in policy-makers. The BSE affair showed the danger of claiming greater certainty than the science warranted. And in the Angela Cannings case, the Court of Appeal observed that, 'not so long ago, experts were suggesting that newborn babies should lie on their tummies. That was advice based on the best-informed analysis. Nowadays, the advice and exhortation is that babies should sleep on their backs – Back to Sleep. This advice is equally drawn from the best possible known sources'. As the Court of Appeal pointedly continued, 'It is obvious that these two views cannot both,

---

[15] On the use and misuse of statistics see Best (2001) and Briscoe (2000). An outstanding analysis of common mistakes in visual representations of statistical data can be found in Tufte (1983, 1990).

simultaneously, be right'. Expert advice is invaluable, but it is not infallible. As Robert May said recently, 'open contention of opinion is exactly what has served science so well over the centuries, and I believe we simply must learn to extend this openness to all forms of science advice in policy-making. Such open publication of advice, and frank admission of areas of uncertainty, ultimately engenders confidence' (May, 2003). That is the spirit in which the Monetary Policy Committee was created and continues to function. The challenge is to extend that spirit to other areas of policy.[16]

Effective communication of the relevant risks is vital. Thinking about how to present information may seem to be the role of a spin doctor. But whereas a spin doctor wishes to claim certainty for the correctness of his master's decision, presenting accurate information about risks to the electorate at large is essential to democratic accountability. As one of the more illustrious members of the profession might have said, 'spin doctors don't do doubt'. That is the trouble with much debate on public policy. As Voltaire put it, 'Doubt is not a pleasant condition, but certainty is an absurd one'.[17]

5. Managing Risks

Managing risks is about sharing risks with others. By finding ways of pooling risks we share both the upside and downside with others. This is beneficial because typically we dislike extremes. In the language of economics, we are 'risk averse'.

A key part of managing risk is the design of institutions that determine the ultimate incidence of risk. The Pensions Commission performed a valuable service in making us aware of the size of the shortfall in our current savings to finance future pension provision. It will report next year on various ways of closing this funding gap. But there is an additional and fundamental question of how to design a pension system which shares risk optimally among individuals and generations.

---

[16] Robert Rubin, former US Treasury Secretary in the Clinton Administration, has stressed the importance of probabilistic thinking in policy-making (Rubin and Weisberg, 2003).
[17] Letter to Frederick the Great, April 16, 1767.

Longevity risks are both individual and collective in nature. Collective risks raise difficult issues of fairness between generations. Individual risks are easier to handle.

In theory, individual longevity risks can be pooled through private annuity markets. As long as the provider of the annuity, typically an insurance company, knows how long we will live on <u>average</u>, then additional payments to the long-lived will be exactly offset by fewer payments to the short-lived. By pooling risks, annuity providers can offer insurance against longevity. But to offer a single annuity is risky. In *Sense and Sensibility*, Fanny Dashwood explains to her husband the drawbacks of providing an annuity to his family: 'if you observe, people always live forever when there is any annuity to be paid them. An annuity is a very serious business; it comes over and over every year, and there is no getting rid of it'.[18]

Only a small fraction of savings are invested in annuities.[19] One explanation for this is that annuity providers find it difficult to distinguish the healthy from the sick *ex ante*, the former purchasing more annuities than the latter, leading insurance companies to offer annuities on less attractive terms thus discouraging their purchase.[20] This 'adverse selection' is important – it is approximately 20% more expensive to buy annuities as an individual than in a group scheme.[21]

Collective risks raise additional problems. Longevity risk to a particular cohort can be shared only across generations. Those generations who are alive today can pool risks through financial markets. But sharing across all generations, including the unborn, requires collective or public insurance through variations in the national debt.[22]

---

[18] Fanny Dashwood is claiming that annuities lead to 'moral hazard' to use the modern jargon: the provision of an annuity changes behaviour in a way that increases life expectancy.

[19] See Yaari (1965) and Davidoff *et al*. (2003). For a discussion of pension arrangements in the US, see Diamond (2003, 2004).

[20] For a theoretical study of annuity provision under adverse selection, see Brugiavini (1993).

[21] Finkelstein and Poterba (2004), for instance, note that 65 year old men who voluntarily buy annuities live about 20% longer than other men. Pauly and Zeng (2003) found that in the United States adverse selection made it impossible for a private market in prescription drug insurance to function.

[22] It is possible that families, or more accurately dynasties, can provide such insurance themselves through changes in the bequest which they make to their descendants: passing on a larger bequest if the current

Consider an annuity provider who faces the risk that an increase in the average longevity of its pool of annuitants will lead to a shortfall in its funds. The annuity provider would clearly like to hedge this risk in some way. One way of doing that would be for it to have assets on its balance sheet that pay out more if the average longevity of its annuitants increases. Financial markets are currently developing 'longevity bonds' along these lines.[23]

The interest rate on longevity bonds is linked to some measure of unexpected changes in average life expectancy. In the case of the recently announced European Investment Bank-BNP Paribas bond, the payout each year for the next 25 years will be a fixed sum multiplied by the proportion of people in England and Wales who were 65 in 2003 that are still living in that particular year. So if five years from now 90% of the cohort are still alive, the payment in that year will be 90% of the fixed sum. The cash flows to a buyer of this bond will decline over time, but decline less rapidly if people live longer.

The demand for such a hedge against longevity risk we might expect to come from pension funds and other providers of annuities. The supply might come from life-insurance companies, who benefit from increased longevity to the extent that fewer early deaths mean lower claims, and financial intermediaries collecting the savings of those younger investors prepared to bear some of the longevity risk of the old.[24]

How close private markets can get to providing the optimal amount of risk sharing is hard to say. Private annuity markets which could pool individual longevity risk are small. For collective longevity risks, the burden of unexpectedly high longevity for a particular cohort should be spread over as many generations as possible. In theory, this provides a potential role for government to share risks across generations.

---

generation dies earlier than expected and a smaller bequest if it survives for longer than expected. But only a small proportion of bequests seems large enough to act as a buffer in this way.

[23] As far as I am aware, the first instrument offering a return specifically to aggregate longevity was issued by Swiss Re, a large reinsurance company, in 2003.

[24] To the extent that insurance companies offer both life insurance and annuities, they may be able partially to hedge the risks.

There are two ways in which government involvement can, in principle, improve risk-sharing. First, collective schemes, whether compulsory or employer-based, avoid the costs of adverse selection associated with individual provision. Second, by combining taxes on current employees with deficit finance (i.e. taxes on future employees), state-funded pensions can transfer risks across generations in ways private markets cannot. Alternatively, the government could issue its own longevity bonds. Private annuity providers could use them to hedge aggregate longevity risk. By assuming longevity risk the government would make it possible for the private sector to purchase such bonds and in turn support a private market in annuities.

What has happened in Britain over the past few years has been that the impossibility of obtaining longevity risk insurance has been a contributory factor to the sharp decline in private sector provision of defined benefit pensions. Over the past ten years, membership of private sector defined benefit pension schemes that are still open is estimated to have fallen by 60% (Pensions Commission, 2004). It is not clear why the market has not developed an annuity or pension product indexed to aggregate longevity which would allow financial institutions to pass the collective longevity risk onto the policy-holders. What is clear is that individuals have now been asked to take on more and more risk because the absence of annuities means that they now bear both individual and collective aspects of longevity risk. There is something to be said for the view that longevity risk is more efficiently shared collectively, leaving the private sector to provide insurance against individual mortality. The question of who should issue and who would buy longevity bonds merits further study.[25]

I have so far discussed only longevity risk. Clearly, any discussion of pension reform must include also the other key collective risks: to productivity, rates of return and fertility among others. We do not, of course, start with a clean slate. Existing government institutions – whether by design or default – already have important risk-

---

[25] The question of who should bear collective longevity risk has been discussed by Blake and Burrows (2001), Blake (2003), Dowd (2003), Cox and Lin (2004a, 2004b), Willets (2004), Auerbach and Hassett (2002), among others.

sharing properties. Means-tested pensions and other benefits, as well as the National Health Service, already give the government a major role in the sharing of risks.

But in reforming our pension system, it is important to separate two issues. First, how should we pay for the cost of present pension commitments? Second, what is the right structure for our pension system: which risks should be borne individually and which collectively? A debate is needed over whether present arrangements imply too much or too little sharing of risk.

We should remember also that insurance is based on ignorance – ignorance of the outcomes facing each individual. But knowledge can turn risk into certainty. The advent of genetic testing could undermine a private market in insurance against longevity and ill-health. Even if insurance companies could not require testing as a condition of insurance, the fact that tests were available could increase even further the degree of adverse selection in such insurance markets. Pooling risks caused by unequal genetic inheritance might no longer be feasible. Insurance would become instead a question about redistribution, and hence a matter for public policy.

6.   Conclusions

In this lecture I have stressed that most public policy decisions are a matter of the balance of risks. There is no certainty. Indeed, it is the illusion of certainty that undermines much public debate. As Bertrand Russell said, 'The whole problem of the world is that fools and fanatics are always so certain of themselves, but wiser people so full of doubts'.

Recognising and modelling risks means understanding the limits to our present knowledge. Communicating risk is about transparency. When information on risks is provided to the public, it is often in a form that is hard to assess. Thinking carefully about communication is important to the level of public debate. The quantity of information is less relevant than its quality – sometimes less is more. And managing risk

means making choices – collective choices – on the basis of a rational shared assessment of the risks involved.

We cannot avoid uncertainty. So let us face up to it.

REFERENCES

Abowd, John and Card, David (1989), 'On the Covariance Structure of Earnings and Hours Changes', *Econometrica*, 57, 411-445.

Auerbach, Alan and Hassett, Kevin (2002), 'Optimal Long-Run Fiscal Policy: Constraints, Preferences and the Resolution of Uncertainty', mimeo, University of California, Berkeley.

Bernstein, Peter (1996), *Against the Gods*, John Wiley & Sons, New York.

Benjamin, Bernard (1982), 'The Span of Life', *Journal of the Institute of Actuaries*, 109, pp319-340.

Best, Joel (2001), *Damned Lies and Statistics*, University of California Press, Berkeley.

Blake, David and Burrows, William (2001), 'Survivor Bonds: Helping to Hedge Mortality Risk', *Journal of Risk and Insurance*, 68(2), 339-348.

Blake, David (2003), 'Reply to "Survivor Bonds: a Comment on Blake and Burrows"', *Journal of Risk and Insurance*, 70(2), 349-351.

Briscoe, Simon (2000), *Interpreting the Economy*, Penguin Books, London.

Brugiavini, Agar (1993), 'Uncertainty Resolution and the Timing of Annuity Purchases', *Journal of Public Economics*, 50, 31-62.

Caselli, Graziella and Vallin, Jacques (2001), 'Demographic Trends: Beyond the Limits?', *Population: An English Selection*, Vol. 13, No. 1, Biodemographic Perspectives on Human Longevity, 41-71.

Cox, Samuel and Lin, Yijia (2004a), 'Securitisation of Mortality Risks in Life Annuities', mimeo, Georgia State University.

Cox, Samuel and Lin, Yijia (2004b), 'Natural Hedging of Life Insurance and Annuity Mortality Risks', mimeo, Georgia State University.

Davidoff, Thomas, Brown, Jeffrey and Diamond, Peter (2003), 'Annuities and Welfare', mimeo, Massachusetts Institute of Technology.

Dawid, Philip (2004), 'Probability and Statistics in the Law', Research Report No.243, Department of Statistical Science, University College, London, mimeo.

Diamond, Peter (2004), 'Social Security', *American Economic Review*, 94, 1-24.

Diamond, Peter (2003), *Taxation, Incomplete Markets, and Social Security*, MIT Press, Cambridge, Massachusetts.

Dowd, Kevin (2003), 'Survivor Bonds: A Comment on Blake and Burrows', *Journal of Risk and Insurance*, 70(2), 339-348.

Eggleston, Richard (1978), *Evidence, Proof and Probability*, Butterworths, London.

Engle, Robert (2004), 'Risk and Volatility: Econometric Models and Financial Practice', *American Economic Review*, 94, 405-420.

Enstrom, James (1989), 'Health Practices and Cancer Mortality Among Active California Mormans', *Journal of the National Cancer Institute*, 81 (23), 1807-1814.

Finkelstein, Amy and Poterba, James (2004), 'Adverse Selection in Insurance Markets: Policyholder Evidence from the U.K. Annuity Market', *Journal of Political Economy*, 112, 183-208.

Gigerenzer, Gerd (2002), *Reckoning with Risk*, Penguin Books, London.

Granger, Clive (2004), 'Time Series Analysis, Cointegration, and Applications', *American Economic Review*, 94, 421-425.

Hacking, Ian (1975), *The Emergence of Probability*, Cambridge University Press, Cambridge.

Hill, Ian (2002), 'Why Sally Clark is, Probably, Innocent', mimeo, University of Salford.

Hill, Ian (2004), 'Multiple Sudden Infant Deaths – Coincidence or Beyond Coincidence', *Paediatric and Perinatal Epidemiology*, 18, 320-326.

Kannisto, Väiniö, Lauritsen, Jens, Thatcher, A. Roger and Vaupel, James W. (1994), 'Reductions in Mortality at Advanced ages: Several Decades of Evidence from 27 Countries', *Population and Development Review*, 20(4), 793-810.

Knight, Frank (1921), *Risk, Uncertainty and Profit*, Houghton, Boston.

MaCurdy, Thomas (1982), 'The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis', *Journal of Econometrics*, 18, 83-114.

Manton, Kenneth G., Stallard, Eric and Tolley, H. Dennis (1991), 'Limits to Human Life Expectancy: Evidence, Prospects, and Implications', *Population and Development Review*, 17(4), 603-637.

May, Robert (2003), 'The Use of Scientific Advice in Making Public Policy', in ed. Rüegg (2003).

May, Robert (2004), 'Uses and Abuses of Mathematics in Biology', *Science*, 303, 790-793.

May, Robert, Gupta, Sunetra and McLean, Angela (2001), 'Infectious Disease Dynamics: What Characterizes a Successful Invader?', *Philosophical Transactions, Royal Society*, 356, 901-910.

Meghir, Costas and Pistaferri, Luigi (2004), 'Income Variance Dynamics and Heterogeneity', *Econometrica*, 72(1), 1-32.

Oeppen, Jim and Vaupel, James (2002) 'Broken Limits to Life Expectancy', *Science*, 296, no 10 May, pp1029-1031.

Olshansky, S.J., Carnes, B.A. and Cassell, C. (1990), 'In Search of Methuselah: Estimating the Upper Limits to Human Longevity', *Science*, 250, 634-640.

Pauly, Mark and Zeng, Yuhui (2003), 'Adverse Selection and the Challenge to Stand-Alone Prescription Drug Insurance', National Bureau of Economic Research Working Paper 9919, Cambridge, Massachusetts.

Pensions Commission (2004), *Pensions: Challenges and Choices*, The Stationery Office, London.

Peterman, T.A. *et al.* (1988), 'Risk of Human Immunodeficiency Virus Transmission From Heterosexual Adults With Transfusion-Associated Infections", *Journal of the 'American Medical Association*, 259, 55-58.

Poterba, James (2004), 'Annuities in Early Modern Europe', forthcoming in W. Goetzmann and K. Geert Rouwenhorts (eds.) *Of Interest and Enterprise: The History of Financial Innovation*, Oxford University Press, New York.

Rubin, Robert and Weisberg, Jacob (2003), *In an Uncertain World: Tough Choices from Wall Street to Washington*, Random House, New York.

Rüegg, Walter (2003), *Meeting the Challenges of the Future*, Balzan Symposium 2002, Leo S. Olschki, London

Taleb, Nassim Nicholas (2001), *Fooled by Randomness*, Texere, New York.

Thompson, W.C. and Schuman, E.L. (1987), 'Interpretation of Statistical Evidence in Criminal trials: the prosecutor's Fallacy and the Defense Attorney's Fallacy', *Law and Human Behaviour*, 11, 167-187.
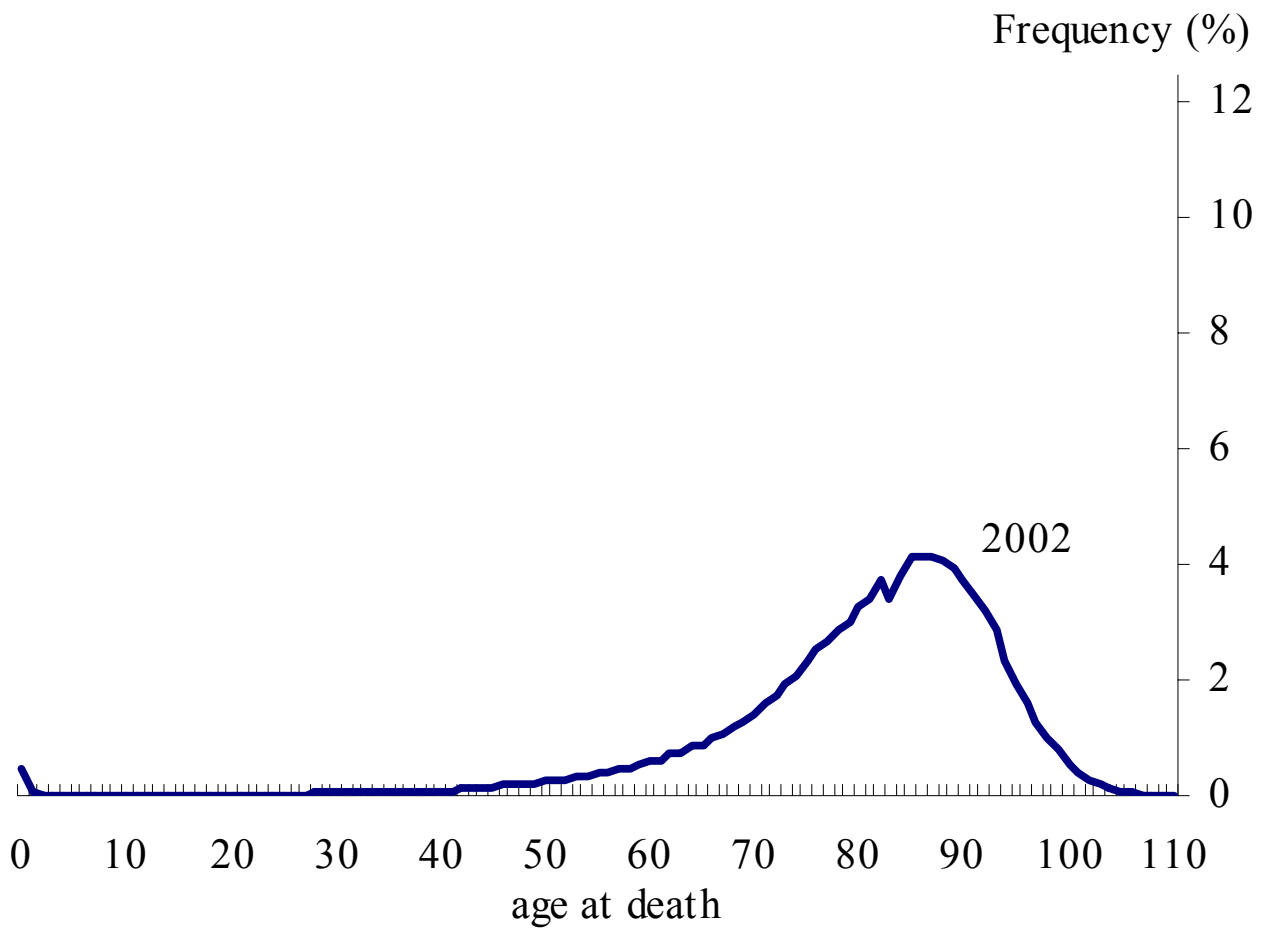
Tufte, Edward (1983), *The Visual Display of Quantitative Information*, Graphics Press, Cheshire, Connecticut.

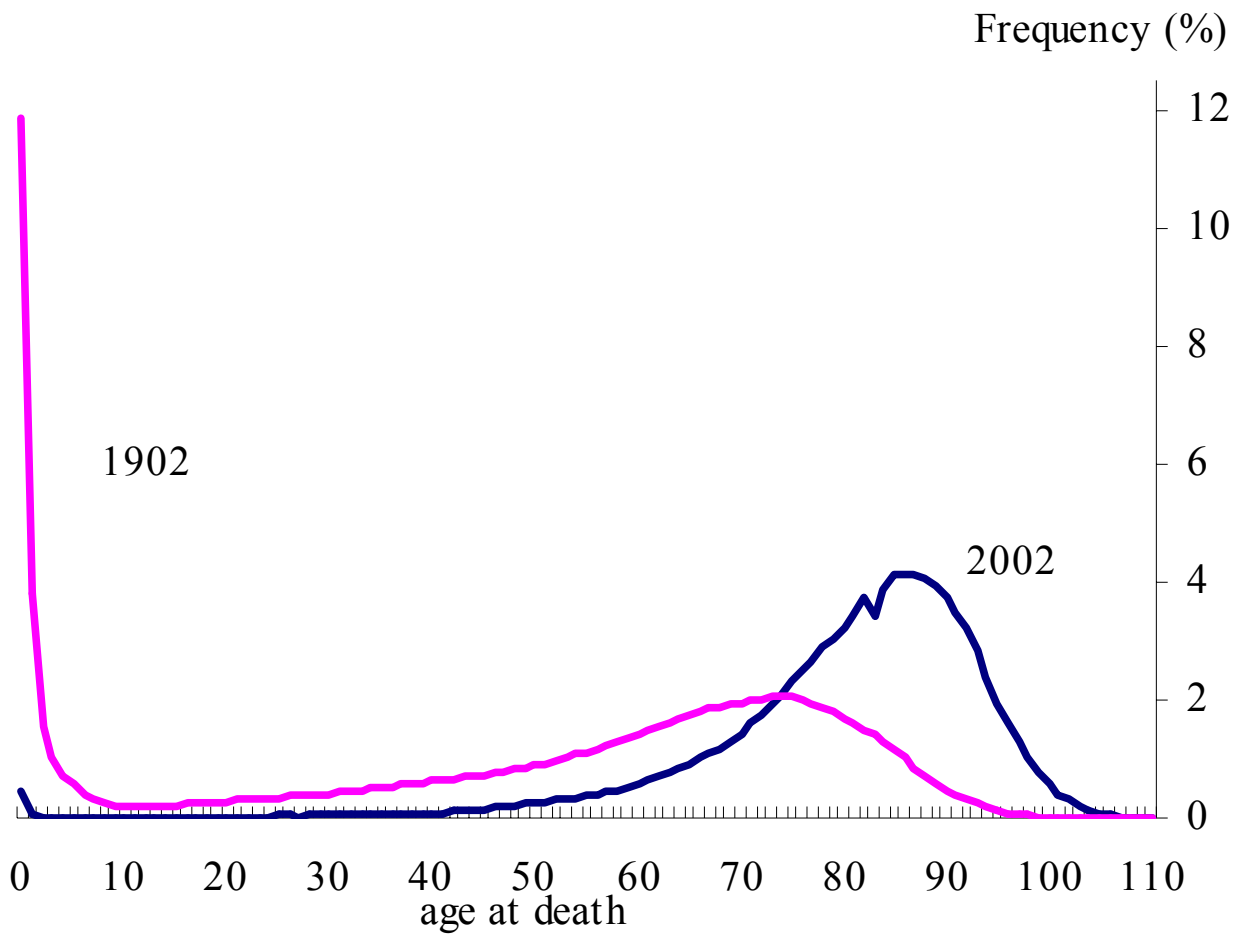Tufte, Edward (1990), *Envisioning Information*, Graphics Press, Cheshire, Connecticut.

Willets, David (2004), *The Pensions Crisis: What it Means and What to do About it*, mimeo.

Wilmoth, John R. and Horiuchi, Shiro (1999), 'Rectangularization Revisited: Variability of Age at Death within Human Populations', *Demography*, 36(4), 475-495.

Yaari, Menahem E. (1965), 'Uncertain Lifetimes, Life Insurance, and the Theory of the Consumer', *Review of Economic Studies*, 32, 137-150.
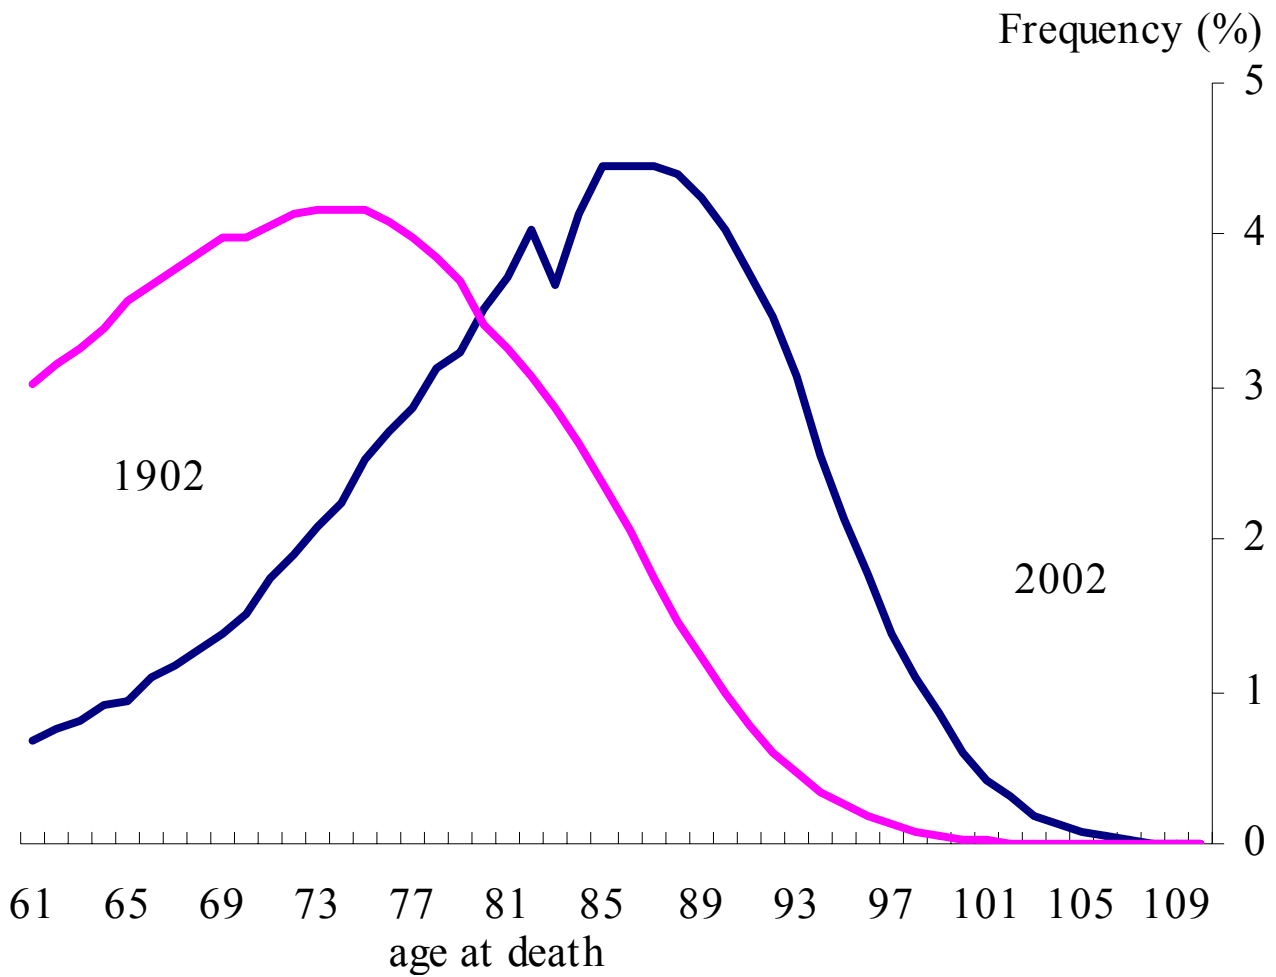
Chart 1: Mortality, England and Wales, 2002

Frequency (%)

2002

age at death

Source: Government Actuary's Department. Data for women.

Chart 2: Mortality, England and Wales, 1902 and 2002
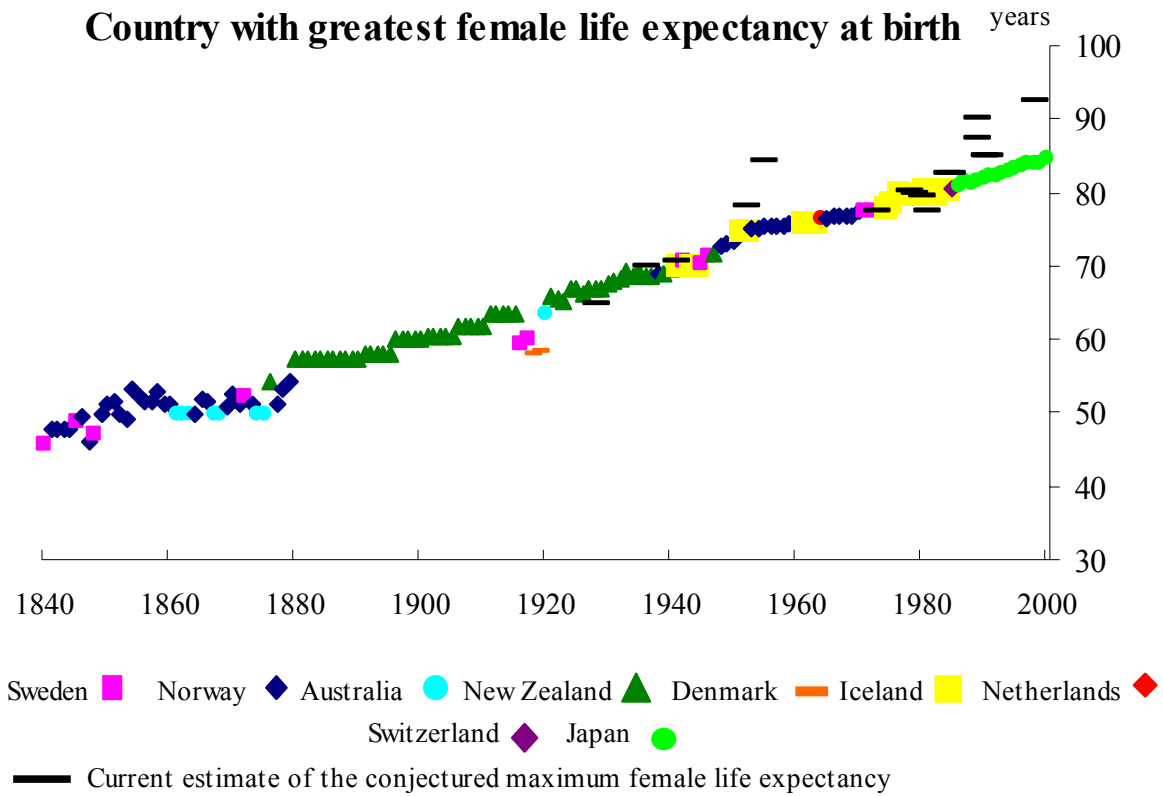
Frequency (%)

1902

2002

age at death

Source: Government Actuary's Department. Data for women.

Chart 3: Mortality after 60, England and Wales, 1902 and 2002

Frequency (%)
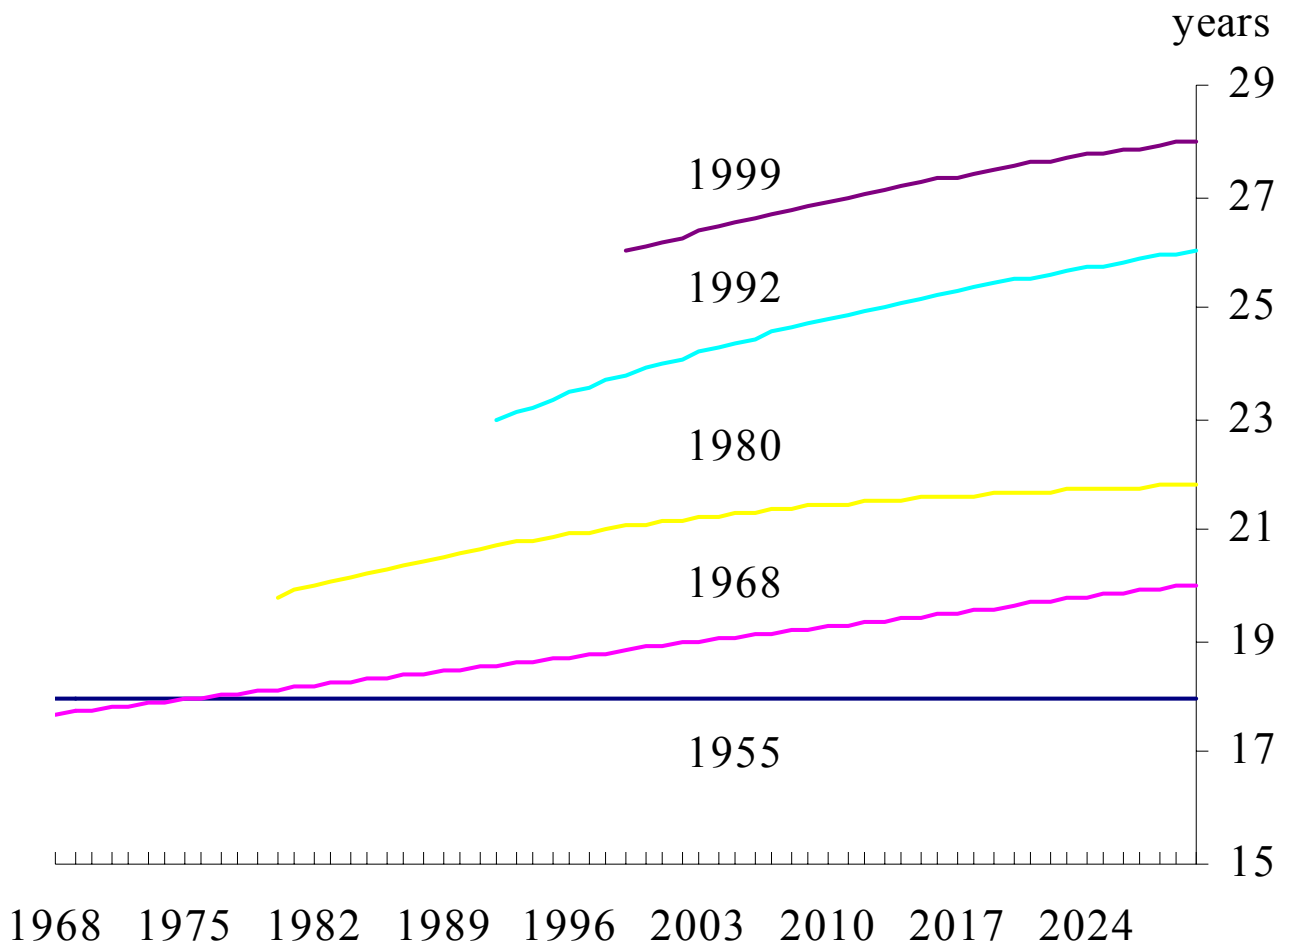

Source: Government Actuary's Department. Data for women.

Chart 4: Estimates of the limit to the female lifespan
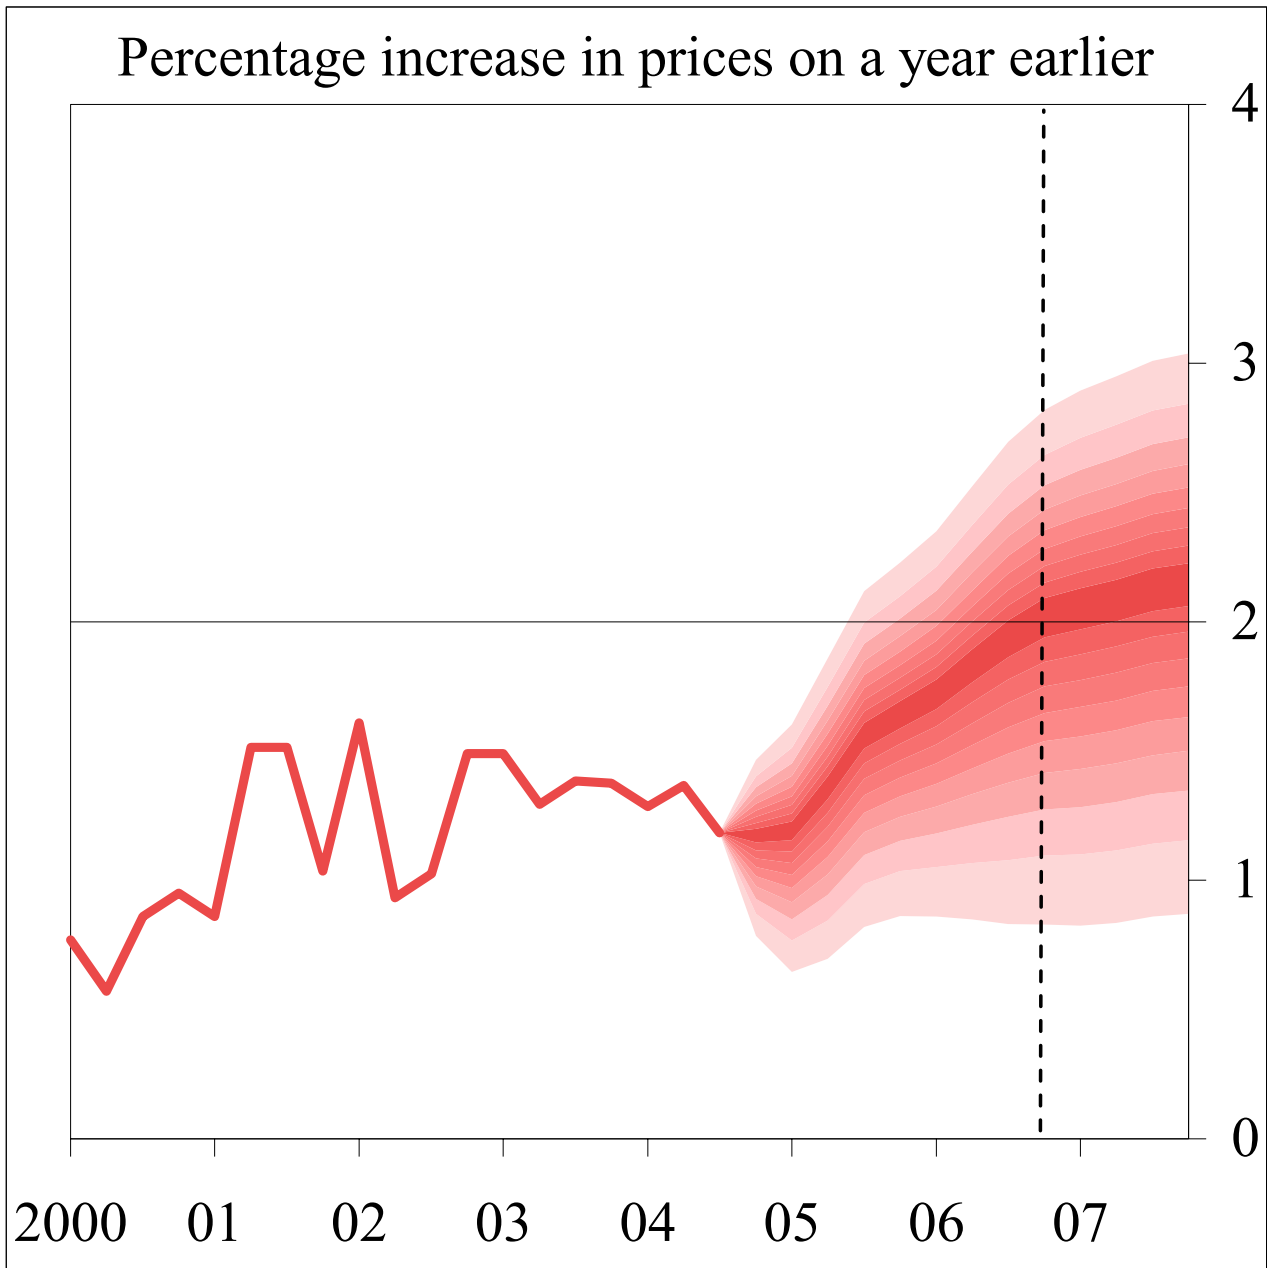
Country with greatest female life expectancy at birth

Source: Reproduced from Oeppen and Vaupel (2002), p1029.

Chart 5: Actuarial Profession projections of male life after 60

years

1999

1992

1980

1968

1955

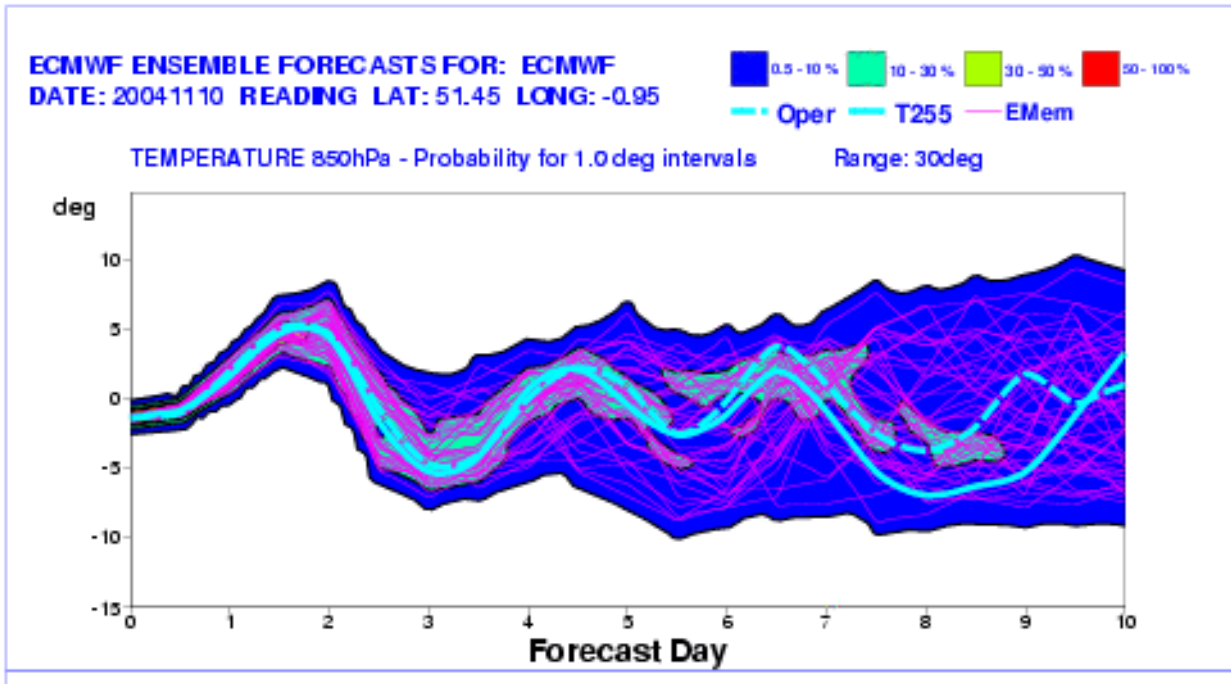1968　1975　1982　1989　1996　2003　2010　2017　2024

Source: Continuous Mortality Investigation, Actuarial Profession. Data are for United Kingdom policy holders of Actuarial Profession members.

Chart 6: Current CPI inflation projection based on market interest rate expectations

Percentage increase in prices on a year earlier

Source: The Bank of England's *Inflation Report* (November 2004), p42.

Chart 7: Temperature forecast on 10 November 2004

ECMWF ENSEMBLE FORECASTS FOR: ECMWF
DATE: 20041110 READING LAT: 51.45 LONG: -0.95

0.5 - 10 %   10 - 30 %   30 - 50 %   50 - 100 %
Oper   T255   EMem

TEMPERATURE 850hPa - Probability for 1.0 deg intervals   Range: 30deg

deg

Forecast Day

Source: Met Office.

Chart 8: Life expectancy fanchart

Years

Source: Bank of England calculations using Government Actuary's Department Data.

Chart 9: March on Moscow

CARTE FIGURATIVE des pertes successives en hommes de l'Armée Française dans la campagne de Russie 1812–1813.

Dressée par M.Minard, Inspecteur Général des Ponts et Chaussées en retraite.

TABLEAU GRAPHIQUE de la température en degrés du thermomètre de Réaumur au dessous de zéro

Xbre = December   9bre = November   8bre = October

Source: Tufte (1983).