

Eliciting Subjective Survival Curves: Lessons from Partial Identification

L. Bissonnette & J. de Bresser

To cite this article: L. Bissonnette & J. de Bresser (2018) Eliciting Subjective Survival Curves: Lessons from Partial Identification, Journal of Business & Economic Statistics, 36:3, 505-515, DOI: 10.1080/07350015.2016.1213635

To link to this article: <https://doi.org/10.1080/07350015.2016.1213635>



© 2018 The Author(s). Published with license by Taylor & Francis© L. Bissonnette and J. de Bresser



Published online: 16 May 2017.



Submit your article to this journal [↗](#)



Article views: 1024



View related articles [↗](#)



View Crossmark data [↗](#)



Citing articles: 2 View citing articles [↗](#)

Eliciting Subjective Survival Curves: Lessons from Partial Identification

L. BISSONNETTE

Departement d'Économique, Université Laval, 1025, Avenue des Sciences-Humaines, Québec (Québec), G1V 0A6, Canada (luc.bissonnette@ecn.ulaval.ca)

J. DE BRESSER

Tilburg University, P.O. Box 90153, Netspar, P.O. Box 90153, 5000 LE Tilburg, The Netherlands (j.r.debresser@uvt.nl)

When analyzing data on subjective expectations of continuous outcomes, researchers have access to a limited number of reported probabilities for each respondent from which to construct complete distribution functions. Moreover, reported probabilities may be rounded and thus not equal to true beliefs. Using survival expectations elicited from a representative sample from the Netherlands, we investigate what can be learned if we take these two sources of missing information into account and expectations are therefore only partially identified. We find novel evidence for rounding by checking whether reported expectations are consistent with a hazard of death that increases weakly with age. Only 39% of reported beliefs are consistent with this under the assumption that all probabilities are reported precisely, while 92% are if we allow for rounding. Using the available information to construct bounds on subjective life expectancy, we show that the data alone are not sufficiently informative to allow for useful inference in partially identified linear models, even in the absence of rounding. We propose to improve precision by interpolation between rounded probabilities. Interpolation in combination with a limited amount of rounding does yield informative intervals.

KEY WORDS: Life expectancy; Partial identification; Subjective expectations.

1. INTRODUCTION

Expectations, especially regarding individual survival, play an important role in inter-temporal models explaining saving and retirement. Early research in the 1990s indicated that expectations can be elicited through probabilistic questions, a single question for a binary event or multiple questions tracing the CDF for a continuous outcome (see Manski (2004) for a general review and Hurd (2009) for a review focused on survival expectations). The two decades that followed have yielded rich descriptions of heterogeneity in and the predictive validity of subjective longevity at the individual level (e.g., Hurd and McGarry 1995, 2002; Smith, Taylor, and Sloan 2001; Bissonnette, Hurd, and Michaud 2014, for the U.S.; and Kutlu and Kalwij 2012, for the Netherlands). Moreover, the common shortcut made in empirical lifecycle models of approximating expectations by means of actuarial tables has been called into question by the finding that especially women expect to die earlier than those figures predict (Perozek 2008; Peracchi and Perotti 2011).

Much of the research on subjective expectations regarding continuous outcomes proceeds by fitting a unique distribution function for each observation, either by means of nonlinear least squares estimation of a parametric distribution (Dominitz and Manski 1997; Dominitz 1998; Dominitz and Manski 2006), or by means of nonparametric splines (Bellemare, Bissonnette, and Kröger 2012). In this article, we propose ways to analyze subjective expectations which take into account that we only have information for some points on the CDF and that the reported probabilities may be subject to rounding, both of which imply that the distributional function of interest is only partially identified. Rounding is the practice of reporting one specific

value whenever a real number lies in an interval (Manski and Molinari 2010), for example, probabilities in $[5, 15)$ are reported as 10%. We either allow for the maximum amount of rounding for each reported probability or infer the extent of rounding from all survival probabilities reported by an individual. Rounding helps to reconcile reported beliefs with the assumption that the true hazard of death should be weakly increasing with age for each respondent: only 39% of responses allow for an increasing hazard if probabilities were reported precisely, while 92% do once we allow for maximal rounding. Given this novel evidence, we construct sets of all functions that are consistent with the data under different assumptions regarding rounding. We use that set to derive bounds on life expectancy (LE). We then illustrate a way to tighten those bounds by applying the spline interpolation technique of Bellemare, Bissonnette, and Kröger (2012) while allowing for rounding. Though our data contain information on CDFs, the methods can also be applied to expectations elicited in terms of the PDF.

This approach of bounding distributions rather than fitting them is inspired by numerous papers by Charles Manski (see, e.g., Manski 2003). In particular, Engelberg, Manski, and

© 2018 L. Bissonnette and J. de Bresser.

Published with license by Taylor & Francis.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License

(<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Business & Economic Statistics

July 2018, Vol. 36, No. 3

DOI: 10.1080/07350015.2016.1213635

Williams (2009) bound measures of the central tendency using data that is similar to ours. Manski and Molinari (2010) studied rounding systematically, but in the context of expectations regarding binary outcomes (a single subjective probability). We depart from these studies in two ways. First, we are interested in bounding an entire survival curve, since the probability to survive to any given age can be a value of interest. Second, we study the implications of rounding in the context of expectations of a continuous variable more thoroughly than is done in Engelberg, Manski, and Williams (2009).

The sample average width of the intervals for LE computed without any assumptions beyond the data and without rounding is 12 years. In partially identified models like Imbens and Manski (2004) and Beresteanu and Molinari (2008), this is too wide to reveal the relationships between LE and the covariates *age* and *health* that are apparent in point-identified models. The combination of interpolation between reported probabilities and limited rounding allows for more precise inference, even to the extent that partially identified models corroborate the findings from point-identified models. Finally, while point-identified models indicate that women expect to die younger than predicted by life tables, we cannot reject that average expectations are in line with life tables once we allow for rounding.

The structure of the article is as follows. Section 2 introduces the type of data that we use and explains the methods that we apply to approximate expectations. The data are described in Section 3 and Section 4 presents our results. Section 5 concludes.

2. METHODS

2.1 Survival Questions

The subjective longevity questions that we analyze are similar to those found in the Health and Retirement Study (HRS). However, in contrast to the two thresholds of the HRS, we consider questions that refer to survival past a maximum of five target ages. The items are phrased as follows:

Indicate on a scale from 0 to 100 how likely you think it is that you live to:

1. [if age < 69] age 70
- ⋮
5. [if age < 89] age 90

Age-eligibility requires a respondent to be at least 2 years younger than a particular target age in order for that question to be presented.

2.2 Point Identification of Survival Functions

We use two methods to fit individual survival curves to the probabilities reported by each survey respondent. The first fits parametric distributions by nonlinear least squares (Dominitz and Manski 1997; Perozek 2008). The resulting survival functions generally do not pass through the reported probabilities, but they are as close as possible according to the least-squares criterion. Following Perozek (2008), we fit Gompertz and Weibull distributions.

The second method, proposed by Bellemare, Bissonnette, and Kröger (2012), uses spline interpolation to construct survival

functions that are not restricted to a certain parametric family. This method uses smooth piecewise polynomial functions to approximate the subjective survival curves. It also enforces monotonicity of the function, an important property of survival curves. The resulting spline function passes *through* all reported probabilities for a given respondent. Hence, it does not require the researcher to choose a bandwidth or other smoothing parameter as would be the case for spline smoothing. Spline interpolation does not impose any restrictions on the shape of the hazard. We apply linear and cubic splines, because the former preserve shape while the latter have been shown to be able to approximate parametric distributions closely (Bellemare, Bissonnette, and Kröger 2012).

In this article, we assume a maximum age of 110 when calculating remaining life expectancy. Sensitivity checks indicate robustness of all our results to alternative maximum ages of 100 and 120 (results available on request).

2.3 Partial Identification of Survival Functions Without Rounding

The methods discussed in the previous subsection use reported probabilities to construct a survival function for each individual in the sample. We next consider admissible regions for survival functions without such smoothing between data points. For now we do maintain the assumption that reported probabilities are not rounded and do not contain measurement error. The gray area in Figure 1 is the admissible set for survival expectations under the assumption that hypothetical reported probabilities lie exactly on the subjective survival curve. Under this assumption, our data identify rectangles within which the subjective survival curve lies, but contain no information on the location of the curve within those rectangles. Note that any function is allowed, even a step function, as long as it passes through the points given in the data.

Calculating bounds on life expectancy is straightforward: we trace the bottom edges of the rectangles in Figure 1 to obtain the most pessimistic survival curve and thus the lower bound on life expectancy that is consistent with the data. Likewise, the upper edges yield the most optimistic curve corresponding to the upper bound on life expectancy. Figure 1




Figure 1. Admissible set for the survival curve (no smoothing between data points and no rounding).

illustrates the method for survival functions, since the survey items we analyze are phrased in terms of survival. However, one could follow the same approach to construct bounds on the CDF.

2.4 Rounding

2.4.1 General and Common Rounding Schemes. The second motivation for our analysis is the possibility of rounding of reported probabilities. Rounded probabilities are not necessarily on the subjective survival curve. Instead, they are informative of intervals within which the true subjective probabilities fall. For instance, a probability equal to 20% that is rounded to a multiple of 5 indicates that the true probability is in the interval [17.5, 22.5), and the same probability rounded to a multiple of ten means that the true probability is in [15, 25). Since a reported probability may result from different degrees of rounding, one has to make an assumption on the extent of rounding that is present in the data.

As a conservative choice, we allow that reported probabilities are rounded one-by-one to the maximum extent. We allow for rounding to multiples of 100, 50, 25, 10, and 5. That is, a reported probability of 100 is interpreted as evidence that the true probability lies in the interval [50, 100] and a reported 35 implies the interval [32.5, 37.5), regardless of the other probabilities reported by that respondent. Probabilities that can only result from rounding to a multiple of 1 are interpreted as indicative of an interval with width 5, so a reported probability of 37% yields the interval [35, 40). Because this scheme allows probabilities reported by an individual to be rounded differently, we call it the *general* rounding scheme.

The general rounding scheme leads to broad intervals for some probabilities, especially for those equal to 0, 50%, and 100%. However, imposing monotonicity on the true probabilities helps to narrow down the bounds. For instance, if we observe a probability equal to 50% for the first age threshold of 70, the general rounding scheme interprets that probability as indicative of an interval equal to [25, 75) for the corresponding true probability. However, if the reported probability for age 75 is 40%, we know that the true probability for age 70 cannot be smaller than the lower bound of 35% (since that is the lower bound on the probability for age 75).

A second assumption is that all probabilities reported by an individual are rounded to the same extent. Under that assumption we can apply the strategy proposed in Manski and Molinari (2010) to infer the degree of rounding from a set of probability questions. That is, we assume that the answers to all survival questions from a given respondent are rounded similarly and select the most conservative rounding rule that is consistent with all those probabilities. We allow for rounding to multiples of 100, 50, 10, 5, and 1 as well as more precise reporting of probabilities close to 0 and 100% (see Manski and Molinari (2010) for a formal definition of this rounding scheme). We call this the *common* rounding rule.

2.4.2 Consistency of Reported Probabilities with Increasing Hazard of Death. As explained in the previous section, most reported probabilities may result from different degrees of rounding. One way to quantify the extent of rounding is to investigate how much rounding is needed to reconcile expectations with plausible assumptions. For instance, we may assume that respondents believe that their survival function has an increasing hazard. This would imply that the probability of dying in a given time interval increases with age. Note that the assumption of increasing hazards does not come from the data and that some respondents may have expectations that are not consistent with it. For instance, a violation would occur if someone believes that if he lives to be 80, he will certainly make it to 100. We can nevertheless check the fraction of respondents that reported probabilities that are consistent with this assumption depending on the extent of rounding that we allow for.

Remember that a continuous survival curve for an individual aged a can be expressed in terms of a hazard function $\lambda(t)$

$$S(t) = \exp\left(-\int_a^t \lambda(u) du\right)$$

or in terms of the integrated hazard function $\Lambda(t)$

$$S(t) = \exp(-\Lambda(t)).$$

While a survival function must be decreasing in t , the hazard function can take many forms. The assumption that a respondent has an increasing hazard function, so that the probability to pass away in a given interval increases with age, would lead to a convex integrated hazard function.

A given rounding scheme yields bounds for the true survival probability at each age threshold, which translate into bounds for

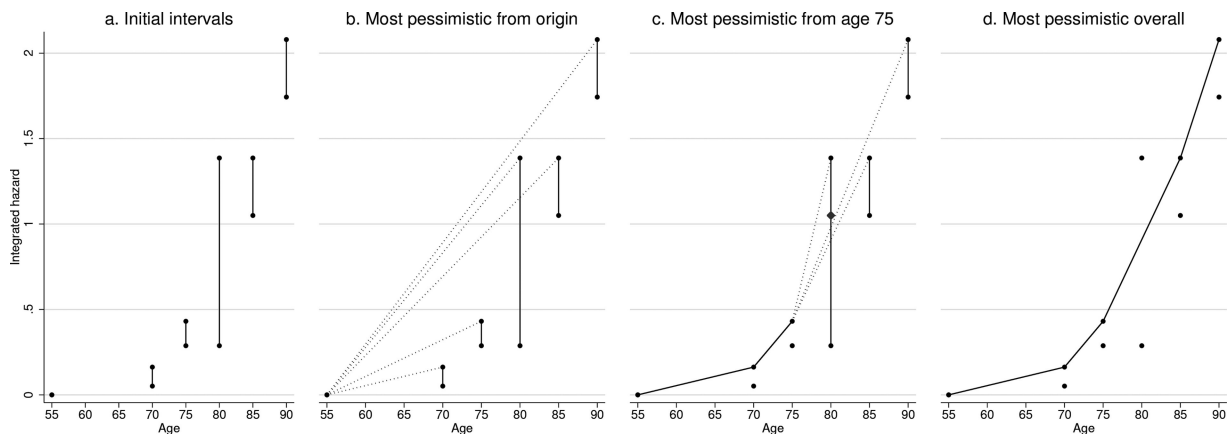


Figure 2. Checking the increasing hazard assumption by constructing a convex integrated hazard function.

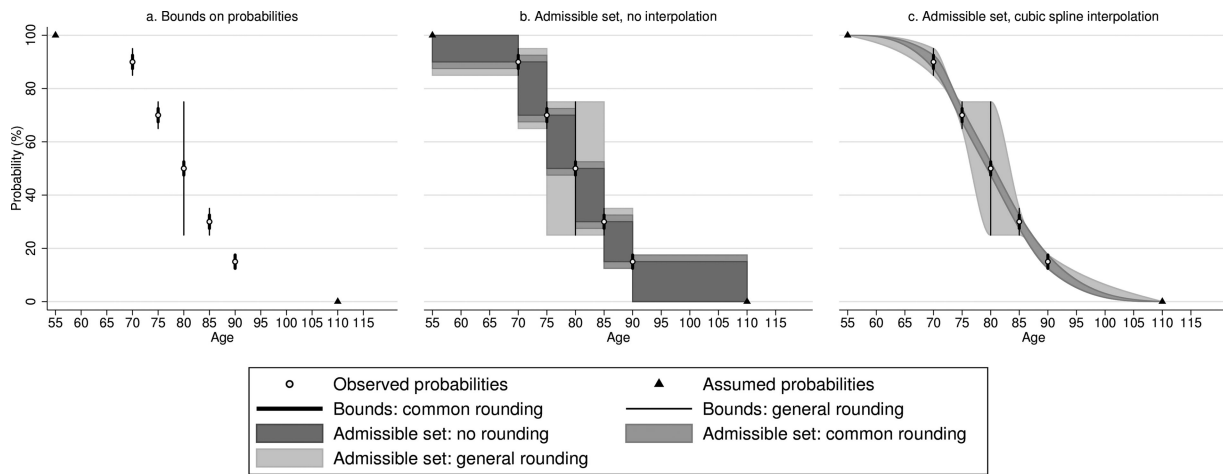


Figure 3. Admissible sets for the survival function under different rounding rules.

the integrated hazard. Figure 2(a) shows these bounds as vertical bars. Since a weakly increasing hazard of death is equivalent to a convex integrated hazard, we check whether there exists a convex piecewise linear function that passes between the bounds for the integrated hazard at all target ages.

We verify whether reported probabilities are consistent with an increasing hazard by constructing the highest integrated hazard, corresponding to the lowest expected lifetime, that is allowed by the data. Finding that the most pessimistic function has an increasing hazard is a necessary and sufficient condition to show that the set of functions compatible with the assumption is not empty. Moreover, our procedure based on the highest integrated hazard is relatively straightforward, as it only requires comparing the slopes of a series of linear functions as discussed below. We present the intuition here and the detailed algorithm in the Appendix. Our method works piecewise, starting from the first interval and concentrating on the largest admissible hazard for each set of bounds. In our example, Figure 2(b) illustrates that the highest integrated hazard follows a steady increase from the origin to the upper bound at age 70. Starting from this point, we can repeat this procedure starting from the maximum hazard at age 70.

Consider now the interval from age 75, as illustrated in Figure 2(c). Over this interval, an increase to the upper bound at age 80 is not admissible. Such an increase in the integrated hazard would lead to a case where the integrated hazard would not increase between age 80 and 85, which violates the convexity assumption. It follows that the highest admissible curve is the linear increase in integrated hazard to the upper bound at age 85. The resulting piecewise linear integrated hazard is presented in Figure 2(d).

In what case would there be no admissible curves? Suppose that the lower bound at age 80 was the same as the one at age 85, as represented by the diamond on Figure 2(c). In this case, there would be no way to find a convex function that would fall within the admissible interval at age 80. Reported probabilities are not consistent with a weakly increasing hazard whenever that is the case. Note how increasing the width of the admissible intervals, for instance by allowing more rounding or reporting error, can reconcile reported probabilities with the increasing hazard assumption.

The alternative approach of studying the most optimistic admissible hazard is more complex, as each interval must now be characterized by two slopes: the slowest admissible increase at the beginning of the interval and the fastest increase admissible at the end. This envelope of admissible optimistic hazards is usually a kinked convex function within each interval, but it is not a convex function taken as a whole. It is therefore simpler and more elegant to focus on the most pessimistic integrated hazard.

2.4.3 Admissible Sets for Survival Functions in the Presence of Rounding. According to the models of rounding explained above, probabilities reported by a respondent are either all generated by a common rounding rule, or they are rounded one-by-one to the maximum extent. We construct admissible sets for the corresponding survival function by tracing the upper and lower bounds of the intervals for the unobserved true probabilities.

Figure 3 illustrates how we construct admissible sets for the survival function under both rounding rules for hypothetical data. Figure 3(a) shows the intervals for the true probabilities under common rounding (thick bars) and general rounding (thin bars). Figure 3(b) explains how the admissible set changes when we introduce rounding in the reported probabilities: the rectangles stretch vertically so they overlap around the reported probabilities. The lightest gray band traces the upper and lower bounds on probabilities constructed using the general and conservative rule that probabilities are rounded individually, while the darker area assumes common rounding. The width of the identified region differs between rounding rules only at those thresholds for which the reported probability is not only a multiple of 5, which is the extent of rounding according to the common rounding scheme, but also a multiple of 10 or even 50 (all ages except age 90). The survival function that yields the highest life expectancy that is consistent with the data is that which traces the upper edge of the identified region, while the lowest life expectancy is obtained by tracing the lower edge.

While we focus on remaining life expectancy, a quantity of interest in the context of subjective survival, bounding the entire distributional function allows us to derive bounds on other aspects of expectations as well. For instance, one can derive bounds on the probability to survive past any age, or equiva-

Table 1. Descriptive statistics of the reported survival probabilities and life table (LT) probabilities

	Men				Women			
	N	Mean LT	Reported probs.		N	Mean LT	Reported probs.	
			Mean	S.D.			Mean	S.D.
Age 70	631	87.0	76.3	21.0	526	90.9	75.3	21.5
Age 75	713	78.2	67.5	23.7	575	85.2	67.7	23.8
Age 80	769	63.4	54.7	25.8	608	75.3	56.2	26.4
Age 85	814	44.0	40.8	26.2	620	58.5	42.7	26.8
Age 90	823	21.8	26.8	25.0	624	35.5	27.5	25.6

NOTE: Probabilities expressed as percentages.

lently on any percentile of the distribution. Having bounded the 25th and 75th percentiles, one can also compute bounds on the interquartile range as a measure of dispersion.

2.4.4 Refinement of the Admissible Sets. The identified regions in Figure 3(b) reflect a worst-case scenario in which we do not assume anything about expectations beyond what is given in the data. As a result, those regions include survival functions that are unlikely to reflect individuals’ beliefs, such as step functions. Therefore, we propose to refine admissible sets by interpolation between rounded probabilities. Rather than constructing rectangles identified by the data, we now trace the upper and lower bounds of the intervals for true probabilities by means of spline functions. This method is shown in Figure 3(c), which illustrates that such smoothed regions are much smaller than the corresponding regions constructed without interpolation shown in panel (b). Moreover, interpolation allows us to focus on the ambiguity caused by rounding of the reported probabilities, as opposed to the inherent coarseness of observing only a few points on the survival function.

3. DATA QUALITY AND DESCRIPTIVES

We use the 2011 wave of the yearly *Pensioenbarometer*, a survey administered to the respondents of the CentERpanel. Data collection was financed by Netspar, Network for Studies on Pensions, Aging and Retirement. The CentERpanel is administrated by CentERdata and is representative of the Dutch adult population. The sample consists of approximately 2500 respondents age 16 and older, but due to the focus on pensions the *Pensioenbarometer* surveys are only elicited from respondents who are older than 24. All CentERpanel questionnaires are administered via the Internet and members of the panel without Internet access are provided with a set-top box to maintain representativeness. Data access can be obtained from CentERdata, via CentERdata@uvt.nl.

The 2011 *Pensioenbarometer* was distributed to 2396 potential respondents and was returned by 1577 of them (66% survey response). Item nonresponse to the subjective survival questions is not an issue: 95% of the panel members who filled out the questionnaire provided an answer to all survival questions. Furthermore, violations of the arithmetic of probabilities are rare despite the fact that no safeguards were applied to ensure logical consistency: reported probabilities decrease weakly with target ages for 97% of the complete responses (note that this does not imply that the hazard of death increases weakly with age). Because of these observations, we do not model item nonresponse or logical inconsistencies. After removing incomplete or logi-

cally inconsistent responses, we are left with 1447 observations (two 89 year old respondents were dropped due to ineligibility for all survival questions).

As for demographic characteristics, about half of the sample is male and the average age is 56. Three quarters of the sample lives with a partner. Respondents are mostly healthy: 75% rate their own health as either “excellent” or “very good,” while only 8% are in “bad” health. 41% have finished some form of higher education (university or an applied college). A further breakdown by sex shows that the sample is better educated than the Dutch average: 44% of men and 38% of women have completed higher education compared to nationwide averages of 31% and 26% respectively in 2009. The average gross personal income in our sample is close to that of the population at large: the economically active within the sample earn 2978 euros per month compared with a national average of 2900 euros in 2010.

Table 1 presents descriptive statistics of reported probabilities and corresponding probabilities from the 2010 life tables assembled by Statistics Netherlands. We match life tables to respondents based on gender and age, so differences between the age distribution in the sample and in the Dutch population do not affect the comparison. Due to the age-eligibility criteria described in Section 2.1, the sample sizes are larger for questions referring to older ages. The average reported probability of survival decreases with the age thresholds from around 75% for age 70 to 27% for age 90. Moreover, the average reported probabilities of men and women are similar for all age thresholds, whereas the average life table probabilities strongly favor women. Table 1 shows that men on average underestimate their probability of living past ages 70, 75, and 80 by roughly 10%-points, while they overestimate their probability of living past age 90 by 5% points. Women, on the other hand, report probabilities that are 8–20% points below the life tables for all target ages. This suggests that the average life expectancy of men is more in line with that reported in life tables than that of women.

3.1 Rounding

The distribution of rounding according to the common rounding scheme is given in Table 2(a). The leftmost column shows that rounding is important: multiples of 5 and 10 account for 51% and 33% of responses, respectively. Cruder forms of rounding do occur but are rare: 6% of respondents round to multiples of 50 or 100. Focal 50/50s are not likely to be an important concern for our data, since no more than 3% of the respondents answer fifty percent to all questions. Confirming the analysis in Manski and Molinari (2010), we find some evidence that re-

Table 2a. Common rounding with different sets of probabilities

	2011 surv. (5 probs.)	2012 surv.; 2011 surv. (10 probs.)	other 2011; 2012 surv.; 2011 surv. (26 probs.)
All 0 or 100	2	1	1
All 0, 50 or 100	4	3	1
All multiples of 10	33	24	18
All multiples of 5	51	57	55
Some in [1, 4] or [96, 100]	7	11	19
Other	3	5	7
Total	100%	100%	100%

NOTE: $N = 1447$ individuals.

Table 2b. General rounding

Multiples of...	Frequency
...100	11
...50	14
...25	9
...10	50
...5	12
...1	3
Total	100%

NOTE: $N = 6703$ probabilities.

spondents may round probabilities near the extremes of zero and one hundred differently: 7% of the sample reports multiples of one near the extremities of the scale. Only 3% of the responses are incompatible with all other forms of rounding and are thus interpreted as exact answers.

In the remaining columns of Table 2(a), we use additional probabilities to determine the extent of (common) rounding. Given that the most precisely reported probability sets the level of rounding for a respondent, these additional probabilities necessarily reduce rounding. We introduce extra probabilities in two steps. First, we add all five survival probabilities that were elicited in 2012, the only other year in which they were asked. Second, we also take into account all other probabilities reported in the 2011 survey, which pertain to future purchasing power and income replacement rates at retirement. Doing so reduces the fraction that rounds to multiples of 10 from 33% to 18% and increases the fraction in the two most precise categories from 10% to 26%. However, 55% of respondents still round to multiples of five.

Table 2(b) describes rounding according to the general rounding scheme. This scheme also indicates substantial rounding: 50% of reported survival probabilities in 2011 are multiples of 10 (but not of 50 or 100). Another 14% of probabilities are equal to 50 and 3% can only result from rounding to multiples of 1. Though not shown in Table 2(b), the variation in frequencies of the different categories across the age thresholds accords with intuition. Respondents express greater certainty near the extreme ends of the age range, while 50/50s are more prevalent at the ages 80 and 85.




Figure 4. Fraction of the sample that is consistent with a weakly increasing hazard of death when allowing for different amounts of noise in reported probabilities.

3.2 Consistency With Increasing Hazard

To quantify the extent to which probabilities are rounded, we use the algorithm discussed in Section 2.4.2 to check whether the data are consistent with an increasing hazard of death under different degrees of rounding. Without rounding, only 39% of respondents report probabilities that are consistent with that assumption. However, that fraction increases to 76% under common rounding, 70% if we use all 26 reported probabilities, and 92% under general rounding. If we maintain that true, unobserved beliefs exhibit a weakly increasing hazard, this is strong evidence that reported probabilities are rounded or that there is another form of measurement error or ambiguity.

To illustrate the implications of such measurement error, we replaced the rounding intervals by intervals created by adding and subtracting an error of given size from all reported answers. Figure 4 shows how the fraction with a weakly increasing hazard increases when we allow for noise in the reported probabilities. Small errors are sufficient to render most of the data compatible: 80% of the sample is compatible with an increasing hazard if we allow for an error of 3.5%-points in both directions around each reported probability. The common and general rounding schemes yield similar rates of increasing hazard-consistent responses as do errors of 3 and 6% points, respectively, suggesting that this is a reasonable range for the magnitude of errors. If we increase all reported probabilities by 5%-points, this would yield a decrease in life expectancy of 1 year over the 20-year interval between the ages 70 and 90. Given the difficulty of predicting one's own survival, this seems like a reasonable error.

4. RESULTS

4.1 Point- and Interval Identification of Life Expectancy

Table 3 contains descriptive statistics of the point-identified expected ages of death calculated from the individual-specific parametric survival curves and the spline survival functions. Depending on the approximation method, men expect to live to age 82–82.5 on average, which is slightly below the average prediction of 83 years found in the life tables. For women,

Table 3. Point estimates of the expected age of death

	N	Life tables	Parametric		Splines	
			Gompertz	Weibull	Linear	Cubic
Men	823	83.0 (2.3)	82.0 (7.6)	82.5 (7.4)	82.4 (8.1)	82.2 (8.2)
Women	624	85.8 (1.5)	81.8 (7.8)	82.4 (7.6)	81.9 (8.3)	81.7 (8.6)

NOTES: Reported numbers are averages; standard deviations in parentheses.

we find a larger discrepancy between the average subjective life expectancy and the average actuarial forecast: women expect to live to age 81.7–82.4, while the average actuarial prediction is 85.8. This larger discrepancy for women was also observed by Perozek (2008) for the U.S. and Kutlu and Kalwij (2012) for the Netherlands. If we divide the sample in age groups, we find that men and women of all ages expect to live shorter than predicted in the actuarial tables. A formal comparison of average subjective life expectancy and the forecasts in life tables can be found in Section 4.3.

Unsurprisingly, expectations exhibit much more variation than the life tables (conditional on gender the latter only vary with age). The standard deviations of the expected age of death are around 8, while that of the actuarial forecasts is 2.3 for men and 1.5 for women. In the next subsection, we analyze whether this variation in expectations is related to covariates such as health and socio-economic status. The life expectancies calculated using the two methods to point identify survival functions are very similar: all correlations between the expected ages of death are above 0.97. Hence, given sufficiently rich data, computed mortality expectations are robust with respect to the choice for a (non-)parametric model for subjective survival.

Now, we turn to the bounds computed using the methodology described in Sections 2.3 and 2.4. Table 4 presents sample averages of the bounds on life expectancy. Panel (a) reports descriptives for the baseline case, while panel (b) applies the refinement from Section 2.4.4. Table 4 contains averages for bounds under the assumption of no rounding and under the common and general rounding rules described in Section 2.4.1. For the common rounding scheme we either infer rounding from

the five survival questions from the 2011 survey (5 probs.) or from all subjective probabilities from the 2011 wave plus the survival questions from the 2012 wave (26 probs.).

According to the baseline bounds, which do not impose any restrictions on expectations beyond being consistent with the data, both men and women expect to live to age 76–88 on average. Hence, the average interval for life expectancy is about 12 years wide. By definition rounding makes the estimated intervals wider and thus less informative. Assuming a common rounding rule based on the five 2011 survival probabilities, we compute bounds with an average width of 15 years. This width is reduced by 1 year if we assume that all 2011 probabilities and the 2012 survival questions are rounded to the same extent. The general rounding scheme yields slightly wider intervals with an average width of 18 years. The number of probabilities observed per respondent is a key determinant of the width of the bounds. For instance, if we restrict the set of probabilities to target ages 75 and 85, as would be the case in the HRS for younger respondents, the average width of the intervals without any rounding increases to 19–20 years. Hence, the intervals computed from two probabilities without rounding are wider on average than the intervals derived from a set of five probabilities under conservative rounding.

We obtain much tighter bounds on life expectancy if we interpolate expectations between the elicited survival probabilities, as can be seen in panel (b). Under the common rounding assumption based on five probabilities, interpolation reduces the average width of the interval from 15 years to 3 years. Including all probabilities in the rounding scheme yields a further reduction in width of 1 year on average. General rounding leads to wider intervals: the average width is 7–8 years. Hence, regardless of whether we smooth expectations, the specific type of rounding that we assume strongly affects the informativeness of the data.

4.2 Life Expectancy and Demographic Variables

Next, we investigate the relationship between life expectancy and demographic variables. We estimate linear regressions with the point-identified life expectancies as dependent variables. For the bounds we apply partially identified models according to the

Table 4. Sample averages of bounds on expected age of death derived in the absence of rounding and under common and general rounding

	Men				Women			
	No rounding	Common 5 probs.	Common 26 probs.	General	No rounding	Common 5 probs.	Common 26 probs.	General
a. Baseline								
LB	76.7	75.4	75.8	73.8	75.9	74.2	74.7	72.4
UB	88.0	89.6	89.1	91.4	88.0	89.7	89.2	91.7
UB – LB	11.3	14.2	13.3	17.6	12.1	15.4	14.4	19.2
N	823	823	823	823	624	624	624	624
b. Refinement: smoothing using cubic splines								
LB	–	80.6	81.2	78.5	–	80.0	80.4	77.6
UB	–	83.6	83.1	85.5	–	83.3	82.8	85.5
UB – LB	–	3.0	1.9	7.0	–	3.3	2.3	7.9
N	–	820	814	643	–	623	619	479

Table 5. Point and partially identified models of remaining life expectancy

	Point-identified models		Partially identified models					
			No smoothing		Cubic splines			
	Weibull	Cubic spline	No rounding		Common rounding		General rounding	
			LB	UB	5 probs.	5 probs.	LB	UB
Age	-1.573*** (0.0927)	-1.562*** (0.110)	-4.277 (-4.524; 1.451)	1.204	-2.271 (-2.491; -0.611)**	-0.831	-3.129 (-3.378; 0.479)	0.229
Age squared/100	0.656*** (0.0791)	0.692*** (0.0934)	-1.704 (-1.912; 3.275)	3.068	0.0641 (-0.121; 1.487)	1.303	-0.830 (-1.036; 2.262)	2.056
Male	-0.724* (0.385)	-0.605 (0.425)	-12.434 (-13.235; 11.871)	11.070	-3.731 (-4.560; 3.436)	2.607	-8.301 (-9.256; 7.707)	6.752
Educ. primary school	-0.733 (0.937)	-0.649 (0.988)	-12.766 (-14.473; 13.317)	11.610	-3.603 (-5.430; 3.813)	1.986	-9.276 (-11.287; 7.169)	5.159
Educ. higher secondary	-0.123 (0.664)	0.0476 (0.742)	-12.709 (-13.950; 13.943)	12.703	-3.441 (-4.823; 4.520)	3.138	-8.541 (-10.066; 9.204)	7.680
Educ. lower vocational	0.0957 (0.610)	0.0358 (0.687)	-13.907 (-15.274; 15.405)	14.038	-4.135 (-5.454; 5.242)	3.922	-9.723 (-11.220; 10.574)	9.077
Educ. higher vocational	-0.295 (0.524)	-0.304 (0.571)	-13.517 (-14.697; 13.954)	12.773	-3.852 (-4.941; 4.061)	2.972	-9.059 (-10.335; 8.956)	7.679
Educ. University	0.520 (0.599)	0.364 (0.648)	-14.087 (-15.508; 16.346)	14.925	-3.736 (-5.021; 5.226)	3.941	-8.994 (-10.413; 10.659)	9.240
Income < 1150 euro/month	-0.371 (0.935)	-0.377 (1.027)	-14.816 (-16.753; 15.927)	13.990	-4.401 (-6.402; 5.579)	3.578	-9.674 (-11.788; 11.195)	9.080
Income 1801–2600 euro/month	0.0523 (0.615)	-0.0288 (0.687)	-12.528 (-13.634; 13.787)	12.682	-3.576 (-4.861; 4.560)	3.275	-7.739 (-9.004; 9.502)	8.237
Income > 2600 euro/month	-0.0280 (0.576)	-0.00580 (0.648)	-13.064 (-14.330; 14.330)	13.064	-3.567 (-4.758; 4.537)	3.346	-8.187 (-9.708; 9.982)	8.462
Bad health	-6.629*** (0.915)	-7.537*** (1.025)	-23.524 (-25.368; 10.653)	8.809	-11.382 (-13.126; -0.771)**	-2.516	-16.070 (-18.050; 6.016)	4.036
Fair health	-5.493*** (0.697)	-6.268*** (0.775)	-21.161 (-22.738; 10.346)	8.769	-10.057 (-11.397; -0.899)**	-2.238	-14.627 (-16.186; 5.008)	3.449
Good health	-1.719*** (0.588)	-2.179*** (0.682)	-15.342 (-16.709; 12.676)	11.308	-5.743 (-7.002; 2.839)	1.580	-10.106 (-11.592; 8.024)	6.537
Constant	95.63*** (2.729)	93.80*** (3.297)	15.615 (7.981; 177.519)**	169.885	73.179 (66.469; 120.343)**	113.633	41.179 (33.807; 145.450)**	138.078
Observations	1446	1446	1446		1443		1122	
R-squared	0.752	0.692						

NOTES: Standard errors in parentheses for point-identified models.
95% confidence sets in parentheses for the partially identified models.
***significant at 1%; **significant at 5%; *significant at 10%.

methods presented in Imbens and Manski (2004). We estimate the interval regressions using the Stata program CIID, presented in Beresteanu, Molinari, and Steeg Morris (2010). Estimation results are presented in Table 5. The model specifications in that table pool men and women, because we could not reject the null hypothesis of equal coefficients for the sexes.

Point-identified models show that age and self-reported health are the most important covariates of subjective life expectancy. As expected, remaining life expectancy decreases nonlinearly with age. Health too is strongly related to subjective life expectancy: compared to the baseline of people in excellent health, those in bad health expect to live 7 years shorter on average. Respondents in fair or good health also expect to live shorter than their healthier peers. Like Kutlu and Kalwij (2012), we do not find significant associations between life expectancy and education or income if we condition on subjective health. This lack

of an association between socio-economic status and subjective mortality is plausible in the context of the Dutch healthcare system, because the quality of medical services is roughly the same for everybody. However, poorly educated and income-poor individuals are especially likely to be in worse health, partly because they are more likely to engage in behaviors that affect their health negatively (like smoking and drinking). Therefore, we find that respondents from the lowest income group and those with the poorest education expect to die younger if we remove subjective health from the estimated equation (estimates available on request). Note that the estimates are not sensitive to the way in which expectations are approximated: similar conclusions emerge whether we fit Weibull distributions or cubic splines.

The rightmost columns of Table 5 present estimates from models with interval-censored life expectancy as the dependent

variable. The estimates from partially identified models estimated on the bounds without smoothing show that little can be learned about variation in life expectancy across the sample if one is unwilling to interpolate expectations between data points. There are clear differences between the identified sets that are in line with the coefficient estimates from point identified models. For instance, the set-estimate of the coefficient of being in bad health, which is associated with a 7-year lower life expectancy relative to the baseline in point identified models, is $(-23.5; 8.8)$. Though that interval suggests a negative correlation, zero is included in all identified sets (and by implication in all 95% confidence sets) except for the constant. Hence, we are not able to draw conclusions about the sign of any of the coefficients: a clear indication that the bounds are too wide for useful inference. Since rounding only makes the bounds on life expectancy wider, this also holds if we allow for rounding. Without additional assumptions there is not enough information in the data to do meaningful inference on differences in life expectancy across socio-demographic groups.

Table 5 shows that if we do smooth expectations by means of cubic splines and allow for rounding according to the common rounding rule, we corroborate the patterns found in point-identified models. In particular, zero is not included in the 95% confidence sets for the coefficients on the indicators for bad and fair health. These results show that the limited number of elicited points on the survival functions is a more important limitation on the informativeness of the data than (common) rounding. Inference can be made slightly more precise if we infer rounding from the full set of 26 probabilities (estimates available on request). However, if we simultaneously interpolate and allow for general, worst-case, rounding, the identified sets again become too wide to draw any conclusions. The informativeness of the data hinges not only on our willingness to smooth beliefs, but also on the particular type of rounding that we assume.

The results for the mean shown in Table 5 also apply to other aspects of the distribution of individual survival, such as the probability of surviving past age 90 and the subjective median. For those features, we confirmed that common, but not general, rounding in combination with interpolation between elicited probabilities allows one to retrieve the sign of associations in partially identified models. For the probability of living past age 90, for instance, we find that respondents in bad health report a 21%-points lower probability on average, with corresponding identified sets of $(-31, -8)$ under common rounding and $(-41, 6)$ under general rounding. The estimated coefficients of models explaining the median are similar to those for the mean. Estimates are available on request.

Finally, note that the sample sizes for the partially identified models of smoothed expectations are smaller than those for the other models. This is because the cubic spline functions that trace the upper and lower bounds for the true, nonrounded probabilities sometimes cross each other, in which case we drop that observation. Such crossing is rare for the case of common rounding, only 3 observations are lost this way, but quite common for general rounding for which we lose 324 observations. Linear splines do not suffer from this complication: if for two consecutive upper bounds we have that $UB_1 \geq UB_2$ and for the corresponding lower bounds we have $LB_1 \geq LB_2$, this implies that any convex combination of the two upper bounds is also at or

above any convex combination of the lower bounds. Therefore, we verified the patterns from Table 5 using linear splines. All estimates are similar to those reported here, results are available on request.

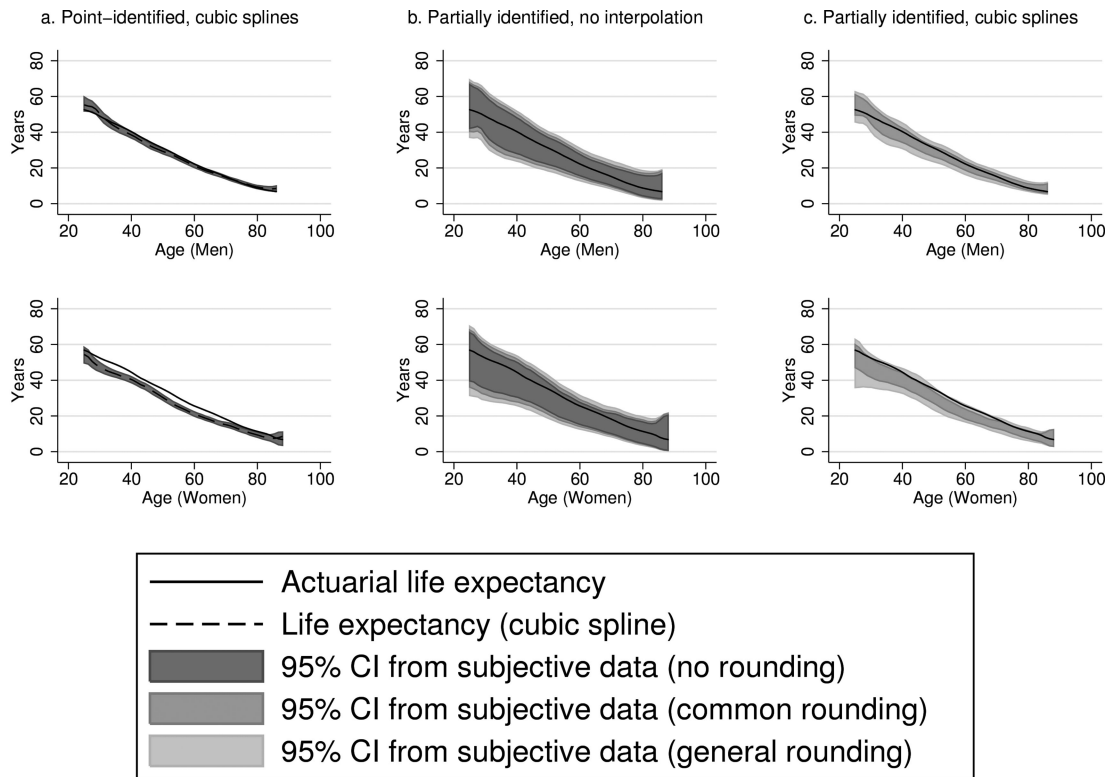
4.3 Expectations and Life Tables

As mentioned in the introduction, comparison of subjective expectations with published life tables allows us to evaluate the use of actuarial figures as proxies for expectations in dynamic economic models. The underlying assumption is that published life tables are adequate proxies for subjective expectations held by the economic agents. One way to assess this assumption is to compare our estimated subject life expectancies with the official tables published by Statistics Netherlands as of December 2010.

We must mention two important limitations before proceeding to this analysis. First, the aim of this exercise is not to determine whether agents are rational. Expecting an earlier demise than what is predicted in the life tables may be justified if an agent has private information regarding his own health and family health history. Second, published life tables are themselves estimates based on observed demographic tendencies. As such, they should be treated as variables with their own prediction uncertainty. However, to our knowledge, Statistics Netherlands does not provide information on the uncertainty regarding these forecasts. For this reason, we will treat them as fixed in the following analysis. It is to be noted that empirical economists rarely include this uncertainty in their models.

Let us first we look at the point-identified life expectancies. Figure 5(a) presents kernel regressions of subjective remaining life expectancy, computed from cubic splines under the assumption of no rounding (dashed line), and the actuarial forecast based on past mortality (solid line). The top graph shows that for men average expectations are close to actuarial forecasts. For women, on the other hand, we find that official forecasts are higher than the 95% confidence band for all ages between 30 and 70. Hence, women indeed expect to die significantly younger than the actuarial estimates predict. The size of this difference is large: close to 5 years around the age of 60. Note, however, that these estimates do not take account of missing information due to the small number of elicited probabilities. Moreover, they assume that reported probabilities are not rounded.

Figure 5(b) plots 95% confidence bands for subjective remaining life expectancy for men (top) and women (bottom). These bands are based on bounds derived without smoothing expectations and span the width from the lower end of the 95% confidence interval for the lower bound on life expectancy to the higher end of the 95% confidence interval for the upper bound. As before, the solid lines are the corresponding predictions from life-tables. Even without rounding, the subjective data are consistent with the actuarial forecasts across all ages for both men and women. Moreover, allowing for rounding does not affect the average bounds much, though the effect is slightly larger at younger ages. Hence, we conclude that without additional assumptions we cannot reject the hypothesis that expectations are on average consistent with the life tables for both men and women in our sample. Robustness checks indicate that these conclusions remain largely unchanged when we lower the maximum lifespan to 100 years: for men the actuarial forecasts



Kernel regressions, bandwidth 3 years

Figure 5. Subjective versus actuarial remaining life expectancy.

remain well within the subjective bounds for all ages, while for women the actuarial forecasts remain within the 95% confidence bands if we allow for rounding. Estimates are available upon request.

Finally, we check whether rounding alone can close the gap between average expectations of women and the actuarial forecasts that is evident in panel (a). Panel (c) shows that common rounding closes the gap for all ages except for 50 to 70 year olds and around the age of 30, for which small differences remain. The more conservative general rounding scheme, on the other hand, removes those discrepancies. Under general rounding, we cannot reject the null hypothesis that the average upper bound of the intervals is equal to the corresponding life table forecast for any age. The same conclusions emerge if we approximate expectations by means of linear splines. Including prediction uncertainty for the official forecasts would reinforce these patterns.

5. CONCLUSION

When investigating subjective expectations regarding a continuous variable, researchers usually point identify expectations parametrically or by means of interpolation (e.g., Dominitz and Manski 1997; Dominitz 1998; Perozek 2008; Kutlu and Kalwij 2012; Bellemare, Bissonnette, and Kröger 2012). Building on the work of Manski (2003); Engelberg, Manski, and Williams (2009); and Manski and Molinari (2010), we analyze what can be learned about mortality expectations under weaker assumptions. In particular, the limited number of elicited probabilities and the possibility that reported probabilities may be rounded

imply that we can only partially identify distributional curves. We find that rounding is prevalent in the data, since only 39% of reported beliefs are consistent with an increasing hazard of death if probabilities are not rounded, while up to 92% are if we do allow for rounding. We construct identified sets for survival functions that neither assume a functional form for expectations nor interpolate between the elicited probabilities. We show that this procedure can easily be generalized to allow for rounding of the subjective probabilities reported in surveys and propose a refinement that narrows down the size of the identified sets by combining spline interpolation with rounding.

In our baseline scenario without refinements, the bounds on life expectancy are 12 years wide on average. This is too wide to be informative: models for interval-censored dependent variables fail to corroborate the associations between life expectancy and the covariates *age* and *health* that are highly statistically significant in noncensored models. If we smooth survival functions between observed points, the intervals can be narrowed to an average width of 3 years, which does allow for meaningful inference. Under the assumption that all probabilities reported by a given respondent are rounded to the same extent, as proposed by Manski and Molinari (2010), partially identified models show the same patterns that emerge from point identified models.

To evaluate the use of actuarial tables as proxies for average expectations we match subjective life expectancies to official life tables constructed by Statistics Netherlands. If we point identify expectations we confirm for the Netherlands the finding from Perozek (2008) that women, but not men, in the U.S. expect to live shorter on average than the life tables predict. However, using the baseline partially identified approach we

can no longer reject that expectations of women are consistent with the forecasts. This emerges even starker if we allow for rounding in the reported probabilities. If we interpolate expectations and simultaneously allow for common rounding, expectations of women are inconsistent with the life tables around the ages of 30 and 60 (the size of the remaining difference is much smaller). If we allow for the more conservative rounding scheme that allows probabilities from a given respondent to be rounded differently, we cannot reject consistency of women's expectations with actuarial forecasts for any age.

The general idea that emerges is that it is possible to learn about subjective expectations without imposing parametric restrictions on beliefs or even point identifying them. The methods we propose are sufficiently flexible to take into account rounding issues that are relevant for survey data of many types. Moreover, the extent of rounding can be quantified by checking whether reported beliefs are consistent with plausible assumptions. Our partial identification framework yields new insights into the influence of parametric assumptions on the analysis of this important and increasingly popular type of data.

APPENDIX: ALGORITHM FOR CONVEX INTEGRATED HAZARD

A given rounding scheme yields bounds for the true survival probability at each age threshold, which translate into bounds for the integrated hazard (see panel (a) of Figure 2). Denote by Λ_i^{\min} and Λ_i^{\max} the bounds on the integrated hazard at target age t_i , where $i = 0, 1, \dots, n$ ($i = 0$ for the current age of the respondent and $i = n$ for the final target age 90). Since a weakly increasing hazard of death is equivalent to a convex integrated hazard, we check whether there exists a convex piecewise linear function that passes between the bounds for the integrated hazard at all target ages. The algorithm, which is illustrated in panels (b) and (c) of Figure 2, consists of the following steps:

1. Start from $i = 0$, the current age of the respondent. Note that $\Lambda_0^{\min} = \Lambda_0^{\max} = 0$.
2. Define the maximum slope over the interval $[t_i, t_{i+1})$ as:

$$m_i = \min \left(\frac{\Lambda_{i+1}^{\max} - \Lambda_i^{\max}}{t_{i+1} - t_i}, \dots, \frac{\Lambda_n^{\max} - \Lambda_i^{\max}}{t_n - t_i} \right)$$

3. If needed, redefine the maximum integrated hazard at some target age t_j , $j > i$, as:

$$\tilde{\Lambda}_j^{\max} = \min (m_i (t_j - t_i) + \Lambda_i^{\max}, \Lambda_j^{\max})$$

If $\tilde{\Lambda}_j^{\max} < \Lambda_j^{\min}$, true beliefs are not consistent with a convex integrated hazard (and hence with an increasing hazard of death).

4. Increase i by 1 and start from Step 2 until reaching $i = n - 1$.

The resulting integrated hazard is shown in panel (d) of Figure 2. Using this algorithm, we check for each respondent whether underlying expectations may satisfy the increasing hazard assumption without rounding and under the general and common rounding schemes.

ACKNOWLEDGMENTS

We thank Netspar for financially supporting the data collection. Moreover, we thank Arthur van Soest, Frederic Vermeulen, Martin Salm, Marcel Das, Pierre-Carl Michaud, Rob Alessie, Charles Bellemare, participants at the 2014 MiSoC Workshop on Subjective Expectations and Probabilities in Economics and Psychology and two anonymous reviewers and the associate editor for their insightful comments. Bissonnette acknowledges financial support from the Industrielle Alliance Research Chair on the Economics of Demographic Change.

[Received September 2014. Accepted July 2016.]

REFERENCES

- Bellemare, C., Bissonnette, L., and Kröger, S. (2012), "Flexible Approximation of Subjective Expectations Using Probability Questions: An Application to the Investment Game," *Journal of Business and Economic Statistics*, 30, 125–131. [505,506,514]
- Beresteanu, A., and Molinari, F. (2008), "Asymptotic Properties for a Class of Partially Identified Models," *Econometrica*, 76, 763–814. [506]
- Beresteanu, A., Molinari, F., and Steeg Morris, D. (2010), "Asymptotics for Partially Identified Models in STATA," Stata Code for BLP on Interval Data, available at <http://www.pitt.edu/~arie/code.htm>. [512]
- Bissonnette, L., Hurd, M. D., and Michaud, P.-C. (2014), "Individual Survival Curves Combining Subjective and Actual Mortality Risks," IZA Discussion Paper Series, No. 8658. [505]
- Dominitz, J. (1998), "Earnings Expectations, Revisions and Realizations," *Review of Economics and Statistics*, 80, 374–388. [505,514]
- Dominitz, J., and Manski, C. F. (1997), "Using Expectations Data to Study Subjective Income Expectations," *Journal of the American Statistical Association*, 92, 855–867. [505,506,514]
- Dominitz, J., and Manski, C. F. (2006), "Measuring Pension Benefit Expectations Probabilistically," *Labour*, 20, 201–236. [505]
- Engelberg, J., Manski, C. F., and Williams, J. (2009), "Comparing the Point Predictions and Subjective Probability Distributions of Professional Forecasters," *Journal of Economic and Business Statistics*, 27, 30–41. [506,514]
- Hurd, M. (2009), "Subjective Probabilities in Household Surveys," *Annual Economic Review*, 1, 543–562. [505]
- Hurd, M., and McGarry, K. (1995), "Evaluation of the Subjective Probabilities of Survival in the Health and Retirement Study," *The Journal of Human Resources*, 30(Special Issue on the Health and Retirement Study: Data Quality and Early Results), S268–S292. [505]
- (2002), "The Predictive Validity of Subjective Probabilities of Survival," *The Economic Journal*, 112, 966–985. [505]
- Imbens, G., and Manski, C. F. (2004), "Confidence Intervals for Partially Identified Parameters," *Econometrica*, 72, 1845–1857. [506,512]
- Kutlu, V., and Kalwij, A. S. (2012), "Individuals' Survival Expectations and Actual Mortality," Utrecht University Working Paper. [505,511,512,514]
- Manski, C. F. (2003), *Partial Identification of Probability Distributions*, New York: Springer. [505,514]
- (2004), "Measuring Expectations," *Econometrica*, 75, 1329–1376. [505]
- Manski, C. F., and Molinari, F. (2010), "Rounding Probabilistic Expectations in Surveys," *Journal of Business and Economic Statistics*, 28, 219–231. [505,507,509,514]
- Peracchi, F., and Perotti, V. (2011), "Subjective Survival Probabilities and Life Tables: Evidence from Europe," EIF working paper 16/10. [505]
- Perozek, M. (2008), "Using Subjective Expectations to Forecast Longevity: Do Survey Respondents Know Something We Don't Know?" *Demography*, 45, 95–113. [505,506,511,514]
- Smith, V. K., Taylor, D. H. Jr, and Sloan, F. A. (2001), "Longevity Expectations and Death: Can People Predict Their Own Demise?" *American Economic Review*, 91, 1126–1134. [505]