

O'Dea, Cormac; Sturrock, David

**Working Paper**

## Subjective expectations of survival and economic behaviour

IFS Working Papers, No. W18/14

**Provided in Cooperation with:**

Institute for Fiscal Studies (IFS), London

*Suggested Citation:* O'Dea, Cormac; Sturrock, David (2018) : Subjective expectations of survival and economic behaviour, IFS Working Papers, No. W18/14, Institute for Fiscal Studies (IFS), London,  
<https://doi.org/10.1920/wp.ifs.2018.W1814>

This Version is available at:

<https://hdl.handle.net/10419/200303>

**Standard-Nutzungsbedingungen:**

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

**Terms of use:**

*Documents in EconStor may be saved and copied for your personal and scholarly purposes.*

*You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.*

*If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.*

# Subjective expectations of survival and economic behaviour

**IFS Working Paper W18/14**  
Cormac O'Dea  
David Sturrock

# Subjective expectations of survival and economic behaviour\*

Cormac O’Dea<sup>†</sup> and David Sturrock<sup>‡</sup>

16 April 2018

*Copy edited by Judith Payne*

## Abstract

This paper investigates individuals’ expectations about their own survival to older ages and compares patterns in average responses about survival chances with actual and projected survival rates. The extent to which individuals have, on average, accurate expectations about survival to older ages is important in a context of increasing personal responsibility for and control over the accumulation and use of savings for retirement.

Keywords: Survival expectations, savings, wealth accumulation, annuitisation

JEL codes: D14, D84, D91, J14

---

\* This paper was funded by the IFS Retirement Savings Consortium and the Economic and Social Research Council (ESRC) through a Knowledge Exchange Grant and by an ESRC Secondary Data Initiative grant (reference ES/N011872/1). The IFS Retirement Savings Consortium comprises Age UK, Association of British Insurers, Chartered Insurance Institute, Department for Work and Pensions, HM Revenue and Customs, HM Treasury, Investment Association, Legal and General Investment Management, Money Advice Service, and Tax Incentivised Savings Association. This work would also not have been possible without support from the ESRC-funded Centre for the Microeconomic Analysis of Public Policy (CPP) at IFS (grant reference ES/M010147/1). The authors are grateful to Rowena Crawford, Carl Emmerson, Gemma Tetlow and members of the IFS Retirement Savings Consortium for helpful comments. Data from the English Longitudinal Study of Ageing (ELSA) were made available through the UK Data Archive (UKDA). ELSA was developed by a team of researchers based at the National Centre for Social Research, University College London and the Institute for Fiscal Studies. The data were collected by the National Centre for Social Research. The data creators, depositors, copyright holders and funders bear no responsibility for the analysis or interpretation of the data presented here. All errors are those of the authors. Correspondence to [david\\_s@ifs.org.uk](mailto:david_s@ifs.org.uk).

<sup>†</sup> Yale University and Institute for Fiscal Studies

<sup>‡</sup> Institute for Fiscal Studies and University College London

# Contents

<b>Executive summary</b>	<b>3</b>
<b>1. Introduction</b>	<b>6</b>
<b>2. Subjective survival expectations</b>	<b>7</b>
2.1 What are survival expectations?	7
2.2 Data on survival expectations	8
2.3 Do individuals give well-formed answers to survival questions?	9
2.4 What information is contained in individuals' answers to survival questions?	11
<b>3. Patterns in subjective survival expectations</b>	<b>17</b>
3.1 Comparison of 'subjective' expectations and official estimates of survival probabilities	18
3.2 Constructing subjective survival curves	20
3.3 Comparison of subjective and objective survival curves for subgroups	23
<b>4. Subjective expectations and economic behaviour</b>	<b>29</b>
4.1 Annuity choice	30
4.2 Deferral of the state pension	33
<b>5. Conclusions: implications for annuities markets and for policy</b>	<b>35</b>
<b>Appendix</b>	<b>37</b>
<b>References</b>	<b>41</b>

---

# Executive summary

This paper investigates individuals' expectations about their own survival to older ages and compares patterns in average responses about survival chances with actual and projected survival rates. The extent to which individuals have, on average, accurate expectations about survival to older ages is important in a context of increasing personal responsibility for and control over the accumulation and use of savings for retirement.

## Subjective survival expectations

---

**Modern surveys ask individuals about their probability of survival to specific older ages. In only a small proportion of cases is there evidence that these questions are not well understood.**

98% of individuals gave an answer to a question asking their chances of surviving to older ages and of these just 14% showed clear evidence of misunderstanding (e.g. by reporting no chance of death in the coming 10-year period).

---

**Individuals' stated beliefs about their probability of survival are correlated with known risk factors such as smoking and the age that their parents died.**

Current smokers report on average 6–8 percentage points lower probability of surviving to an age 11–15 years ahead than do non-smokers.

---

**Beliefs about probability of survival are also correlated with the individual's actual age of death and respond to new diagnoses of health conditions.**

A new cancer diagnosis was associated with a 5 percentage point reduction in the stated probability of surviving to an age 11–15 years ahead.

---

## Comparing subjective expectations with 'objective' data

---

**Individuals from a range of ages and birth cohorts underestimate their chances of survival to ages 75, 80 and 85, on average.**

Those in their 50s underestimate their chances of survival to age 75 by around 20 percentage points and to 85 by around 10 percentage points.

---

---

---

**Individuals in their late 70s and 80s are, on average, mildly optimistic about surviving to ages 90, 95 and above.**

This optimism becomes larger at older ages (10–15 percentage points when looking at age 95) and is larger for men than for women, amongst those born in the 1920s and 1930s.

---

**Subjective survival curves, estimated using individuals' stated survival expectations, capture the general patterns in expectations.**

These show growing pessimism relative to official life tables for ages up to the mid 70s before turning to mild optimism from the late 90s onwards.

---

**Actual survival probabilities differ significantly according to individuals' education, wealth and marital status.**

Women of 60 in the bottom household wealth quintile had a 65% chance of surviving to age 80, compared with 87% for those in the top quintile, based on mortality data.

---

**Subjective survival curves reflect differences in actual mortality rates between groups to varying degrees.**

Widowed men and women aged 60 show the greatest survival pessimism, underestimating their probability of survival to age 80 or above by almost 30 percentage points on average.

---

## **Subjective expectations and economic behaviour**

---

**Survival pessimism is a potential driver of the unpopularity of annuities.**

Around 65% of individuals could perceive an annuity priced according to average survival chances as offering an unfairly low income.

---

**Deferral of the state pension, a choice analogous to annuity purchase, is rarely taken up despite being offered at a favourable rate.**

While individuals are roughly twice as likely to defer the state pension if it represents a 'fair deal' given their survival expectations, the overall level of deferral is sufficiently low that we cannot make strong statements about this relationship.

---

---

---

**The divergence between subjective and objective survival expectations due to survival pessimism may have important implications for behaviour and policy.**

As individuals are given more control over saving for retirement and use of accumulated wealth, survival pessimism may mean individuals are more at risk of having insufficient retirement resources.

---

---

# 1. Introduction

In the UK, individuals have increasing levels of choice and responsibility in relation to providing for their retirement. The elimination of any earnings-related component in the state pension, the decline of occupational defined benefit pensions in the private sector and the introduction of 'pension freedoms' for those with defined benefit pensions have changed the UK pensions landscape. To avoid a drop-off in income at retirement, individuals will increasingly have to decide how to accumulate savings and how to use these in retirement.

One important component in decisions around savings and the use of resources in retirement is individuals' expectations about how long they are going to live. Beliefs about longevity will shape the amount that individuals think they need to save and the manner in which they use savings once in retirement. Much of economic and policy analysis in this area proceeds on the assumption that individuals are well informed about the chances they face of survival to older ages. This paper seeks to examine in detail the nature of individuals' expectations about survival to older ages in order to determine whether individuals have well-formed expectations about their own future survival chances and whether these are 'accurate' – at least on average – or show systematic biases. We then explore the implications of individuals' expectations for economic behaviour.

This paper draws on data from the English Longitudinal Study of Ageing (ELSA). Section 2 describes the available information about individual subjective mortality expectations, exploring whether or not individuals have well-formed expectations and the extent to which expectations respond to factors relevant to mortality. Section 3 then explores the divergence between average subjective expectations and the 'objective' mortality data given by official life tables and mortality rates in ELSA, looking at variation across sexes and according to education, wealth and marital status. Section 4 considers the implications for economic behaviour of the patterns in subjective expectations, focusing on annuitisation and deferral of the state pension. In conclusion, Section 5 discusses the implications for annuities markets and for policy.



---

## 2. Subjective survival expectations

### Key findings

**Modern surveys ask individuals about their probability of survival to specific older ages. In only a small proportion of cases is there evidence that these questions are not well understood.**

98% of individuals gave an answer to a question asking their chances of surviving to older ages and of these just 14% showed clear evidence of misunderstanding (e.g. by reporting no chance of death in the coming 10-year period).

**Individuals' stated beliefs about their probability of survival are correlated with known risk factors such as smoking and the age that their parents died.**

Current smokers report on average 6–8 percentage points lower probability of surviving to an age 11–15 years ahead than do non-smokers.

**Beliefs about probability of survival are also correlated with the individual's actual age of death and respond to new diagnoses of health conditions.**

A new cancer diagnosis was associated with a 5 percentage point reduction in the stated probability of surviving to an age 11–15 years ahead.

### 2.1 What are survival expectations?

Many different economic actors and analysts may have expectations about the chances of certain individuals, or groups of individuals, living to specific ages. In this paper, we examine the expectations of individuals about *their own* chances of survival to older ages. We now set out and examine the information we have about these individual expectations.

We will ask what information we actually obtain when we ask individuals about their chances of surviving to certain ages. We may wonder whether or not individuals have well-formed survival expectations. This concern could be about whether individuals understand these sorts of probabilistic objects at a conceptual level. We might also ask whether, given this understanding, individuals actually have a well-formed, reflective judgement of their personal situation. We will explore these questions in this section and provide evidence that suggests that, for most individuals, their reports about survival expectations do represent a well-formed answer to a question they understand. Of course, an individual having a well-formed expectation of their survival probability does not preclude them being inaccurate – either due to cognitive biases or poor information. We will return to this issue in Section 3.

There are a number of key concepts in the context of survival and mortality, including mortality rates, survival probabilities and life expectancy. These are set out in detail in Section A.1 of the appendix. The concept that we will employ most often in this paper is a survival probability. This is a number that tells us the probability that a certain individual (or type of individual) of a certain age will survive to some older age. It is possible to summarise a full set of survival probabilities for someone of a particular age using a survival curve. Survival curves slope downwards, tracing out the declining probability of survival to each successively older age.

## 2.2 Data on survival expectations


The data that we use on individuals' expectations come from the English Longitudinal Study of Ageing (ELSA). This is a survey carried out every two years which covers a representative sample of the English household population aged 50 and over. We use waves 1 to 7, which cover the years from 2002–03 to 2014–15. Where possible, individuals are followed from one 'wave' of the survey to the next, so that we can see what happens to them over time. ELSA contains a wide range of information about individuals' demographic characteristics, financial situation, health status, activities and attitudes.

As part of ELSA, individuals are asked questions about how likely they think it is that certain events will happen in the future. These events include whether or not the individual will still be working at a certain older age and whether or not they will leave an inheritance. Also included as part of this expectations section of the survey are questions about the probability that an individual survives to certain older ages.

At the beginning of the section on expectations, the individual is shown a show card, as per Figure 2.1, and read the following statement:

*Now I have some questions about how likely you think various events might be. When I ask a question I'd like you to give me a number from 0 to 100, where 0 means that you think there is absolutely no chance an event will happen, and 100 means that you think the event is absolutely certain to happen.*

**Figure 2.1. Show card from expectations section of ELSA survey**



Source: ELSA wave 7 questionnaire.

The questions on survival expectations take the form: 'What are the chances that you will live to be age X or more?'. The exact age asked about depends on the age of the individual being interviewed. Table 2.1 details the questions asked of individuals of different ages. The questions are such that an individual is asked about the probability of survival to an age that is between 11 and 15 years ahead of their current age. From the third wave of

ELSA onwards, individuals aged below 70 were also asked a second question about the probability that they would survive to age 85 or more.

**Table 2.1. Survival questions asked of individuals**

Age of respondent	Age asked about in first question	Age asked about in second question (wave 3 onwards)
65 and under	75	85
66–69	80	85
70–74	85	-
75–79	90	-
80–84	95	-
85–89	100	-

Source: ELSA waves 1–7.

As part of the analysis in this paper, we show results for different subgroups of individuals, splitting individuals by their position in the wealth distribution, their marital status or their level of education. In the case of education, we use a categorisation based on the age that the individual finished full-time education. The first category includes all those who left education up to and including the school-leaving age for someone of their birth cohort (47% of the sample), the second category includes those who left above the compulsory age and below the age of 19 (37% of the sample), and the third educational category includes all those who remained in full-time education to the age of 19 or above (16% of the sample).

### 2.3 Do individuals give well-formed answers to survival questions?

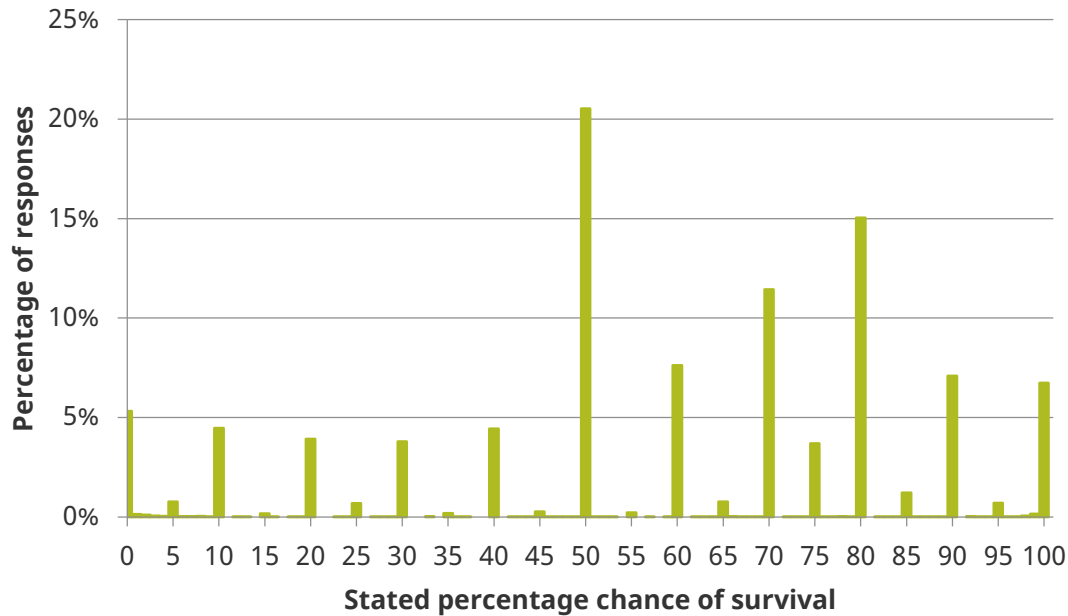
Much of economic analysis in relation to savings, retirement and other decisions at the end of life assumes that individuals have well-formed beliefs about the probability they will die at each given age in the future (e.g. Hurd, 1989; French, 2005) and that these beliefs are instrumental in driving decision-making in these areas. We may, however, have reservations about whether individuals have well-formed beliefs about survival. Before treating our data as a representation of individuals' beliefs, we try to test whether their beliefs about survival are well formed.

Figure 2.2 shows the distribution of answers given to the survival questions.<sup>1</sup> We see that individuals tend to give an answer corresponding to a multiple of 5 or 10. These 'round number' answers are perhaps not surprising as individuals might see a choice of a non-round number as suggesting spurious precision. 7% of answers given are '100%'. Given that individuals are being asked about an age 11–15 years ahead, it must be the case that there is at least some chance that the individual will die (even due to an accident, for example) during that period, and so we can be fairly confident that these answers reflect a

<sup>1</sup> Note that 1.5% of individuals answer 'don't know' to survival questions. Older individuals are more likely to answer 'don't know'. For example, 1% of those aged 60–69 answer 'don't know', compared with 5% of those aged 85 and above.

lack of understanding of the question, or of probability, in some way. It is perhaps not fair to draw the same conclusion about those who answer '0%', as they may – due to terminal illness, for example – reasonably answer that there is absolutely no chance that they will survive to some older age.

**Figure 2.2. Distribution of answers to survival questions**



Source: ELSA waves 1–7. 66,210 answers from 16,203 unique individuals.

There is a second group of individuals who show a lack of understanding of the nature of the survival questions. These are the individuals who are asked two survival questions and say that they have a higher chance of surviving to some older age than they do of surviving to a younger age. This is clearly impossible because in every case where an individual survives to 85, for example, they must also have survived to 75.

In total, of those who answer the survival questions, 14% of individuals give an answer that is 'impossible' for one of the two reasons identified. When conducting our analysis, we exclude these individuals from our sample on the basis that they have not given a well-formed answer and so appear not to fully understand the questions or concepts involved. Section A.2 of the appendix details the distribution of these individuals by age and education.<sup>2</sup>

Considering Figure 2.2 again, we might be concerned about the 'spike' of 21% of responses being '50%'. We might worry that some individuals do not understand what is being asked of them and so simply say '50%' as it is a focal answer. We assess whether this response reflects general lack of understanding of probability by looking at the distribution of responses to other probability questions given by those who answered '50%' to one or both survival questions. We see that individuals who answer '50%' almost always give a range of answers to other questions and are no more likely to answer '50%' to other probability questions than are the general sample. Figure A.1 in the appendix

<sup>2</sup> Differences across education groups are not significant. While there are differences according to age, our analysis is conditional on age and so this is not of concern.

---

compares the distribution of the number of '50%' answers for those who answered '50%' to survival questions with that for all individuals. Of the 16,778 individuals who are asked one or more survival questions, only 41 individuals (0.2%) answered '50%' to all survival questions in all waves. Given this, we choose to keep in our sample those cases where individuals answer '50%'. Note that the results in the remainder of this paper are unchanged, qualitatively, if we instead exclude answers of '50%', and so our analysis does not hinge in any way on our decision here.

In summary, once we have removed from our sample those individuals who give answers that are 'impossible', we are left with the vast majority of individuals and, furthermore, a sample of individuals who show no sign of failing to understand the questions being asked (even if they do not necessarily give accurate answers).

## **2.4 What information is contained in individuals' answers to survival questions?**

In this section, we look at the risk factors and outcomes that are associated with having higher or lower expectations of survival to older ages, as a means to assessing whether individual responses seem to reflect individuals' well-considered judgements about their own future mortality.

It is possible that individuals' answers to survival expectations questions could be uncorrelated with actual outcomes – both relevant risk factors and subsequent mortality – and still represent individuals' 'actual' expectations. However, if individuals' expectations were completely uninformative in this way, we might be concerned about whether they are in fact reflective, well-established views (and hence likely to persist over time and be relevant to future behaviour) or are just numbers picked with little consideration and hence likely to be revised. Conversely, if we see that expectations do correlate with outcomes relevant to mortality, this can be seen as additional evidence that they are well-formed views that are more likely to drive economic decision-making. Furthermore, it is an interesting question in its own right to ask whether individual survival expectations contain information – in terms of predicting future mortality – beyond other facts known about the individual, such as their demographic characteristics and health status.

We first look at whether risk factors that are known to be correlated with mortality (and are also known to the individual) are correlated with individuals' reported survival expectations. There are established causal channels that link certain health behaviours, including smoking and drinking (e.g. Lew and Garfinkel, 1987), and certain health conditions, such as heart attack and cancer, to mortality. For some other outcomes, such as the age at which a parent died, there are plausible channels through which they could be linked to survival probabilities (e.g. the early death of a parent may capture information about a shared environment – such as past diet – which does not manifest itself in the other health conditions examined, but nevertheless has an impact on how long an individual lives).

**Table 2.2. Relationship between stated survival probabilities and risk factors**

Risk factor	Excl. self-reported health		Incl. self-reported health	
	Male	Female	Male	Female
<b>Smoking (relative to non-smoker)</b>				
Ex-occasional smoker	-3.0***	-0.7	-2.8**	-0.2
Ex-regular smoker	-1.9***	-0.4	-1.4**	-0.1
Ex-smoker (frequency: don't know)	-3.0***	-1.9	-2.6**	-1.4
Current smoker	-8.1***	-7.5***	-6.7***	-6.4***
<b>Alcohol consumption (relative to once or twice a month)</b>				
At least 3-4 days a week	0.1	0.1	-0.6	-0.4
Once or twice a week	0.1	0.2	-0.2	-0.1
A few times a year	-1.2	-1.3**	-0.6	-1.0*
Not at all	-2.5**	-2.1***	-1.5	-1.2
<b>Age mother died (relative to 60-64)</b>				
<50	1.6	1.3	1.7	1.6
50-59	1.0	-2.6*	0.7	-2.6*
65-69	3.3**	-1.8	3.1**	-1.5
70-74	2.2	-0.7	2.1	-0.5
75-79	1.9	0.2	1.4	0.3
80-84	3.1**	2.2*	2.7**	2.5**
85+	5.6***	6.9***	5.2***	7.0***
<b>Age father died (relative to 60-64)</b>				
<50	0.9	0.1	1.1	0.1
50-59	-0.1	0.3	0.2	0.5
65-69	0.9	-0.9	0.8	-0.6
70-74	2.7**	1.8*	2.6**	1.8*
75-79	4.5***	2.9***	4.5***	2.9***
80-84	4.4***	2.5**	4.2***	2.4**
85+	7.1***	5.0***	7.0***	4.9***
<b>Health conditions</b>				
Cancer	-6.2***	-3.4***	-4.2***	-2.2***
Heart condition	-3.1***	-2.3***	-1.5*	-1.4*
Hypertension	-2.1***	-2.3***	-0.8*	-1.3***
Stroke	-3.3***	-1.0	-1.4	0.7

---

Note to Table 2.2: Coefficients represent percentage point deviations in mean response. Statistical significance at the 10%/5%/1% level is denoted by \*/\*\*/\*\*\*. Standard errors are clustered at the individual level. Other control variables, for which coefficients are not reported, are whether in a couple, income and wealth quintile, education level, whether working and dummy variables for whether diagnosed with Alzheimer's, angina, arthritis, diabetes, lung disease, osteoporosis, Parkinson's and psychiatric disorders, whether the individual is white or non-white and a full set of dummy variables for each single year-of-age and wave interaction.

Source to Table 2.2: ELSA waves 1–7. 52,170 observations of 14,092 unique individuals.

Table 2.2 shows the results of an analysis looking at the correlation between these risk factors and reported survival probabilities. We conduct an econometric analysis that attempts to answer the question: 'For someone of a particular age interviewed at a particular point in time, what is the difference in average reported survival probability for a person who has risk factor X compared with a person who does not (holding other risk factors constant)?'. The reported numbers, for selected risk factors, represent the mean percentage point difference in response for someone with the risk factor compared with someone without it. We run one analysis in which individual self-reported health is not included as a risk factor and one in which it is included, and produce both analyses separately for men and women to allow for potentially different impacts of these risk factors between sexes.

There are strong correlations between a wide number of risk factors and self-reported survival probabilities. Moreover, these correlations are in the direction that one would expect, given other evidence on the role of these risk factors. For example, male ex-smokers report a 1–3 percentage points lower probability of survival (conditional on their age, wave of interview and other risk factors) than male non-smokers. Being a current smoker has an even larger association with survival reports, with those currently smoking reporting 6–8 percentage points lower probability of survival on average than non-smokers. We see that having been diagnosed with one of the four major health conditions we highlight is associated with a significantly lower reported survival probability for both men and women. While in the case of a stroke, once one controls for the individual's self-reported health, there is no statistically significant correlation between having ever been diagnosed and reported survival probability, in the case of cancer the 'impact' of diagnosis on survival report is diminished, but still strongly statistically significant, when controlling for self-reported health. It is worth noting that not all of the results accord with our understanding of the relevant causal channel. In the case of alcohol consumption, those who do not drink at all have lower survival expectation than those who drink once or twice a month, though this difference is not significant when controlling for self-reported health.

Individuals' view of their own current level of health is arguably a primary channel through which these health conditions and risk factors could affect reported survival probabilities. Carrying out the analysis both without and with controls for self-reported health tells us the extent to which these risk factors are correlated with reported survival probabilities both with and without the part of the correlation that is due to this channel. We find that when controlling for self-reported health, the risk factors are still correlated, albeit to a lesser extent, with self-reported survival probabilities. This suggests individuals' reported survival probabilities reflect not just their current health state but also an awareness of the level of risk that they face in the future due to their health behaviours and current health status. Again, this is consistent with a reflective and considered assessment of future mortality prospects based on relevant information.

The above analysis is reassuring initial evidence that individuals' responses to survival questions reflect factors relevant to mortality. We may be further interested in the question of whether the relationship between risk factors and survival expectations is one where knowledge of exposure to risk factors has a causal impact on individual expectations or whether, for example, expectations are driven by general pessimism about the future that also drives low care for personal health. If the latter were the case, expectations would correlate with survival without there being any feedback from real factors to the formation of beliefs. We might be concerned that, in this case, the expectations questions could be giving us information about individuals' optimism and pessimism, but little else.

We can explore the above question further by assessing whether individuals' expectations change over time in response to news, e.g. a new diagnosis of a health condition. We do this by looking at the average change in responses to survival questions between waves, controlling for the individual's age and sex, and looking at how this change is affected by whether or not an individual has been diagnosed with a new health condition between waves. The results of this analysis are reported in Table 2.3. We find that individuals do respond to new information by revising their survival expectations. For example, a new diagnosis of cancer or a case of a stroke cause large and statistically significant downward revisions in survival expectations of 5 and 6 percentage points respectively. While the estimated revisions for other conditions are in some cases marginally significant, or not significant at all, in almost all cases the average revision in response to negative news about health is negative. This analysis provides strong evidence that answers to expectations questions represent, to some extent at least, an assessment that is founded on information that an individual has about factors affecting their likelihood of surviving to certain ages.

**Table 2.3. Average revision to survival expectations following diagnosis with major health conditions**

Condition newly diagnosed	Percentage point revision in expectation	
	1 <sup>st</sup> survival question	2 <sup>nd</sup> survival question
Alzheimer's disease	-7.8*	-1.9
Cancer	-4.5***	-3.0*
Dementia	7.4**	-2.6
Heart attack	-2.5	-1.3
Lung disease	-2.2	1.4
Parkinson's disease	-2.9	-8.6
Psychiatric problems	-2.0	-1.2
Stroke	-5.7***	-5.0*

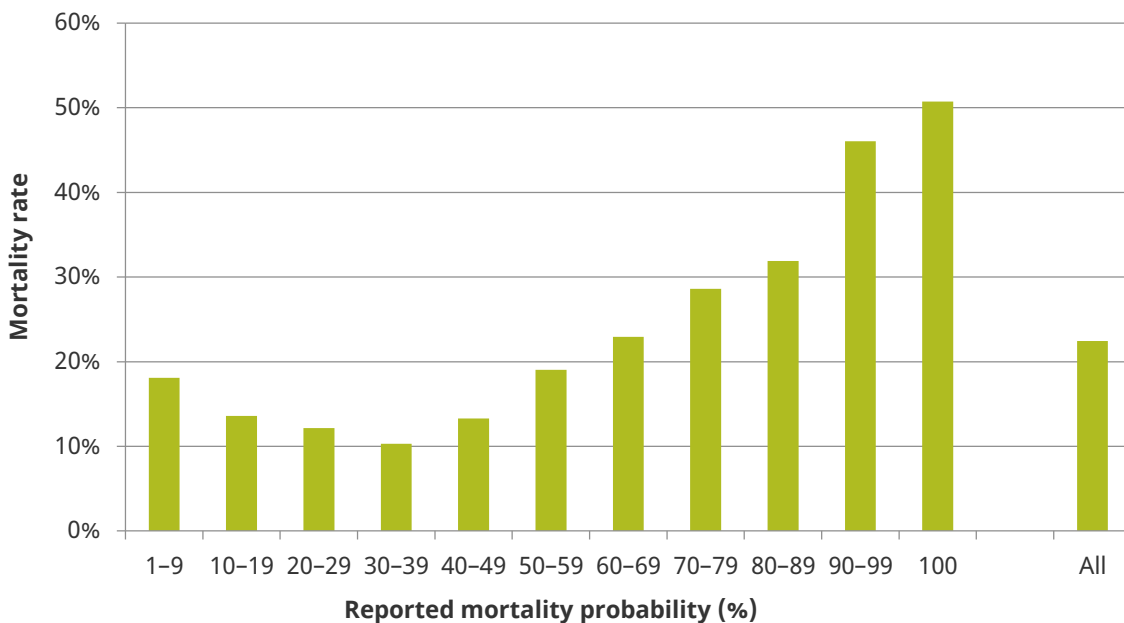
Note: Coefficients represent percentage point deviations in mean response. Statistical significance at the 10%/5%/1% level is denoted by \*/\*\*/\*\*\*. Standard errors are clustered at the individual level.

Source: ELSA waves 3–7. 37,760 observations of 12,027 unique individuals.



Finally, we show evidence that survival expectations are correlated not only with known mortality risk factors but also with actual mortality outcomes. We are able to use information about the date of death of individuals surveyed in ELSA to determine whether or not a person had died within a given time after their answer to a survival question. As we know, the age about which individuals are asked their chances of survival in ELSA is 11–15 years greater than their current age. ELSA is linked to NHS mortality data for most of our sample, which include the date of death of individuals who died before February 2013. As the first wave of ELSA interviews was carried out in 2002–03, this means that the longest time horizon over which we can assess individual mortality following response to a survival question is 10 years. While we cannot therefore compare, for each individual, their reported survival probability and their survival outcome for exactly the same age, we can get reasonably close for respondents to wave 1. Figure 2.3 shows the average 10-year mortality rate for different groups of individuals according to their reported mortality probability. We see a strong relationship between mortality expectations and actual mortality rates.

**Figure 2.3. Ten-year mortality rates by reported probability of death**



Note: Reported probability of death is 100 minus the reported probability of survival in the first survival question the individual is asked. Distribution of responses in all waves is shown in Figure 2.2.

Source: ELSA wave 1 and linked death records. 6,225 individuals.

We test rigorously for the correlation between mortality expectations and outcomes when controlling for individuals’ age-specific and sex-specific mortality risk. Using two alternative econometric models, we find a robust relationship between mortality expectations and outcomes. This relationship remains even when we control for the same set of risk factors explored in the earlier analysis. This means that responses to individual survival expectations questions give us additional information about the probability that an individual will survive to older age even once we have accounted for their age, sex, health behaviours, existing health conditions and self-reported health. This is strong evidence that these responses are not simply answers picked with little thought but in fact represent well-considered judgements about survival prospects reflecting a wide range of information – both objective and subjective – that an individual has about their mortality.

---

In summary, we have seen in this section that reported survival expectations correlate with known mortality risk factors. This holds true (though more weakly) when controlling for self-reported health. These correlations seem to be driven by expectations responding to information, rather than some other factor driving both of these things. Furthermore, these expectations contain 'information' about actual mortality over and above that given by objective risk factors. This is encouraging from the perspective of interpreting these reported expectations as reflective and well-considered judgements and is in alignment with similar evaluations of self-reported survival probabilities in other settings.<sup>3</sup> However, the fact that survival expectations correlate with risk factors, are predictive of mortality and respond to information in this way does not necessarily imply that individuals' expectations are accurate. It is to this issue that we turn in Section 3.

---

<sup>3</sup> See, for example, Hurd and McGarry (1995) for analysis of answers to equivalent questions in the US Health and Retirement Survey and Manski (2004) for a review of subjective expectations in relation to economic decisions more widely.

---

### 3. Patterns in subjective survival expectations

#### Key findings

**Individuals from a range of ages and birth cohorts underestimate their chances of survival to ages 75, 80 and 85, on average.**

Those in their 50s underestimate their chances of survival to age 75 by around 20 percentage points and to 85 by around 10 percentage points.

**Individuals in their late 70s and 80s are, on average, mildly optimistic about surviving to ages 90, 95 and above.**

This optimism becomes larger at older ages (10–15 percentage points when looking at age 95) and is larger for men than for women, amongst those born in the 1920s and 1930s.

**Subjective survival curves, estimated using individuals' stated survival expectations, capture the general patterns in expectations.**

These show growing pessimism relative to official life tables for ages up to the mid 70s before turning to mild optimism from the late 90s onwards.

**Actual survival probabilities differ significantly according to individuals' education, wealth and marital status.**

Women of 60 in the bottom household wealth quintile had a 65% chance of surviving to age 80, compared with 87% for those in the top quintile, based on mortality data.

**Subjective survival curves reflect differences in actual mortality rates between groups to varying degrees.**

Widowed men and women aged 60 show the greatest survival pessimism, underestimating their probability of survival to age 80 or above by almost 30 percentage points on average.

In this section, we examine the question of whether there are systematic biases in individuals' expectations about their own survival. It would be impossible to assess whether any one individual had 'accurate' expectations about their own chances of

---

survival. The chances that an individual lives to a particular older age depends on a whole host of factors, many of which cannot be observed, change over time, or interact in ways we do not fully understand. However, we can assess whether a large group of people are, on average, close to having accurate expectations about their own chances of survival. We compare the average of individuals' reported expectations with official life tables and then look at how the average expectations of particular groups compare with mortality rates in the corresponding ELSA sample.

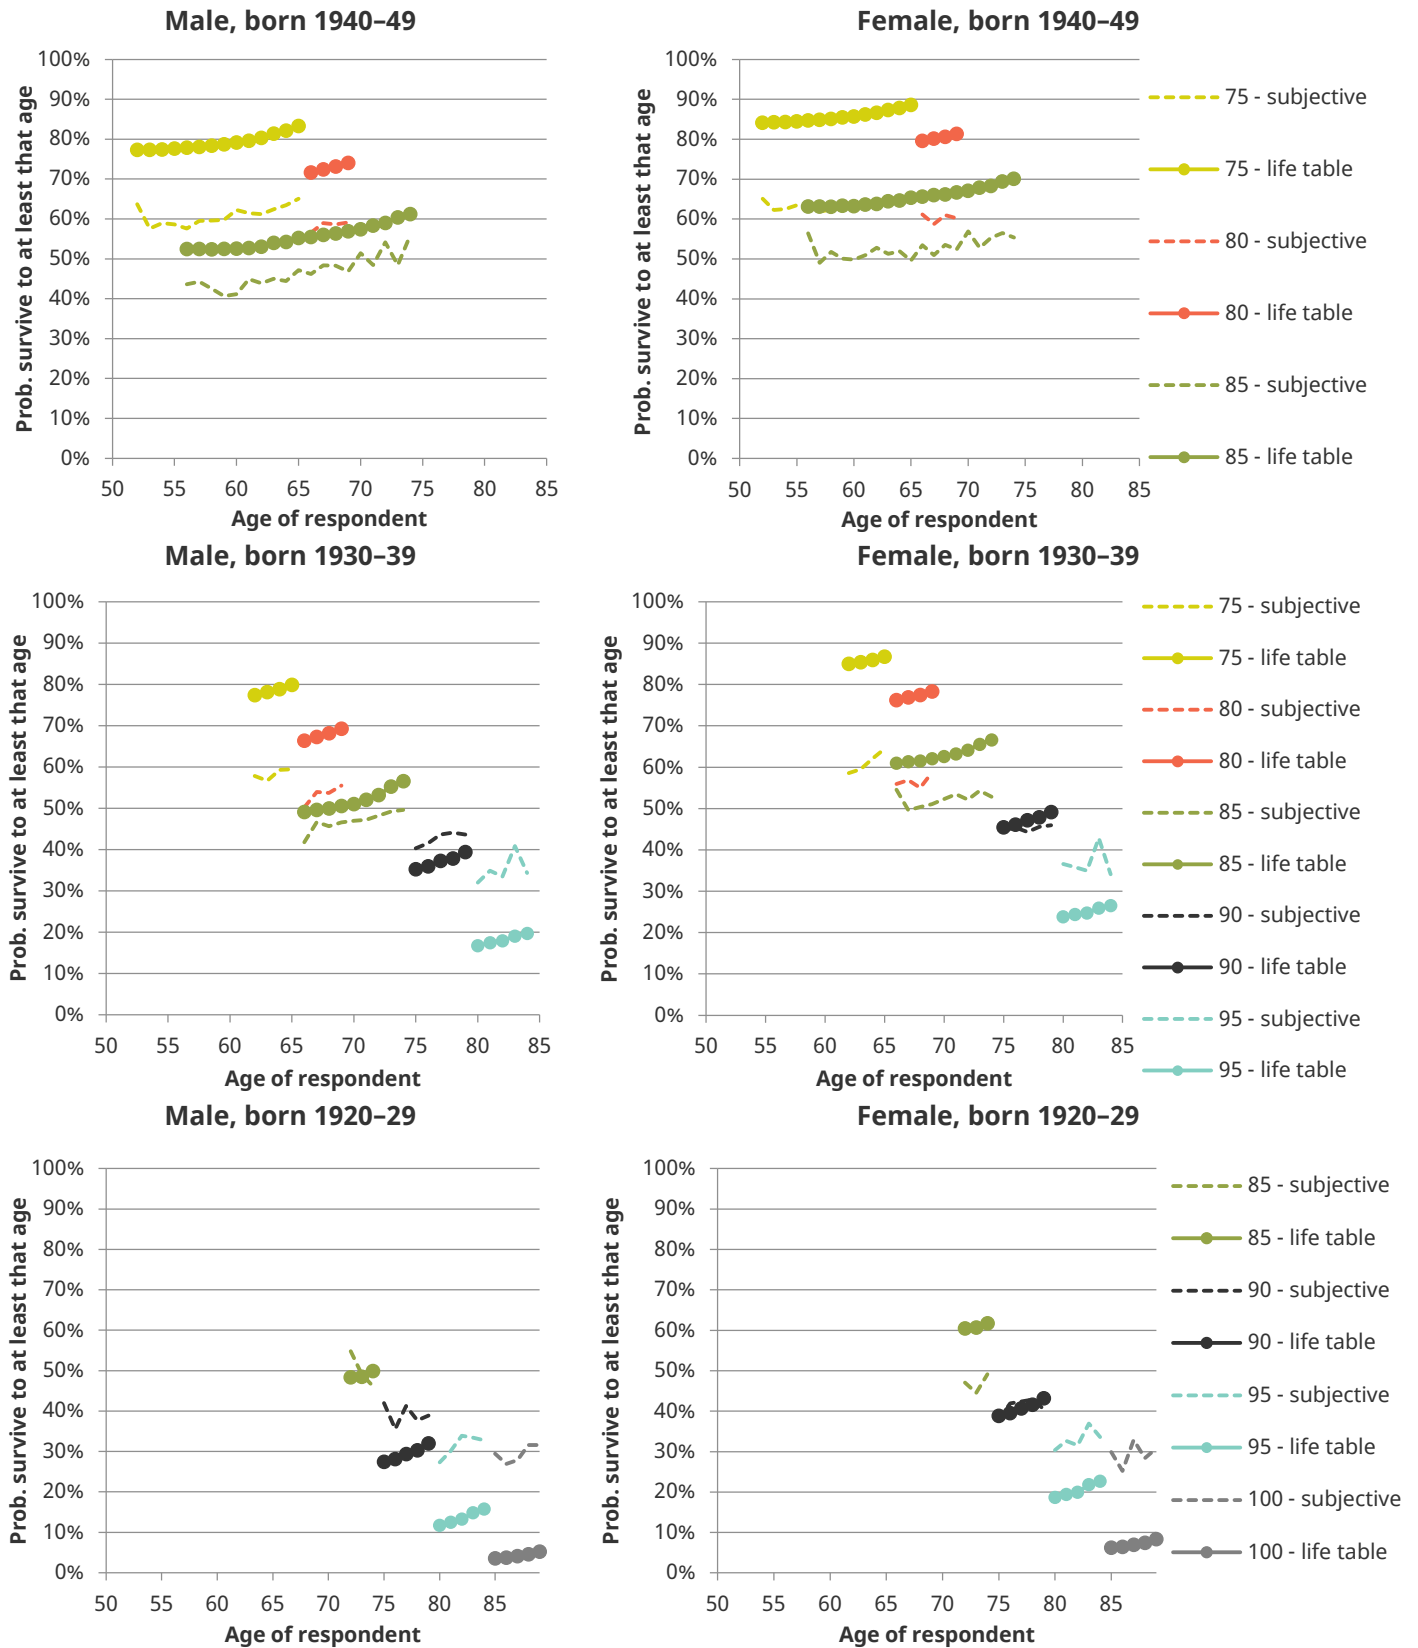
### **3.1 Comparison of 'subjective' expectations and official estimates of survival probabilities**

We assess whether individuals are, on average, accurate about their expectations of survival to older ages by comparing the average reported survival probability for individuals of a particular group with the survival probability given by the relevant official life table. Official life table estimates are an attempt to give an unbiased estimate of the proportion of individuals of a particular group who will survive to a given older age. In other words, they are designed such that the risks of either overestimation or underestimation of survival are balanced and hence they can be seen as a best unbiased estimate of survival probabilities, against which average subjective reports can be judged.

Life tables for England are produced by the Office for National Statistics (ONS). The ONS produces life tables giving historic and projected survival rates that are specific to individuals of each sex and single birth-year cohort. Section A.1 in the appendix gives details of the information contained in life tables and the definitions of statistics – such as life expectancy – derived from these. With these data, it is possible to answer questions such as: 'What is the official estimate of the probability that an English man, aged 60 and born in 1957, will survive to the age of 90 or above?'. We are therefore able to make comparisons between average subjective expectations and life tables at the level of individual sex and birth cohort. In Section 3.3, we attempt to make comparisons between average subjective expectations and actual survival rates, within more specific subgroups, by using the ELSA data on mortality.

Previous research comparing subjective expectations with life table survival probabilities has focused on US data. Hurd and McGarry (1995) use data from the US Health and Retirement Survey (HRS) to document that the average of individuals' expectations is not far away from period life table estimates. However, Elder (2013) uses more recent data and compares individuals' reports with life tables and with actual mortality rates in the HRS. He finds that individuals are, on average, pessimistic about survival up to around age 85 and gradually more optimistic thereafter. Further evidence of survival pessimism at younger ages, combined with increasing survival optimism at older ages, is found by Wu, Stevens and Thorp (2015). Their study uses Australian data about retirement plans to disentangle and model age and cohort effects. However, Bissonnette, Hurd and Michaud (2017), attempting to correct for mortality differences between sample respondents and the population average, find evidence of mild survival optimism at ages above 70, with variation across groups.

**Figure 3.1. Comparison of average survival rate in subjective responses and life tables**



Note: Different coloured series correspond to different ages about which respondents are asked questions.

Source: ELSA waves 1-7 and ONS 2014-based cohort life tables for England and Wales.

---

As outlined in Section 2.2, ELSA asks individuals about survival to particular ages that are multiples of five. Therefore, for each single birth-year cohort, we have responses to questions about two or three different older ages. For these particular ages, we are able to make direct comparisons between the average subjective survival report and the corresponding life table survival probability. In Figure 3.1, we group single-year cohorts together to form 10-year cohorts, in order to make comparisons easier. The fact that we have pooled together single-year cohorts in this way does not change the qualitative conclusions of our analysis.

We find that individuals are on average ‘pessimistic’ about reaching ages towards the lower end of those asked about and are ‘optimistic’ about reaching older ages. More precisely, we see that for men in all birth cohorts, average reported probabilities of surviving to ages 75 and 80 are significantly lower than the corresponding life table figures. Men in more recent cohorts are also ‘pessimistic’ about living to 85, while those in earlier birth cohorts are, on average, accurate about their expectations of living to this age. Looking at ages 90, 95 and 100, we see that men become increasingly ‘optimistic’ about their survival, relative to life tables, when looking at progressively older ages. For women, the qualitative pattern of pessimism about survival to younger ages turning to optimism about survival to older ages also holds true in each cohort, although, relative to men, women are more pessimistic (or less optimistic) at each age.

The ELSA data are particularly well set up to help us identify relative ‘optimism’ and ‘pessimism’ that can change depending on the age being examined. By asking about survival to particular ages – and, in some cases, multiple different ages for the same individual – as opposed to life expectancy, we gain more precise information about the nature of divergences between life tables and expectations. For example, our analysis suggests that it is likely that a 65-year-old man believes he has a lower chance of surviving to age 75 than is given by the relevant life table, but believes he has a chance of living to 85 that is much closer to that given by the life table. It is plausible that for individuals of some ages, when considering their life expectancy, these two effects would ‘cancel out’ such that they give, on average, an accurate answer to the question about life expectancy. This highlights the fact that life expectancy is a summary measure of a large amount of information about survival, and so an analysis based on such a measure might fail to reveal, or analyse less fully, trends in survival expectations with important economic implications.

### 3.2 Constructing subjective survival curves


In Section 3.1, we compared individuals’ expectations about survival to particular ages with the relevant life table survival probability. Ideally, we would like to compare individuals’ expectations about survival to each possible older age with the relevant life table probability. That is, we would like to compare a subjective-based and an objective-based survival curve.

What can we infer about an individual’s subjective survival curve based on the responses that we have in ELSA? For individuals for whom we have a response to only one question about survival expectations, there are an infinite number of possible survival curves consistent with this response and there is little we can say about their subjective survival curve. However, with responses to two questions, we can produce, with plausible further assumptions, a survival curve consistent with these answers. Section A.3 in the appendix

gives a full description of the method that we use to construct subjective survival curves for those individuals who give two responses to survival questions (i.e. all individuals under the age of 70). In essence, we need to make an assumption about the type of shape taken by individual survival curves. We can then select the curve that is of this general shape and passes closest to the responses given by the individual. We make the particular assumption, widely used in the epidemiological literature and for modelling ageing processes generally, that subjective survival curves follow a Weibull distribution (see the appendix for further details). In order to produce survival curves that imply plausible beliefs about survival to the oldest ages, we also make the further assumption that individuals' subjective expectations about survival to 110 match the life table value and treat this as a third 'response' when fitting the survival curve.

Figure 3.2 shows an example of a survival curve based on individuals' subjective reports about their survival expectations and compare this with the relevant life table survival curve. We show the self-reported-based survival curve produced both with and without the assumption that individuals know their probability of survival to 110. We can see that this additional assumption gives far more plausible beliefs about survival to the oldest ages. In the example shown, this assumption means that the average implied belief about the probability of survival to age 100 or older is 14%, rather than 28%. These particular curves are based on the median response amongst women born 1940–49 and are compared with the life-table-based survival curve for women born in 1945.

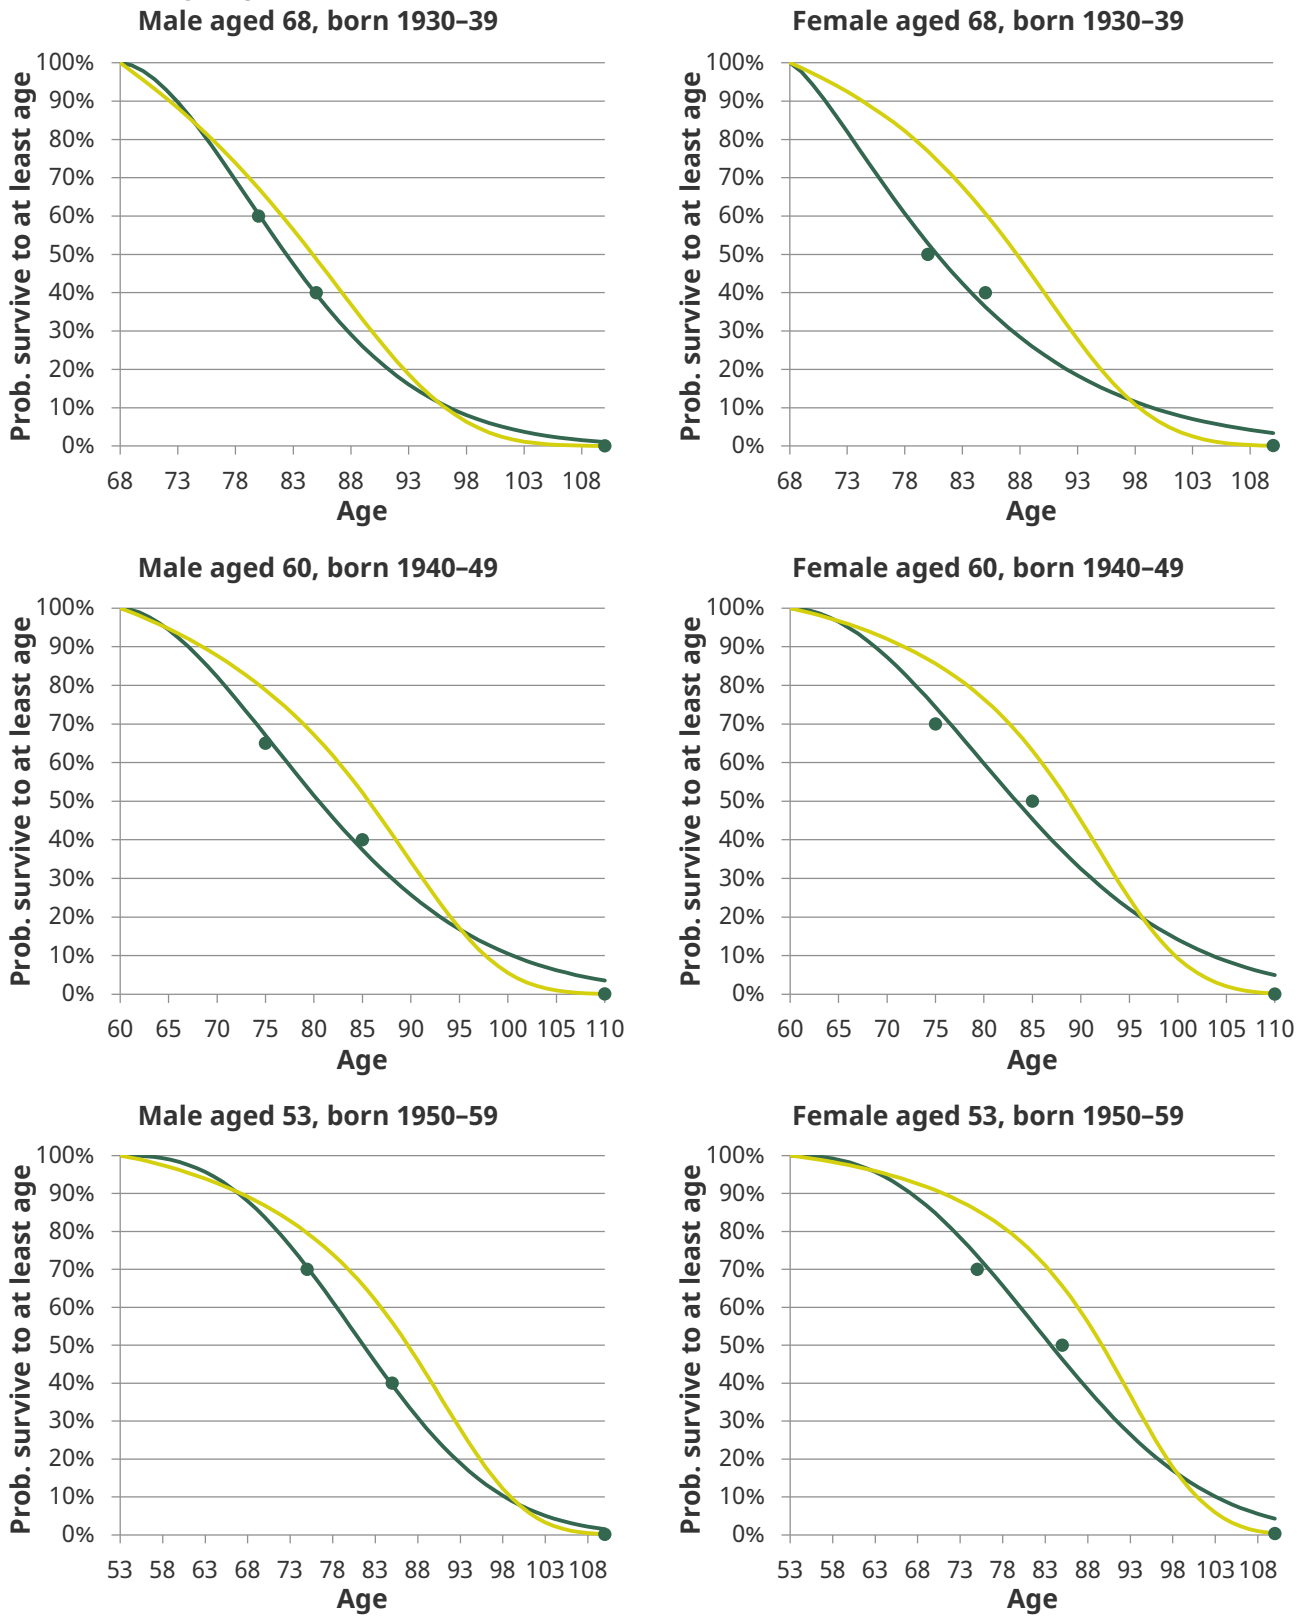
**Figure 3.2. Comparison of self-reported- and life-table-based survival curves for women born 1940–49**



Note: Life table curve is the survival curve based on actual and forecast mortality data for 60-year-old women born in 1945.

Source: ELSA wave 3 and ONS 2014-based cohort life tables for England and Wales.

**Figure 3.3. Comparison of subjective and objective survival curves, by sex and birth cohort groups**



Note: Each curve shows, given the initial age in each (i.e. 68, 60 and 53), the probability of survival to older ages.

Source: ELSA wave 3 and ONS 2014-based cohort life tables for England and Wales.



---

We can see that these survival curves encapsulate the main qualitative features that we saw when comparing the subjective responses and life table survival probabilities – namely, significant survival ‘pessimism’ when looking at survival to ages 75, 80 and 85, which turns to modest ‘optimism’ when looking at survival to ages from the mid 90s onwards.

Figure 3.3 compares objective and subjective survival curves for three different cohorts for both men and women, based on the median reported subjective survival probabilities in each group. We see that the pattern of pessimism turning to mild optimism is generally found across cohorts and sexes.

### 3.3 Comparison of subjective and objective survival curves for subgroups

In this section, we compare subjective and objective survival curves for different types of individuals, focusing on differences between groups according to wealth, marital status and level of education. There are differences in individuals’ expectations of how long they will live along each of these dimensions and these reflect actual differences in mortality experiences between groups. Comparing expectations with reality within each subgroup shows that each, on average, follows the same qualitative patterns of relative pessimism and optimism exhibited in the overall population.

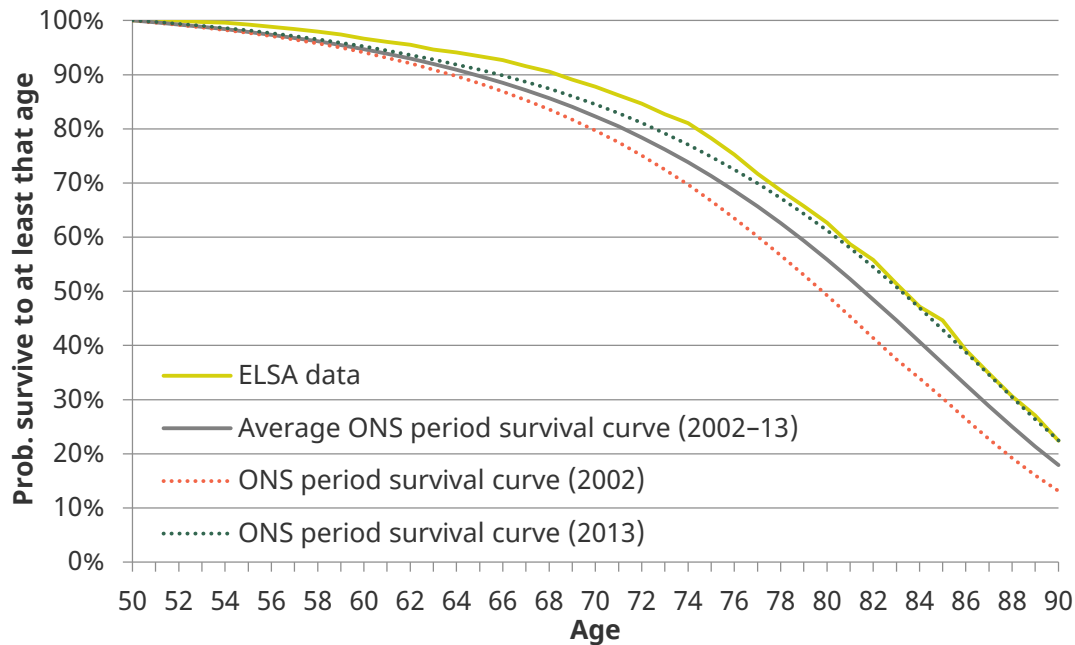
The ONS does not produce life tables for the subgroups that we examine in this section. In order to generate ‘objective’ survival curves, we therefore use the data on the actual mortality of ELSA respondents. To calculate a survival curve for a particular subgroup – e.g. married women – we take, for each age, all of the observations in our data of married women and calculate how many died at that age to obtain a mortality rate. ‘Stitching together’ these mortality rates gives us a survival curve.

In an ideal world, we would calculate objective survival curves for each cohort of each subgroup being examined, given that increasing longevity means survival curves have generally changed over time. With the data available, we can only track each cohort for around 10 years and so this is not feasible. What we therefore calculate is a subgroup survival curve that combines the mortality rates for all of the cohorts that are interviewed in ELSA. This survival curve can be thought of as somewhat like a period survival curve: it reflects the mortality experiences at different ages at a particular point in time. In this case, that point in time is the period 2002–13 rather than a single year.

Figure 3.4 compares the objective survival curve created using the ELSA male mortality data with the average of the ONS period life curves for 2002 to 2013. There is a slight difference in the survival curves, due to marginally lower mortality rates amongst the ELSA sample of men than in the whole population, for most of the ages up to 74. These differences are not very large, reaching at most 8 percentage points difference in survival probability. We should expect a slightly lower mortality rate amongst individuals in ELSA than in the general population because ELSA is representative only of the household population and not of individuals in institutions, including residential care and hospitals,

for whom we would expect higher subsequent mortality rates.<sup>4</sup> For the sake of comparison, we show the ONS period survival curves for 2002 and 2013, to demonstrate the magnitude of changes in mortality rates over time and so give some sense of the significance of the difference between the ELSA and ONS survival curves. This exercise can therefore reassure us that the ELSA mortality rates are a reasonably close approximation of those in the population at large.

**Figure 3.4. Comparison of period survival curves in ELSA data and ONS life tables, constructed using mortality rates for men aged 50–90 in 2002–13**




Note: The ELSA-calculated survival curve is based on 61,873 observations of 16,800 unique individuals. The average ONS period survival curve is an equally weighted average of the 2002–13 curves.

Source: ELSA waves 1–6 and ONS 2014-based period life tables for England and Wales.

We now turn to compare survival curves within certain subgroups. Figure 3.5 shows the subjective and objective survival curves for men and women with different levels of education. Individuals are classified as having a ‘low’, ‘mid’ or ‘high’ level of education. The low category includes all those who left education up to and including the school-leaving age for someone of their birth cohort, the mid category includes those who left above the compulsory age and below the age of 19, and the high educational category includes all those who remained in full-time education to the age of 19 or above. The ‘objective’ curves are based on the actual mortality rates of individuals of the relevant sex and education level in the ELSA data. The ‘subjective’ curves are based on the average answers to the survival expectations questions amongst those of the relevant sex and education level at age 60.

<sup>4</sup> To be clear, ELSA does record whether individuals in the sample who then move into residential care or into hospital subsequently die. It is at the point of initial selection of the sample that those in institutions are not captured.

**Figure 3.5. Comparison of objective and subjective survival curves, by education level**



Source: ELSA waves 1–6.

The first thing to note from Figure 3.5 is that, for both men and women, the objective survival curves for those with higher levels of education are higher at all ages, indicating a higher chance of survival to older ages, on average. For instance, the figure says that low-educated men at age 60 face a 61% chance of surviving to at least age 80, whereas the chance for high-educated men is 77%. The subjective curves show differences in expectations by education group. For women, there is a clear ordering whereby those who stayed longer in full-time education expect a higher chance of surviving to older ages than do those who left education at a younger age. For men, the average responses of the mid- and high-educated groups correspond to very similar survival curves, although both show greater expected longevity than the low-educated group. For example, the subjective curve for low-educated men implies a 47% chance of survival to 80 or above, compared with 58% for the high-educated men.

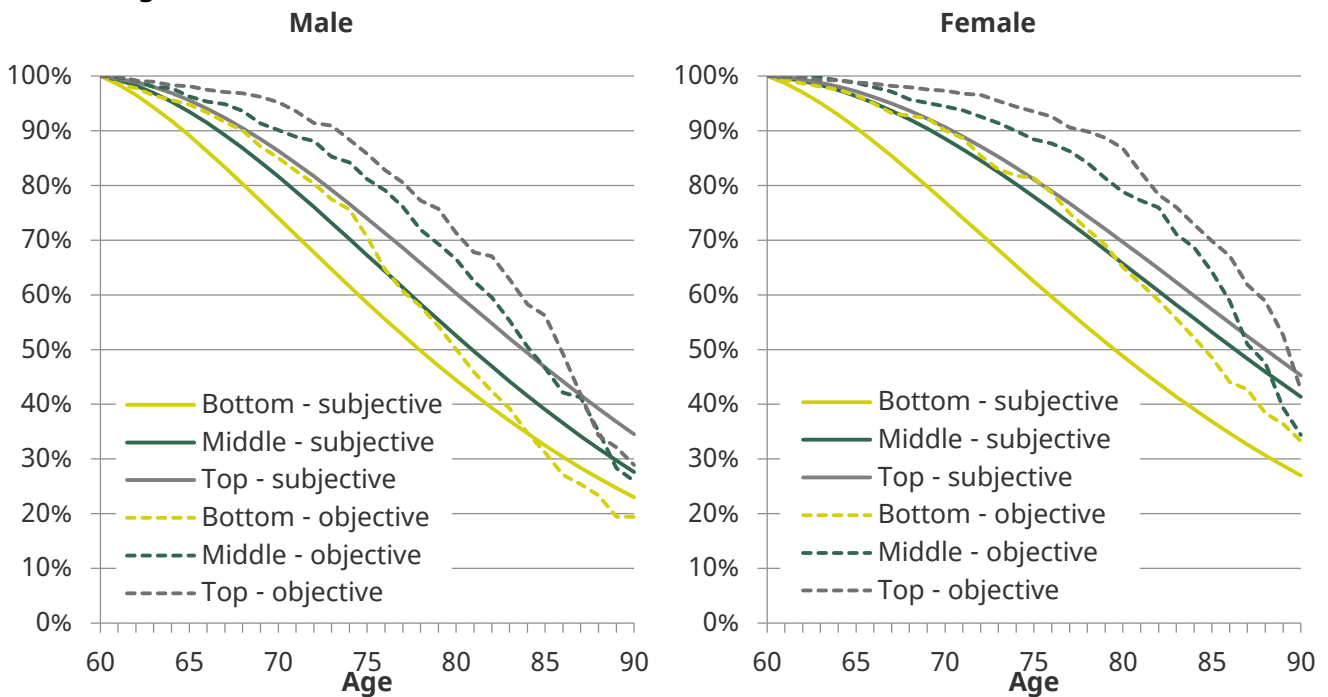
The second point to note from Figure 3.5 is that each subgroup for both sexes shows the qualitative pattern that the average expectation of survival is lower than the objective probability below around age 90. This relative ‘pessimism’ grows larger until around the mid 70s and then becomes smaller until objective and subjective curves cross at around age 90. Amongst women, there is a clear pattern whereby those with lower levels of education are more ‘pessimistic’ on average than those with higher levels of education. The gap between objective and subjective curves reaches 20 percentage points among low-educated women but does not exceed 11 percentage points amongst the high-educated. Amongst men, the differences are less stark, and the mid-educated group is least pessimistic at most ages.

In each case, the objective and subjective curves are not a perfect ‘like-for-like’ comparison. As already described, the ELSA mortality data approximate a period survival curve for the years 2002 to 2013, but the subjective curves are based on average expectations amongst those born in 1945–49 when interviewed in wave 3 (2006–07), for

whom their actual 'objective' survival curve would be the cohort survival curve for that group.<sup>5</sup> This means that the objective curve will not take into account any changes in mortality rates that ought to be expected for each subgroup. However, any such changes in mortality rates across time are very unlikely to be large enough to alter any of the qualitative conclusions of this analysis.

Figure 3.6 compares objective and subjective survival curves for individuals with different levels of wealth. Each individual is classified into one of five quintiles, according to their position in the household wealth distribution at age 60 (or, for individuals not interviewed when aged 60, their position in the wealth distribution when first interviewed aged over 60). To aid comparison, the figure shows curves for the bottom, middle and top wealth quintiles only. Those with higher levels of wealth experienced higher survival rates amongst both sexes. For a 60-year-old woman in the lowest wealth quintile who experienced average mortality rates for women in this group over the period of our data, the chance of surviving to the age of 80 or older is 65%. For a woman in the top wealth quintile, it is 87%.

**Figure 3.6. Comparison of objective and subjective survival curves, by wealth quintile at age 60**



Note: Wealth quintiles are based on position in the household wealth distribution at age 60 (or, for individuals not interviewed when aged 60, their position in the wealth distribution when first interviewed aged over 60) and calculated separately for each wave.

Source: ELSA waves 1–6.

Subjective expectations also show differences between groups, with chances of survival to age 80 being 49% and 70% in the bottom and top wealth quintiles respectively. Again, the pattern of survival pessimism is found in each subgroup for both sexes up until their mid-to-late 80s. Only the women in the bottom quintile still show (mild) survival pessimism at

<sup>5</sup> This group has an average age of 60 when interviewed.

---

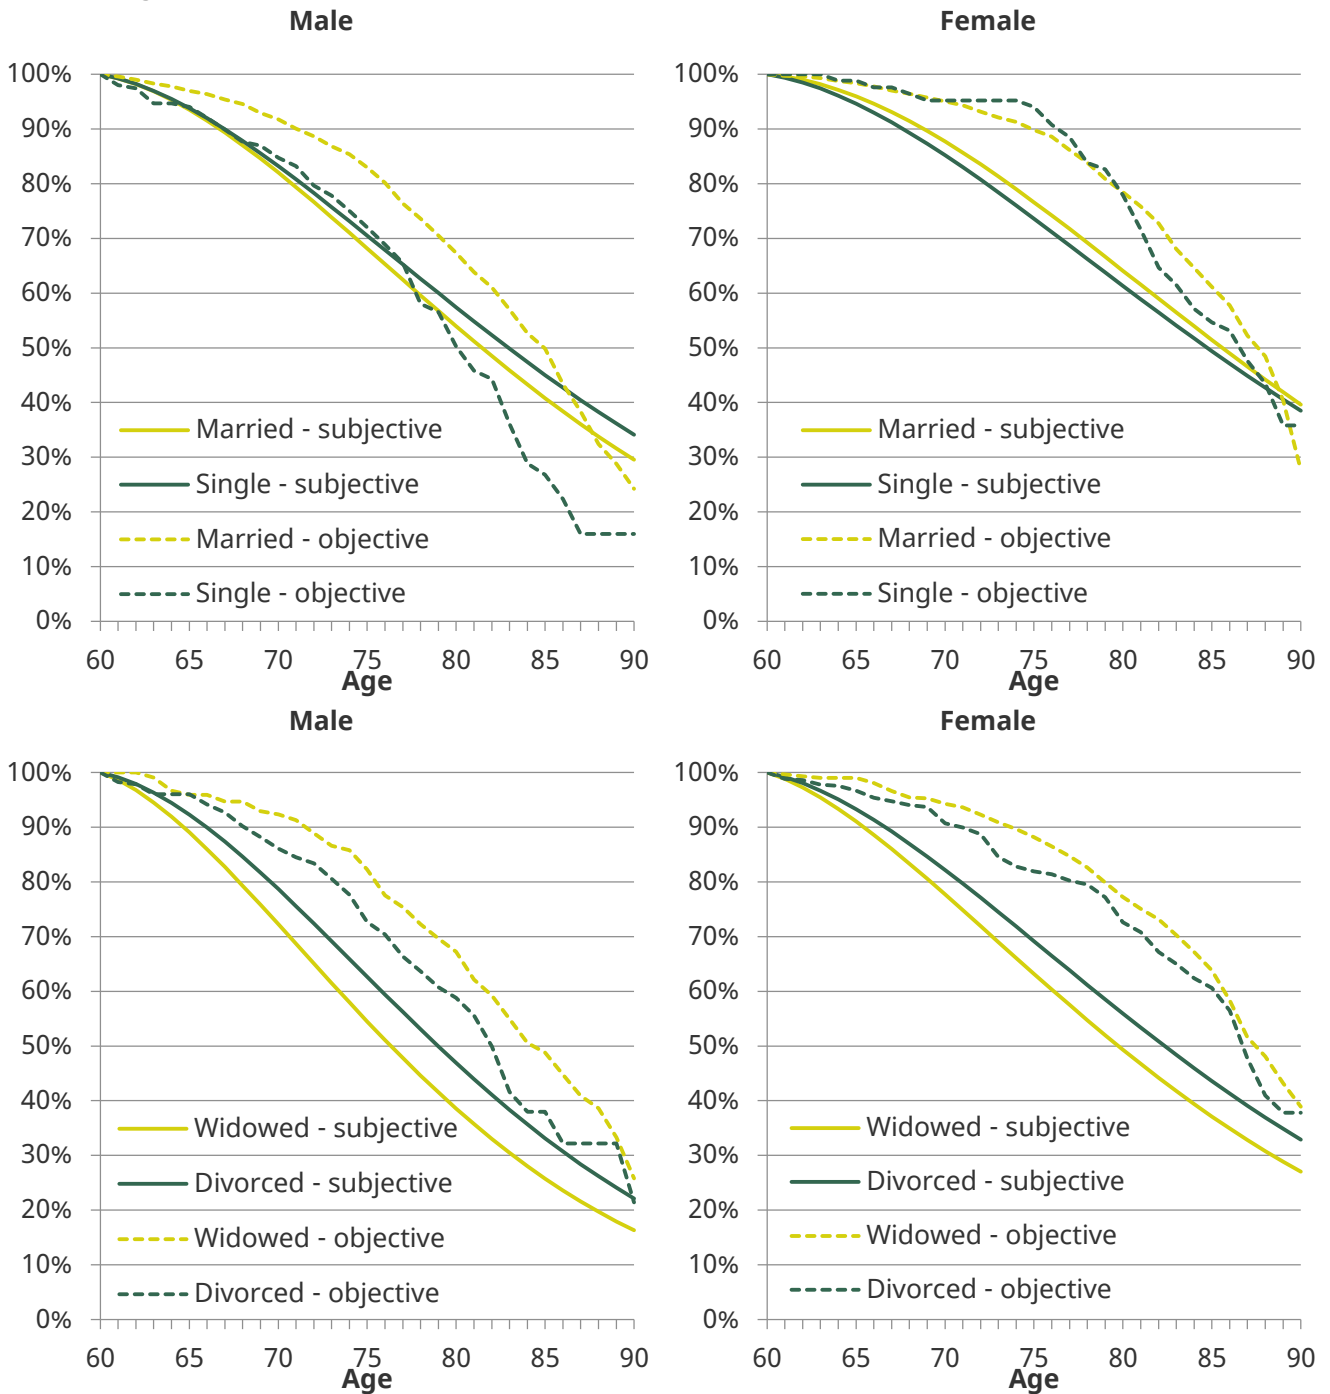
age 90. The degree of pessimism is similar amongst men with different levels of wealth when considering survival up to around age 75, but relatively high mortality rates for the lowest-wealth men at older ages means that this group is much less pessimistic than high-wealth individuals about survival through their 80s. Women in the lowest wealth quintile are most pessimistic about survival through their 60s and 70s, but their 'pessimism' converges somewhat with that of high-wealth women at older ages.

It is important to stress that the differences in objective and subjective survival curves between groups do not imply a causal link between, for instance, having lower wealth and having a higher mortality rate. While there could be mechanisms whereby the characteristics explored do cause mortality differences – e.g. wealthier individuals might be able to purchase products that improve their health and higher-educated individuals might find it easier to understand information about the risks associated with certain health behaviours – it is also possible that mortality and these other characteristics are both affected by some other factor. For example, those with a long-term illness might die younger and also tend to be in a lower wealth quintile because they have been less able to work and faced higher costs of living.

Finally, we explore patterns by marital status. Figure 3.7 shows the subjective and objective survival curves for four subgroups – married or cohabiting, single (never married), divorced (including separated) and widowed. The differences in subjective expectations of survival between groups are not as large as was seen in the case of education or wealth. Nevertheless, expectations of survival to older ages are lower for divorced individuals than for those who are single or married, and the expectations of widowed individuals are lower still. For example, married and single women's subjective expectations imply a 40% and 39% chance respectively of survival to age 90 or older (conditional on surviving to age 60), whereas the probabilities for divorced and widowed women are 33% and 27% respectively.

Turning to the objective survival curves, there are not large differences in survival by marital status for women. However, we see significantly lower survival rates for single men than for the other male groups, especially from around the age of 80 onwards. Interestingly, the subjective expectations of single men do not reflect this lower objective survival rate. In fact, single men show significant survival optimism from their late 70s onwards. For the other groups, and for single women, we see the established pattern of survival pessimism throughout their 70s and until their late 80s. In the case of divorced and widowed individuals, this pessimism is somewhat more pronounced for both sexes. And even looking at the age of 90, there is significant survival pessimism amongst widowed men and women. For example, widowed men aged 60 believed they had a 16% chance of surviving to age 90, whereas the objective probability is 26%. For widowed women, the corresponding figures are 27% and 39%. The pessimism shown by widowed men and women is significantly greater than amongst any of the other groups by marital status. Both underestimate their chances of survival to their early 80s by just under 30 percentage points, on average. Amongst the other groups, the objective probability of survival does not exceed the subjective curve by more than 20 percentage points at any age.

**Figure 3.7. Comparison of objective and subjective survival curves, by marital status at age 60**



Source: ELSA waves 1–6.

---

## 4. Subjective expectations and economic behaviour

### Key findings

#### **Survival pessimism is a potential driver of the unpopularity of annuities.**

Around 65% of individuals could perceive an annuity priced according to average survival chances as offering an unfairly low income.

#### **Deferral of the state pension, a choice analogous to annuity purchase, is rarely taken up despite being offered at a favourable rate.**

While individuals are roughly twice as likely to defer the state pension if it represents a 'fair deal' given their survival expectations, the overall level of deferral is sufficiently low that we cannot make strong statements about this relationship.

Individuals' expectations about how long they are going to live matter for a whole range of important economic choices. These choices include when and how much to save for retirement, when to retire, how to use wealth at older ages, whether to buy an annuity and whether to buy life insurance. Given the divergences seen between subjective and objective expectations of survival, individual expectations may have important implications for decision-making in these areas. For example, the lower an individual believes their chances of living to older ages to be, the less worthwhile they may think it is to save for their retirement and the faster they may choose to spend their savings once retired. If their true chances of survival to older ages are in fact higher than they believe, then they may face a higher-than-optimal chance of being left with insufficient financial resources to maintain their desired standard of living.

Research focusing on the US has previously explored the implications of subjective survival expectations for a range of economic behaviours. Hurd, McFadden and Gan (1998) find that higher reported survival probabilities are associated with higher savings rates. Hurd, Smith and Zissimopoulos (2004) find that those with particularly low expectations of survival are more likely to retire earlier and to claim social security benefits earlier. Bloom et al. (2006) find that a higher subjective probability of survival is associated with higher wealth levels but has no impact on the length of working life. Using data from a range of European countries, Post and Hanewald (2013) find that greater uncertainty over longevity does not lead to greater saving. Examining the voluntary market for annuities in England using data from ELSA, Inkmann, Lopes and Michaelides (2011) find that those with higher expectations of survival are more likely to purchase an annuity.

---

In this section, we explore in detail the consequences of individual survival expectations for two economic decisions in particular: whether or not to purchase an annuity and whether or not to defer receipt of the state pension.

## 4.1 Annuity choice

An annuity is a product that individuals can purchase and which guarantees them a fixed income until death. Historically, in a UK context, an individual can choose for their annuity to pay a constant cash income each year or to have their income increase year-on-year.<sup>6</sup> Individuals can also purchase a ‘joint life’ annuity which continues to pay out, at some fraction of the original rate, to a surviving partner (or some other individual) until their death.<sup>7</sup>

Individuals or couples who retire with a certain level of savings face a choice of how to spend these over time. If a couple simply tries to spend some amount of their wealth each month or each year (perhaps using an income drawdown arrangement), then they must decide how to weigh up the risks of spending down their wealth quickly – and potentially running out of savings if they live for a long time – or spending more slowly – and having a lower standard of living and perhaps having resources left over when they die. Purchasing an annuity allows an individual to guard against any risk of outliving their financial resources without having to be very cautious in their rate of spending. Another way of expressing this is to say that an annuity is insurance against the risk of living for a long time and so requiring a large amount of financial resources to maintain living standards.

An annuity therefore has the potential to be a useful and welfare-enhancing product for those who arrive at retirement with a stock of savings and wish to maintain a steady standard of living for the rest of their life. If an annuity is priced ‘fairly’ – in the sense that the annuity provider charges a price equal to the total amount they expect to pay out as annuity income – then economic theory predicts that risk-averse individuals will be made better off using their savings to purchase an annuity.<sup>8</sup>

### The annuities ‘puzzle’

In reality, annuities have proved much less popular with consumers than this economic theory would suggest. In the US, only around 7% of individuals aged over 70 have any annuity income, and the ending of compulsory annuitisation that had existed for many individuals in the UK before 2015 was widely seen as popular with consumers.<sup>9</sup> This unpopularity and lack of take-up of annuities has received attention from policymakers and academics. Understanding the reasons for this phenomenon is important when considering how to structure policies around annuities and retirement generally. Uncovering the drivers of decision-making in this area is particularly pertinent as individuals make decisions in the era of pension freedoms (discussed further in Section 5).

---

<sup>6</sup> Products that increase annuity income in line with inflation or some other fixed percentage are not commonly purchased. Association of British Insurers data suggest that 10% or less of annuities increase in line with inflation or some fixed percentage.

<sup>7</sup> Some annuity products also include a guarantee period during which the annuity continues to be paid, even if the holder dies. For further details on recent developments in annuities products in the UK, see Association of British Insurers (2017).

<sup>8</sup> Yaari, 1965.

<sup>9</sup> Finkelstein and Poterba, 2004; Financial Conduct Authority, 2017.



---

Previous work has identified some potential explanations for this so-called ‘annuities puzzle’. These explanations include that individuals may want to keep some wealth in a liquid form in case they face unexpectedly high costs in some particular years, that annuities may not be ‘fairly’ priced due to costs of administration, regulation or a poorly functioning annuities market, or that individuals want to keep some wealth in a form that they can pass on to their children.<sup>10</sup> In the specific context of the UK in the past 10 years, low bond yields have also been perceived as drivers of unattractive annuity rates.<sup>11</sup>

Could the pattern of survival pessimism documented in this paper also be an explanation for the unpopularity of annuities? To answer this question, we ask whether an annuity priced on the basis of the ‘objective’ (i.e. life table) estimate of survival probabilities for an individual will look like a worse deal than one priced on the basis of their subjective beliefs about their probability of survival.<sup>12</sup> If, given what they believe about their chances of survival, many individuals think that they are being offered an annuity at a less than ‘fair’ rate, then this may explain why, for some, an annuity seems to be a bad option.

### ‘Fair’ annuity rates: perception and reality

To explain our approach to answering this question in slightly more detail, it is necessary to briefly set out the concept of an annuity rate and an ‘actuarially fair’ annuity rate. An annuity rate is the rate at which the price of an annuity corresponds to a stream of income. For example, an annuity purchased with a pension pot of £100,000 that guaranteed an annual income of £5,000 has an annuity rate of 5%. An *actuarially fair* annuity rate is the rate that implies that the expected total amount of annuity income received until death is equal to the annuity price. In other words, the actuarially fair annuity rate ensures that, averaging across all possible future lengths of life and weighting these by how likely they are, the total income received from the annuity is equal to the lump sum used to purchase it.

What the actuarially fair annuity rate is depends on the probabilities of an individual surviving to each possible future age. The annuity rate offered to an individual is tailored only to a limited extent to their individual characteristics. For each individual, we calculate an ‘objective’ annuity rate, which is the actuarially fair annuity rate for an individual of their age, sex and cohort. This is a rough approximation of the annuity rate that they might be offered in a well-functioning annuities market.<sup>13</sup> We calculate a second, ‘subjective’ annuity rate, which is based on the subjective survival curve fitted to the individual’s own responses to the survival expectations questions, constructed as described in Section 3.2. This subjective annuity rate is that which the individual would *perceive* as an actuarially fair rate. In both cases, we assume that individuals purchase an

---

<sup>10</sup> Finkelstein and Poterba, 2004; Brown, 2007; Pashchenko, 2013.

<sup>11</sup> Clearly, to the extent that low yields affect the returns to alternative ways of holding wealth in retirement, they do not imply that an annuity *actually* represents worse value for money than these alternatives. This may nevertheless be the perception. See Lowe (2014) for further discussion.

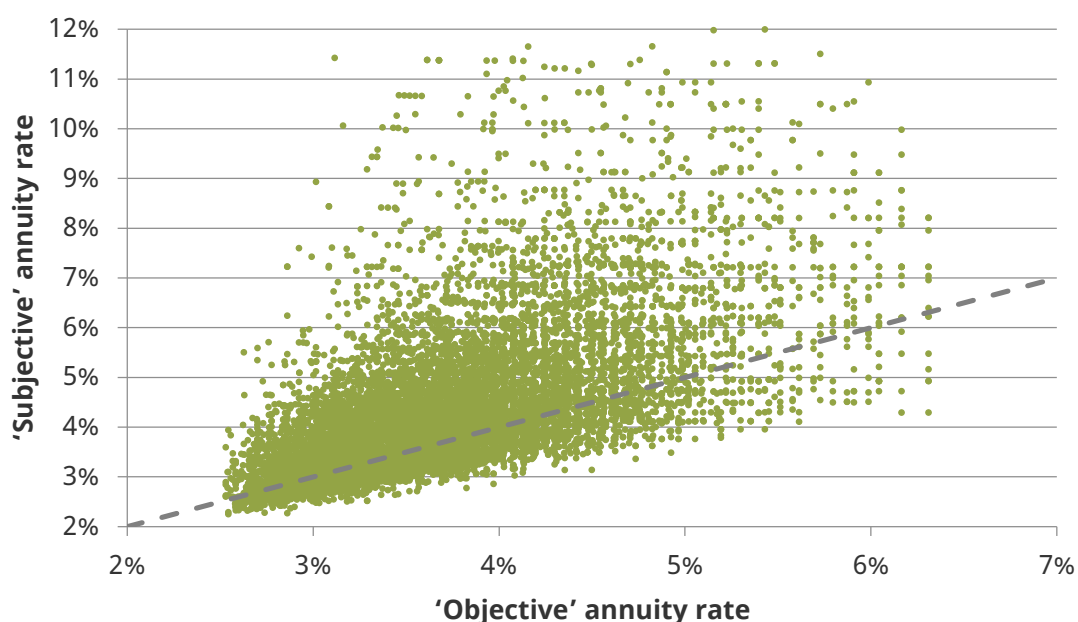
<sup>12</sup> Research has shown that those who purchase annuities tend to live longer than average. If annuities are priced on the latter basis, then those offered will be an even worse deal than those priced on the basis of life tables.

<sup>13</sup> Lowe (2014) uses the average of the three best fixed annuity rates available in February 2014 for men and women aged 55, 60 and 65 and estimates that these are between 90% and 102% of the actuarially fair rate. These calculations assume survival according to ONS life tables and rates of return in line with that on government bonds. In this period, at least, consumers were able to purchase an annuity that was approximately actuarially fair, if they shopped around. The ‘worst’ annuity rates available were estimated to be between 80% and 92% of the actuarially fair rate.

annuity that pays out an income that is constant in real terms. For those in couples, we assume that they purchase an annuity that continues to pay out at a 50% rate to a surviving partner.<sup>14</sup>

We have seen that individuals tend to be pessimistic about survival to ages up to 85. Pessimism will tend to increase what they would perceive to be a fair annuity rate. But we have also seen that individuals are optimistic on average about survival to older ages, and this would tend to decrease what they perceive to be a fair annuity rate. By explicitly calculating an annuity rate based on individuals' subjective survival curves and comparing this with the objective rate, we are able to see whether the pessimism or optimism dominates, at least in its effect on what is perceived to be a fair annuity rate.<sup>15</sup>

**Figure 4.1. Comparison of annuity rates based on 'subjective' and 'objective' survival curves, assuming 0% real interest rate**



Note: 'Subjective' annuity rates are the actuarially fair rate implied by the subjective survival curve constructed from the individual's responses to the survival expectations questions. 'Objective' annuity rates are the actuarially fair rate implied by the survival curve derived from the cohort life table for the individual's sex, age and year of birth.

Source: ELSA waves 3–7 and ONS 2014-based cohort life tables for England and Wales.

Figure 4.1 compares, for each observation of an individual for whom we can calculate a subjective survival curve (all those aged under 70 in ELSA waves 3–7), their subjective and objective annuity rates. Where an individual has a higher subjective rate than objective rate (those observations above the grey dashed line), this implies that they would perceive

<sup>14</sup> The actuarially fair annuity rate therefore depends on the survival curve of both members of the couple. We construct and use the objective and subjective survival curves for both partners in a way analogous to that for an individual.

<sup>15</sup> Calculating an actuarially fair annuity rate also requires an assumption about the discount rate applied to future income. We assume a zero discount rate in line with the approximate real return on 15-year government bonds. All of the qualitative conclusions of our analysis continue to hold if we make alternative assumptions.

---

an actuarially fairly priced annuity as paying out at less than a fair rate. The subjective rate is higher than the objective rate in 64% of cases. We can see that, in many cases, what the individual would perceive to be a fair annuity rate is several percentage points higher than the 'objective' annuity rate. This could, therefore, be an important cause of the unpopularity of annuities.

It is important to note that even if an individual perceives an annuity rate to be unfair, they may still want to purchase the annuity. If individuals are sufficiently risk averse, then paying this extra amount may be worth it because of the insurance value provided. Nevertheless, individuals considering buying an annuity presumably have some minimum rate at which they would be willing to buy, which will depend on what they perceive to be a fair deal. If the rate they are offered is lower than this, they will think it not worthwhile buying the annuity.

## 4.2 Deferral of the state pension

In the UK, individuals eligible for the state pension can choose to begin receiving it from their state pension age, but also have the option to defer receipt until a later date. For each month that an individual defers the receipt of their state pension, they receive an increase in the level of their pension once they claim. For the individuals in our sample, deferral of receipt of the state pension by one year beyond their state pension age would have brought them a state pension that was 10.4% higher on average.

While the choice of whether or not to defer receipt of the state pension may seem to be quite a different economic decision from an annuity choice, it is in essence the same. An individual chooses to give up a lump-sum amount in the present day – in this case, the state pension for the months that they defer – and in return they receive a fixed, guaranteed flow of income until they die – in this case, the additional state pension payments. Looking at whether an individual takes up this offer, and whether take-up is related to the individual's subjective annuity rate, is potentially a way of assessing whether survival optimism and pessimism play a role in this economic decision.

Very few individuals choose to defer their state pension at all. In our sample, 4% of individuals who are eligible for the state pension report having deferred take-up. On the face of it, this is consistent with survival pessimism playing a role. One way to assess whether survival pessimism may be driving this low level of take-up is to separate individuals into those for whom the deferral of the state pension looks like an actuarially 'fair deal' and those for whom it looks like an actuarially 'unfair deal'. That is to say, we divide individuals into two groups based on whether their subjective annuity rate is below or above 10.4%.

Table 4.1 shows the rates of deferral of the state pension amongst these two groups. We can see that while the rates are low in both groups, the deferral rate among those who would perceive the deferral to be actuarially 'fair' (3.6%) is almost double that among those who would perceive the deferral rate to be actuarially 'unfair' (1.9%). However, this difference is not statistically significant.<sup>16</sup>

---

<sup>16</sup> We have further examined the propensity to defer the state pension in both linear and probit regression frameworks when controlling for a range of relevant covariates including age, sex, health diagnoses,

**Table 4.1. Rates of deferral of the state pension by whether subjective annuity rate is above or below 10.4%**

<b>Group</b>	<b>Percentage that defer</b>	<b>Percentage that do not defer</b>	<b>Percentage of sample</b>
Subjective annuity rate 10.4% or less ('fair deal')	3.6%	96.4%	96.5%
Subjective annuity rate greater than 10.4% ('unfair deal')	1.9%	98.1%	3.5%
All	3.5%	96.5%	100%

Note: Sample includes all individuals over their state pension age.

Source: ELSA waves 3–7.

While deferral of the state pension is higher amongst those who perceive it to be at least actuarially fair, the overall levels of deferral are too low for us to form any robust conclusions about the relationship between subjective expectations of survival and deferral of the state pension.<sup>17</sup>

There are, of course, a number of potential explanations for this low take-up rate of state pension deferral, despite it appearing to be a good deal for a number of people. If deferring the state pension would mean that an individual would be left without sufficient income to cover their expenditure in the short term (e.g. if they are not in paid work and do not have another source of income or savings), then they may choose not to defer, even if they perceive the rate on offer to be favourable. Social norms and lack of information around claiming the state pension, a preference for cash in the immediate term, and unwillingness to deal with the complexity or hassle of thinking about deferral may also explain this behaviour.

---

education and wealth levels. When controlling for these other variables, the change in deferral rates associated with perceiving deferral as at least actuarially fair is not statistically significantly different from zero.

<sup>17</sup> Hurd, Smith and Zissimopoulos (2004) conduct a similar study looking at deferral of social security in the US. They find that those who report zero probability of survival to older ages are less likely to defer receipt but, otherwise, subjective expectations are not predictive of deferral, even accounting for other relevant factors including wealth.

---

## 5. Conclusions: implications for annuities markets and for policy

This paper has explored information about individuals' subjective expectations of their own survival to older ages, and documented a pattern of significant 'pessimism', on average, about survival to ages up to the late 80s and of mild 'optimism' about the chances of survival to ages from the mid 90s onwards.

Individuals' expectations about their chances of survival to older ages are a crucial component in a range of economic decisions – such as how to save for retirement and how to spend savings once retired – that are of increasing significance as individuals are given more responsibility for and control over their retirement provision.

One significant recent change in the environment faced by individuals in retirement is the introduction of so-called 'pensions freedoms'.<sup>18</sup> What do our findings about individuals' survival expectations suggest about decision-making in annuities markets under these new arrangements? The analysis presented in Section 4 suggests that the overwhelming majority of individuals in this situation are likely to view the annuities on offer – even those that are close to being actuarially fairly priced – as a bad deal. Consequently, many may choose not to purchase annuities.

Evidence that has been gathered since the freedoms were announced has indeed shown a sharp decline in the use of annuities. Before the freedoms, around 90% of pension pots were used to purchase an annuity. Now, individuals are twice as likely to move into income drawdown as to purchase an annuity.<sup>19</sup> Overall, annuity demand was estimated to have fallen by 75% between 2012 and 2016, and several providers have moved out of the annuities market.<sup>20</sup> These trends will undoubtedly have been driven by a number of factors, but survival pessimism may well have played a contributing role.

What might be the implications of survival pessimism in other areas? Looking into the future, individuals are increasingly being expected to play a larger role in providing for their own retirement. As defined contribution (DC) pensions replace defined benefit (DB) pensions for almost all workers, individuals have greater choice about how much to save and how these savings should be invested. An important factor in these choices will be how many years of retirement individuals expect to have to provide for. Survival pessimism could be a concern if it leads individuals to be insufficiently prepared for retirement. As more and more individuals arrive at age 55 with DC pension wealth that they are eligible to access, the role of individuals' expectations about their future and the need to provide for themselves will become increasingly important, particularly if the option to purchase an annuity is no longer available for many individuals. Finally, optimism about survival at the very oldest ages may lead to reluctance to spend remaining wealth if an individual survives through their 80s and into their 90s or beyond.

---

<sup>18</sup> Before 2015, individuals with defined contribution (DC) pension pots larger than £20,000 (but not so large as to guarantee a minimum secure level of income) were taxed at a rate of 55% on all lump-sum withdrawals beyond the one-off 25% tax-free amount. This meant that, in essence, the majority of individuals with DC wealth had to use this to purchase an annuity. Under 'pension freedoms', individuals can draw down accumulated pension wealth and be taxed at their marginal rate of income tax.

<sup>19</sup> Financial Conduct Authority, 2017.

<sup>20</sup> Cannon, Tonks and Yuille, 2016.

---

The findings of this paper are therefore highly pertinent as recent changes to the private pensions environment mature. Policymakers should continue to closely monitor developments in annuities markets and the adequacy of both savings and incomes through the course of retirement.

# Appendix

## A.1 Life tables and associated terminology

The ‘basic’ element of the life table is a mortality probability. This is the probability of dying at age X, conditional on reaching age X. This is shown in column 1 of the example life table in Table A.1. 100% minus this probability is the survival probability conditional on reaching that age (column 2), i.e. the probability of surviving to age X+1 conditional on surviving to age X. A survival curve is the probability, conditional on reaching age X, of surviving to at least each older age. In the example in the table, the individual faces some chance of death at ages 0, 50, 80 and 100 (and no chance of death at other ages). Given a mortality probability of 1% at age 0, someone of age 0 has a 99% chance of surviving to age 50 and an 85.1% (= 99%×86%) chance of surviving to age 80 (column 3). The probability of death at each age is the difference between the survival probability at that age and the next (column 4). Multiplying by that age and summing the results yields life expectancy, as per the example calculation at the bottom of the table.

**Table A.1. Example ‘life table’**

Age	(1) Pr{Die} at age if alive	(2) Pr{Don’t die} if alive	(3) Pr{Survive} to age	(4) Pr{Die} at particular age
0	1.0%	99.0%	100.0%	1.0%
50	14.0%	86.0%	99.0%	13.9%
80	97.0%	3.0%	85.1%	82.5%
100	100.0%	0.0%	2.6%	2.6%

Life expectancy = 1%×0 + 13.9%×50 + 82.5%×80 + 2.6%×100 = 75.5

Source: Authors’ example.

Life tables are available in both ‘cohort’ and ‘period’ form. A cohort life table is based on the mortality probabilities for individuals born in a particular year, e.g. the 1958 table uses mortality for 1-year-olds in 1959, 2-year-olds in 1960 and so on. For ages that the cohort has not yet reached, the table must be constructed using projected mortality rates. These are based on extrapolation of trends in mortality as well as assumptions about how these trends may change in future.

Period life tables are also available but are arguably less relevant in the context of considering individuals’ chances of surviving to older ages. Period life tables are based on the mortality probabilities for individuals at the relevant age in a particular year, e.g. the 2010 table uses mortality probabilities for those who were aged 1 in 2010, aged 2 in 2010 and so on. The survival curves it yields tell us the chances that an individual would face of living to older ages if they experienced the mortality probabilities being experienced by people alive in that year.

## A.2 Individuals giving well-formed answers to survival questions

Here, we provide some additional detail on the issue of whether individuals give well-formed answers to survival questions. Tables A.2 and A.3 show the distribution of individuals giving ‘impossible’ answers, by their level of education and their age respectively.

**Table A.2. Distribution of individuals giving ‘impossible’ answers, by education level**

Education level	100% chance of survival	Higher answer for second question	Total
Low	7.6%	8.1%	14.6%
Mid	6.8%	8.9%	14.3%
High	4.8%	7.8%	11.7%
All	6.8%	8.3%	13.9%

Note: Individuals aged 70 and above are asked only one survival question. Sample includes only those who did not answer ‘Don’t know’ to the first survival question.

Source: ELSA waves 1–7. 66,210 answers from 16,203 unique individuals.

**Table A.3. Distribution of individuals giving ‘impossible’ answers, by age**

Age	100% chance of survival	Higher answer for second question	Total
50s	7.7%	9.3%	15.7%
60s	8.3%	16.0%	22.0%
70s	4.8%	N/A	4.8%
80s	3.8%	N/A	3.8%
90s	8.0%	N/A	8.0%
All	6.8%	8.3%	13.9%

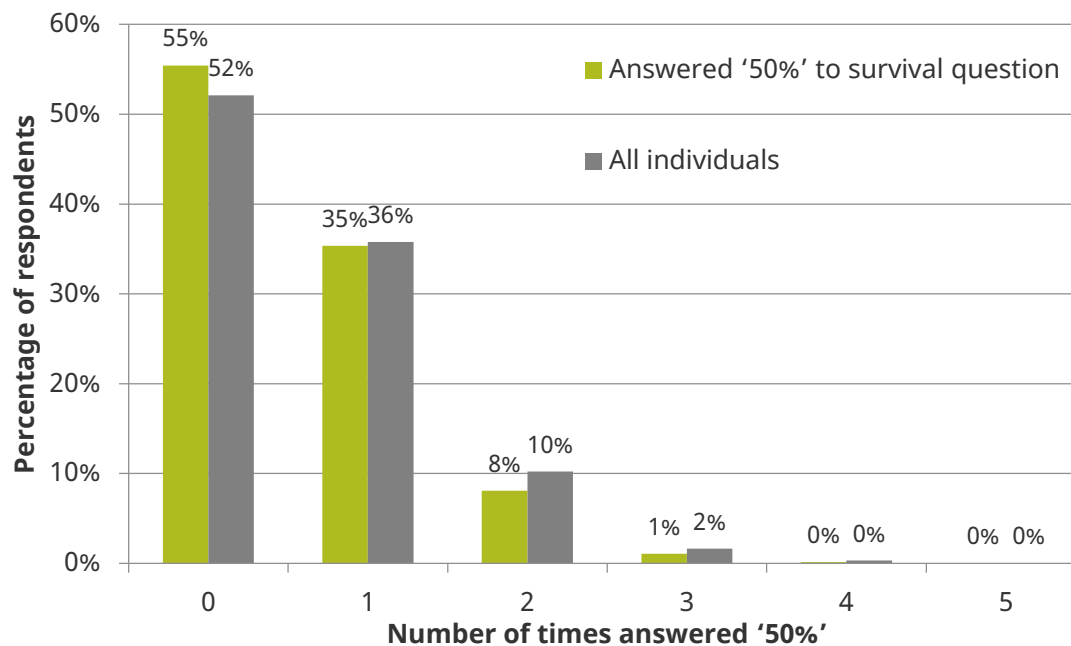
Note: Individuals aged 70 and above are asked only one survival question. Sample includes only those who did not answer ‘Don’t know’ to the first survival question.

Source: ELSA waves 1–7. 66,210 answers from 16,203 unique individuals.

Figure A.1 shows the distribution of the number of times that individuals give an answer of ‘50%’ to the other probability questions in the survey, split by whether or not they answered ‘50%’ to the first survival question.



**Figure A.1. Number of times individuals answered '50%' to other probability questions**



Note: Other probability questions include those related to the probability of moving out of one's home in the future, of being in work in a number of years' time, of having insufficient financial resources to meet needs at some point in the future, of it raining tomorrow and of giving and receiving an inheritance.

Source: ELSA waves 1–7. 66,210 answers from 16,203 unique individuals.

### A.3 Constructing subjective survival curves

For individuals who give responses about the probability of surviving to multiple ages, it is possible to construct a subjective survival curve, if we are prepared to make some further assumptions.

It seems reasonable to assume that the perceived probability of dying at some particular age, conditional on reaching that age, is increasing smoothly with age. In other words, an individual's perceived probability of dying within a year of their 75<sup>th</sup> birthday (assuming they reached it) is slightly higher than their perceived probability of dying within a year of their 74<sup>th</sup> birthday, and so on. It also seems reasonable to assume that the risk of death increases by more with each passing year of age. Objective life tables have these two properties and it seems a weak assumption to suppose that individuals' expectations will behave in this way too.

Given we are willing to make these assumptions, in order to be able to extrapolate to a full survival curve given knowledge of two points on this curve, we must make an assumption about the precise way in which the perceived probability of death increases with age. We make the particular assumption, widely used in the epidemiological literature and for modelling ageing processes generally, that subjective expectations about the probability of death follow a Weibull distribution. Technically, this means that the probability of death at each older age (the 'hazard rate') is the age being considered raised to some power and then multiplied by a parameter. For an individual considering the probability that they

---

will die at age  $A$ , if they reach that age, we assume that their perceived probability of death at age  $A$  is given by

$$P(A, \lambda, k) = \frac{\lambda}{k} \left(\frac{A}{k}\right)^{\lambda-1}.$$

We can think of the parameter  $\lambda$  as determining the ‘rate’ at which the chances of death increase with age (or, equivalently, the ‘shape’ of the survival curve), whereas the parameter  $k$  determines the overall level of the probability of death. This form of ‘hazard rate’ implies a particular form of cumulative distribution function for age of death (and therefore a particular form of the survival curve).

With two observations of perceived survival probabilities, it is possible to infer values for the two parameters of the Weibull distribution and therefore construct a subjective survival curve. However, the estimated survival curves that are constructed when we use this method imply beliefs about survival to very old ages that are implausible. For example, for many individuals, the implied Weibull curve gives them a significant chance of living beyond 110. We think that it is safe to say that this does not represent individuals’ true expectations of survival to this age. We therefore, for each individual, assume that they have an accurate belief about their probability of survival to age 110 (this is a probability below 1% for all individuals in our data). This gives us three ‘points’ from which to infer a survival curve. While two points imply a unique subjective survival curve, there are several subjective survival curves that could be taken to be ‘consistent’ with three points. We select the curve that minimises the sum of squared differences between the curve and the three points used. We estimate the Weibull parameters that minimise this criterion by using an iterative procedure, with the Weibull parameters from the two-point-based curve as initial values.

---

## References

- Association of British Insurers (2017), *The New Retirement Market: Challenges and Opportunities*,  
[https://www.abi.org.uk/globalassets/files/publications/public/lts/abi\\_bro3598\\_retirement\\_market\\_v7.pdfmel4.pdf](https://www.abi.org.uk/globalassets/files/publications/public/lts/abi_bro3598_retirement_market_v7.pdfmel4.pdf).
- Bissonnette, L., Hurd, M. D. and Michaud, P-C. (2017), 'Individual survival curves comparing subjective and observed mortality risks', *Health Economics*, 26, e285–303.
- Bloom, D. E., Canning, D., Moore, M. and Song, Y. (2006), 'The effect of subjective survival probabilities on retirement and wealth in the United States', National Bureau of Economic Research (NBER), Working Paper 12688,  
<http://www.nber.org/papers/w12688>.
- Brown, J. R. (2007), 'Rational and behavioral perspectives on the role of annuities in retirement planning', National Bureau of Economic Research (NBER), Working Paper 13537, <http://www.nber.org/papers/w13537>.
- Cannon, E., Tonks, I. and Yuille, R. (2016), 'The effect of the reforms to compulsion on annuity demand', *National Institute Economic Review*, 237, R47–54.
- Elder, T. E. (2013), 'The predictive validity of subjective mortality expectations: evidence from the Health and Retirement Study', *Demography*, 50, 569–89.
- Financial Conduct Authority (2017), *Retirement Outcomes Review: Interim Report*,  
<https://www.fca.org.uk/publications/market-studies/retirement-outcomes-review>.
- Finkelstein, A. and Poterba, J. (2004), 'Adverse selection in insurance markets: policyholder evidence from the U.K. annuity market', *Journal of Political Economy*, 112, 183–208.
- French, E. (2005), 'The effects of health, wealth, and wages on labour supply and retirement behaviour', *Review of Economic Studies*, 72, 395–427.
- Hurd, M. D. (1989), 'Mortality risk and bequests', *Econometrica*, 57, 779–813.
- Hurd, M., McFadden, D. and Gan, L. (1998), 'Subjective survival curves and life cycle behavior', in D. A. Wise (ed.), *Inquiries in the Economics of Aging*, University of Chicago Press, <http://www.nber.org/chapters/c7089.pdf>.
- Hurd, M. D. and McGarry, K. (1995), 'Evaluation of the subjective probabilities of survival in the Health and Retirement Study', *Journal of Human Resources*, 30, S268–92.
- Hurd, M. D., Smith, J. P. and Zissimopoulos, J. M. (2004), 'The effects of subjective survival on retirement and social security claiming', *Journal of Applied Econometrics*, 19, 761–75.
- Inkmann, J., Lopes, P. and Michaelides, A. (2011), 'How deep is the annuity market participation puzzle?', *Review of Financial Studies*, 24, 279–319.

- 
- Lew, E. A. and Garfinkel, L. (1987), 'Differences in mortality and longevity by sex, smoking habits and health status', *Transactions of Society of Actuaries*, 39, 107–30, <http://www.soa.org/library/research/transactions-of-society-of-actuaries/1987/january/tsa87v397.pdf>.
- Lowe, J. (2014), 'Whither UK annuities?', [http://oro.open.ac.uk/40642/1/ILC%20Whither%20annuities\\_FINAL.pdf](http://oro.open.ac.uk/40642/1/ILC%20Whither%20annuities_FINAL.pdf).
- Manski, C. F. (2004), 'Measuring expectations', *Econometrica*, 72, 1329–76.
- Marmot, M., Oldfield, Z., Clemens, S., Blake, M., Phelps, A., Nazroo, J., Steptoe, A., Rogers, N., Banks, J. and Oskala, A. (2017). *English Longitudinal Study of Ageing: Waves 0-7, 1998-2015*. [data collection]. 27th Edition. UK Data Service. SN: 5050, <http://doi.org/10.5255/UKDA-SN-5050-14>
- Pashchenko, S. (2013), 'Accounting for non-annuitization', *Journal of Public Economics*, 98, 53–67.
- Post, T. and Hanewald, K. (2013), 'Longevity risk, subjective survival expectations, and individual saving behavior', *Journal of Economic Behavior and Organization*, 86, 200–20.
- Wu, S., Stevens, R. and Thorp, S. (2015), 'Cohort and target age effects on subjective survival probabilities: implications for models of the retirement phase', *Journal of Economic Dynamics and Control*, 55, 39–56.
- Yaari, M. E. (1965), 'Uncertain lifetime, life insurance, and the theory of the consumer', *Review of Economic Studies*, 32, 137–50.