Volume Title: The Economics of Aging

Volume Author/Editor: David A. Wise, editor

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-90295-1

Volume URL: http://www.nber.org/books/wise89-1

Conference Date: March 19-22, 1987

Publication Date: 1989

Chapter Title: The Timing of Retirement: A Comparison of Expectations and Realizations

Chapter Author: B. Douglas Bernheim

Chapter URL: http://www.nber.org/chapters/c11587

Chapter pages in book: (p. 335 - 358)

# 11 The Timing of Retirement: A Comparison of Expectations and Realizations

B. Douglas Bernheim

## 11.1 Introduction

Modern life-cycle theory is based upon the premise that consumers think seriously and coherently about the relatively distant and uncertain future. While the empirical validity of this premise is controversial, existing evidence is either highly indirect or anecdotal. To resolve this controversy, it is necessary to conduct direct comparisons of consumers' plans and expectations with eventual realizations.

Previous empirical work on household expectations has focused primarily on short-run inflation (see Huizinga 1980, Curtin 1982, Gramlich 1983, and Papadia 1982; Aiginger 1979 considers a somewhat broader range of variables). Accordingly, these studies shed very little light on the issue of whether consumers form accurate expectations and successful economic plans over relatively long time horizons.

In a separate paper (Bernheim 1988), I have studied the accuracy of pre-retirement expectations concerning Social Security benefits. I found that while survey responses to questions about expected benefits contain a relatively high level of noise, there is nevertheless strong evidence that consumers do think seriously about future events. While consumers do not form expectations on the basis of all available information, they do appear to be reasonably competent at making relatively accurate forecasts conditional upon the information that they do use. Indeed, the data broadly suggest that consumers correctly

B. Douglas Bernheim is the Harold J. Hines Professor of Risk Management at the J. L. Kellogg graduate School of Management, Northwestern University, and a Research Associate of the National Bureau of Economic Research.

The author would like to thank the National Institute on Aging for financial support. Laurence Levin provided invaluable research assistance. The opinions expressed here are those of the author and should not be attributed to any other individual or organization.

anticipated the general effects of legislative action during the early 1970s, contrary to the supposition of most previous authors (see, e.g., Hurd and Boskin 1981; Anderson, Burkhauser, and Quinn 1986; and Burtless 1986).

In the current paper, I employ data drawn from the Social Security Administration's Retirement History Survey (RHS) to study the accuracy of expectations concerning the timing of retirement. This is an important complement to my earlier work, in that Social Security benefits are largely determined by exogenous events, while retirement is a fundamental decision variable. Accordingly, the emphasis here is on the accuracy of economic *plans*, rather than pure expectations. While many authors have previously studied determinants of the retirement decision (see the surveys by Hurd 1983, and Mitchell and Fields 1982), all have simply assumed that workers make systematic and viable retirement plans. There is no previous test of this proposition.

The major findings of this study are as follows.

1. Survey responses to questions about expected dates of retirement reflect modes, rather than means. That is, consumers report the most likely date of retirement, rather than the mean date, given subjective probabilities. This distinction is extremely important, since the distribution of actual retirement dates for a fixed expectation is highly skewed. Unlike the case of Social Security benefits, the evidence does not support the view that consumers report noisy measures of expectations.

2. Most individuals are reasonably competent at forming relatively accurate expectations about the timing of retirement. Alternatively, consumers apparently form serious economic plans and ordinarily stick to them. Perhaps surprisingly, there is once again very little evidence to support the view that expectations were abnormally inaccurate during periods in which Social Security benefits enjoyed significant statutory increases.

3. The accuracy of expectations differs systematically by population subgroup. In contrast to my findings for Social Security benefits, I find that men form more accurate retirement expectations than women, although single women do appear to gain relative accuracy as retirement approaches. Married women are particularly prone to discover that they must work longer than expected. Comparatively wealthy individuals tend to make somewhat more accurate forecasts, but education does not improve predictive skill. Some evidence also suggests that workers with mandatory retirement dates typically retire much sooner than expected, perhaps because they suppose erroneously that alternative employment will be easy to find.

Work by Anderson, Burkhauser, and Quinn (1986) has also employed the RHS data on retirement expectations. Their object was to use this data in an analysis of behavior rather than to identify properties of reported expectations, as in the current paper. My findings are at odds with the implicit assumptions upon which these authors based their behavioral analysis, and therefore they call their results into question. Hall and Johnson (1980) have also studied retirement expectations, but their object was to model the formation of plans rather than to compare these plans with later realizations.

This paper is organized as follows. Section 11.2 discusses some alternative hypotheses about the nature of reported expectations concerning the timing of retirement. I describe the data in section 11.3. Section 11.4 tests the view that individuals report mean realizations given probabilistic beliefs. In section 11.5, I consider the hypothesis that respondents report modal beliefs, i.e., most likely dates of retirement. Finding the evidence favorable, I proceed to a comparison of various population subgroups in different survey years.

## 11.2 The Alternative Hypotheses

When an individual is asked to report his expected date of retirement, what does his answer represent? Survey questions about expectations are unfortunately ambiguous and admit several possible interpretations. Yet if we are to make valid use of these data in any behavioral context, it is essential to resolve this issue.

One possibility is that the typical individual reports the mean of some subjective probability distribution. It is useful to set up this hypothesis formally. Let $R$ be the individual's actual date of retirement. At some time $t$, he has access to information, $I(t)$, which he uses to form subjective beliefs about the timing of retirement. Let $p[r|I(t)]$ denote the subjective probability that the individual will retire at date $r$, given available information at time $t$, and let $ER(t)$ denote his reported expectation at time $t$. The first hypothesis is that

$$ER(t) = E[R \mid I(t)], \tag{1}$$

where $E[.]$ indicates a mathematical expectation based on the probability distribution $p[.]$.

Unless we place some additional restriction on the subjective probability distribution, this proposition is not testable. My strategy is to test it jointly with the hypothesis of rational expectations. Specifically, if one assumes that the subjective probability distribution $p[.]$ coincides with the objective distribution, then equation (1) suggests a regression of $R$ on an intercept, $ER(t)$, and $I(t)$. Under the joint hypotheses, the intercept and coefficients of $I(t)$ should be 0, while the coefficient of $ER(t)$ should be 1. It is, of course, essential that one only include informational variables that the individual actually used in forming his expectations. Since this is difficult to establish a priori, it is advisable to conduct a weak form of this test by omitting the $I(t)$ entirely.

My study of expected Social Security benefits provided strong support for the analog of this first hypothesis, and one might therefore expect to find the data supportive here as well. Yet it is essential to understand that retirement is a very different kind of event than the realization of Social Security benefits. Many workers form extremely specific retirement plans, which they intend to follow barring unforeseen circumstances. In contrast, workers may have only "ballpark" notions about their Social Security benefits. Accordingly, it is easily conceivable that, when asked about their expectations, individuals report means for Social Security benefits but report the most likely outcome for date of retirement.

This discussion leads naturally to the second hypothesis, which is that measured expected dates of retirement reflect modes of subjective distributions. Formally,

$$ER(t) = \text{argmax}_r \; p[r \mid I(t)]. \tag{2}$$

Once again, this proposition is not testable in the absence of further restrictions on the subjective distribution. As before, my strategy is to test it jointly with the hypothesis of rational expectations. Assuming that $p[.]$ coincides with the true objective distribution, one can compare measured expectations with modal realizations.

While these two hypotheses certainly do not exhaust all conceivable alternatives (e.g., perhaps individuals report medians, or pure noise),[1] I take them to be the most interesting possibilities.

## 11.3 Data

The data for this study are drawn from the Social Security Administration's Retirement History Survey (RHS), which followed a sample of retirement-aged households (58 to 63 years old in 1969) for a period of 10 years, beginning in 1969. Each household was surveyed once every two years (1969, 1971, 1973, 1975, and 1979). Although the initial wave included more than 11,000 households, there was substantial attrition over successive waves.

Each working respondent reported his or her expected age of retirement in every survey year, with the exception of married women, who were not asked this question in 1973. Using the respondent's age, I transformed this variable into ERET, the expected date of retirement. Survey responses to questions about expected retirement were extremely sparse in 1977 and 1979 (primarily because most of the sample had already retired by 1977); I therefore focus on expectations reported in the first four survey waves.

The primary advantage of the RHS is that it allows the analyst to identify realizations by employing data from subsequent survey waves. While the identification of a date of retirement is usually problematic, here it poses few difficulties. In the current context, it is not necessary or even desirable to obtain a conceptually "correct" measure of retirement. When an individual reports an expected date of retirement, he may well have in mind some idiosyncratic notion of what retirement means. However, unless he changes his notion over time, one can assume that self-reported retirement refers to the same potentially idiosyncratic event. Accordingly, I use self-reported retirement to construct RET, my measure of the eventual realization.

Unfortunately, data on self-reported retirement are somewhat incomplete. Although individuals do report whether or not they consider themselves retired in each survey year, they are not asked to indicate exactly when retirement took place. This creates a problem in that surveys were administered in alternate years. In practice, I calculate RET as follows. First, I identify the first survey year in which the respondent reported himself to be retired. Second, for this same survey year, I determine the date at which the respondent left his last job. If this falls within the previous two years, I take it to be his date of retirement. If it does not (typically because of missing information), I determine the date at which the respondent began to receive Social Security benefits. If this falls within the previous two years, I take it to be his date of retirement. If it does not, I simply assume that he retired midway between the surveys.

In conjunction with testing the first set of joint hypotheses, I relate forecast errors to available information in order to identify the kinds of information that individuals either ignore or process incorrectly. I consider a large number of informational variables, which I group into three categories.

The first category contains variables which measure other reported expectations. The inclusion of these variables allows me to determine whether or not individuals have internally consistent expectations, in the sense that they base all expectations on the same set of information. Definitions of specific variables follow:

ESS: expected Social Security benefits
EOI: expected retirement income, other than Social Security

Data on expectations are, of course, incomplete—many individuals who report an expected date of retirement do not, for example, report expected Social Security benefits. Accordingly, I also use dummy variables, which equal 1 if the individual reports the associated expectation, and 0 otherwise. I refer to the dummies corresponding to the two variables listed above as DSS and DOI.

The second category includes a single variable, which is the individual's current Social Security entitlement, CSS, defined as the level of benefits he would receive under current law if he retired immediately. CSS is, theoretically, part of each individual's information set, in that it depends only upon his own past earnings history and upon current law (which is public information). My previous study of Social Security benefits suggested that individuals fail to use much of the information contained in CSS; since it is natural to suppose that workers adjust retirement plans upon learning more about Social Security entitlements, this information could be correlated with the forecast error for date of retirement as well.

The third and final category includes various demographic variables and other household characteristics which might be useful in predicting retirement. The list of variables includes:

| | |
|---|---|
| MAR: | a dummy variable, indicating whether or not the respondent is married (1 = married, 0 = other) |
| DIV: | a dummy variable, indicating whether or not the respondent is divorced (1 = divorced, 0 = other) |
| WID: | a dummy variable, indicating whether or not the respondent is a widow or widower (1 = widow or widower, 0 = other) |
| AGE: | the respondent's age |
| SAGE: | the respondent's spouse's age |
| ED: | the respondent's level of educational attainment (measured in number of years) |
| SED: | the respondent's spouse's level of educational attainment |
| W: | the household's net wealth (including financial assets, businesses, and real property) |
| GH: | a dummy variable, indicating whether or not the respondent reports his health as being better than average for his age (1 = better, 0 = other) |
| BH: | a dummy variable, indicating whether or not the respondent reports his health as being worse than average for his age (1 = worse, 0 = other) |
| KIDS: | number of children |
| COMPRET: | a dummy variable, indicating whether or not the respondent's employer maintains a compulsory retirement age (1 = yes, 0 = no) |
| MOVE: | a dummy variable, indicating whether or not the respondent has moved within the past two years. |

Before passing on to analysis of the data, it is important to discuss two potential problems. The first concerns sample selection biases. I drop observations from the analysis for three reasons: 1) the respondent fails to report an expected date of retirement, 2) the reported date is obviously nonsensical (e.g., it precedes the date at which it was reported), or 3) the household disappeared from the RHS prior to retirement. Note that the first two items both reflect household characteristics that are known when the respondent makes his forecast. Since the forecasts are then presumably conditioned on this information, no sample selection biases arise. The third item (subsequent attrition) is potentially problematic. I return to this issue in section 11.4, where I propose and implement a statistical correction.

The second problem concerns the non-independence of realizations. In a short panel such as the RHS, forecast errors are probably correlated across observations due to "macro" events. Since the 1970s witnessed several large and potentially unexpected real increases in Social Security benefits, this problem is potentially severe. In particular, real Social Security benefits increased by 4.2 percent in January 1970, 4.8 percent in January 1971, and 14.1 percent in September 1972. In addition, benefits were "double indexed" for inflation from 1975 to 1977. If, as suggested by many analysts, unanticipated increases in Social Security benefits caused many workers to retire unexpectedly early, then we might well find that expectations were systematically off during this period. On the other hand, the major benefit increases were primarily concentrated in a few years (especially 1972). It should be possible to shed some light on the question of whether these changes were indeed unanticipated by looking for evidence of systematic forecast errors at those points in time.

## 11.4 The Mean Value Hypothesis

I begin formal analysis of the data by comparing expectations to mean realization in order to test the first hypothesis discussed in section 11.2. Table 11.1 contains some highly revealing summary statistics for married men. For each survey year, I have grouped observations by common values of ERET. For each group, I report four things: the difference between the average date of actual retirement and ERET, the standard deviation of the retirement date, the mean squared forecast error, and the number of observations.

The most striking feature of table 11.1 is that there is very little relationship between ERET and the average date of retirement. To be sure, those with higher values of ERET tend to retire later, on average. However, the mean date of retirement coincides with ERET in few if any cells. Indeed, in 19 out of 20 cells one can reject the hypothesis

Table 11.1 **Expectations and Mean Realizations for Married Men**

| | Survey Year | | | |
|---|---|---|---|---|
| ERET | 1969 | 1971 | 1973 | 1975 |
| 1969 | 1.9 | — | — | — |
| | 2.0 | | | |
| | 7.6 | | | |
| | 157 | | | |
| 1970 | 1.3 | — | — | — |
| | 1.9 | | | |
| | 5.3 | | | |
| | 311 | | | |
| 1971 | 0.9 | 1.2 | | |
| | 1.9 | 1.6 | | |
| | 4.5 | 4.1 | | |
| | 411 | 281 | | |
| 1972 | 0.7 | 0.9 | | |
| | 2.1 | 1.6 | | |
| | 5.0 | 3.5 | | |
| | 375 | 367 | | |
| 1973 | −0.1 | 0.2 | 1.0 | — |
| | 2.0 | 1.5 | 1.3 | |
| | 4.0 | 2.4 | 2.6 | |
| | 290 | 309 | 198 | |
| 1974 | −0.5 | −0.2 | 0.6 | |
| | 2.1 | 1.6 | 1.4 | |
| | 4.9 | 2.8 | 2.2 | |
| | 240 | 241 | 225 | |
| 1975 | −1.1 | −0.5 | 0.3 | 0.7 |
| | 2.1 | 1.8 | 1.4 | 0.9 |
| | 5.7 | 3.3 | 1.9 | 1.3 |
| | 263 | 255 | 253 | 135 |
| 1976 | −2.0 | −1.7 | −0.3 | 0.3 |
| | 2.3 | 1.8 | 1.3 | 0.9 |
| | 9.4 | 5.9 | 1.7 | 0.9 |
| | 112 | 93 | 78 | 76 |

that the mean date equals ERET with at least 95 percent confidence.[2] Roughly speaking, it appears that a one-year change in the expected date is associated with slightly less than a one-half year change in the average realized date. The implications of equation (1) are strongly contradicted.

Other aspects of table 11.1 are also puzzling. The standard deviation of RET does not appear to be higher for groups that intend to retire in the more distant future, despite the fact that information should improve as retirement grows more imminent. Similarly, mean squared forecast errors do not rise monotonically with ERET. Yet standard errors and mean squared forecast errors both fall monotonically between successive survey years. The mean value hypothesis provides no clue as to the source of this trend.

As remarked in section 11.3, these calculations suffer from potential sample selection biases. Specifically, I have dropped from my sample all individuals who leave the survey before retiring. Unless attrition is associated with earlier-than-normal retirement, the (objective) expected date of retirement for such individuals, conditional upon ERET and observed behavior, exceeds the expectation based upon ERET alone. Accordingly, the omission of these observations probably biases the estimated mean retirement date downward.

To correct for this problem, one must know something about the retirement behavior of individuals after they leave the sample. By definition, this is unobservable. Consequently, it is necessary to maintain an ancillary hypothesis. In order to make some illustrative calculations, I assume that attrition is not systematically related to subsequent retirement.[3] This assumption allows me to correct for sample selection as follows. For each subsample (characterized by survey year and ERET), I calculate hazard rates for retirement in each year, i.e., the number of individuals retiring in that year divided by the total number of individuals remaining from the original subsample (including those who subsequently left the sample before retiring). Under my maintained hypothesis, this yields a consistent estimate of the true population hazard rate. From these rates, one can then reconstruct the true distribution of retirement dates.

In practice, relatively few individuals who met my other selection criteria actually left the sample before retiring. As a result, the impact of this correction was extremely small. For most cells, the mean of the corrected distribution exceeded the uncorrected mean by 0.1 year; in a few cases the difference was 0.2 years, and in a few others it was virtually zero. The corrected distributions strongly resembled the uncorrected distributions, and indeed the modes did not differ in any cell. Thus, I conclude that the sample selection bias is of little consequence. Furthermore, I suspect that the correction used here overstates the bias, in that attrition is probably correlated with earlier-than-normal retirement.

In light of the results in table 11.1, it should hardly be surprising that a regression of RET on ERET produces extremely negative results. Coefficient estimates appear in equation 1 of table 11.2. These results are based on expectations reported in 1971 but are representative of other years as well. I have chosen to report results for 1971 only because the data for that year are somewhat superior (in 1969, the ESS variable, used below, is flawed; in 1973, ERET is not available for married women; in 1975, the total data sample is much smaller). Note that the intercept is non-zero and dwarfs its standard error. The coefficient of ERET is

Table 11.2 Regression Results for 1971

| | Equation Number | | | |
|---|---|---|---|---|
| | 1 | 2 | 3 | 4 |
| Technique | OLS | IV | IV | IV |
| Instruments | None | Set #1 | Set #2 | Set #3 |
| Intercept | 56.5 | 45.2 | 20.7 | 37.1 |
| | (1.3) | (13.1) | (6.22) | (3.44) |
| ERET | 0.234 | 0.374 | 0.722 | 0.499 |
| | (0.018) | (0.179) | (0.085) | (0.047) |
| $R^2$ | 0.080 | 0.002 | 0.036 | 0.055 |

far below unity and is estimated very precisely. Formally, this signals a resounding rejection of the null hypothesis.

Yet one should not be too hasty in discarding the mean value hypothesis. I obtained similar negative findings in my analysis of expectations concerning Social Security benefits, but noted that these could be attributable to "noisy" measurement of the expectations variable. Formal analysis bore this conjecture out. It is therefore advisable to investigate the same possibility in the current context.

The classical remedy for measurement error is instrumental variables. In the current context, a variable is a valid instrument if it belongs to the information set on which the individual based his expectation. Unfortunately, the identity of this set is known only to the individual. Accordingly, one must maintain the hypothesis that individuals do use certain kinds of information in order to conduct the test.

The evidence in my previous study supported the view that individuals use the same information to form all of their expectations. This suggests that other expectations (ESS, EOI) are valid instruments. Of course, these variables may also be measured with error, but this is of no consequence as long as the measurement errors are uncorrelated. Equation 2 in table 11.2 provides estimated coefficients, where the expectational variables have been used as instruments. While the estimates are somewhat less precise than those obtained through OLS, the overall picture is unchanged.

For completeness, I have included two additional regressions, using the other two sets of informational variables as instruments. One can think of these regressions as reflecting alternative hypotheses about the kinds of information that workers actually use when constructing their forecasts. The results are uniformly negative. I obtain the most favorable estimates by using CSS as an instrument (equation 3). However, my previous study clearly established that individuals do not make use of all the information contained in CSS; it is therefore an unsuitable instrument.

These results contrast with my findings for expectations about Social Security benefits. The statistical failure of the mean value hypothesis cannot in this case be traced to the presence of measurement error. Upon reflection, this is hardly surprising. Since individuals probably do not have very precise notions about their future Social Security benefits, it stands to reason that they will report "ballpark" figures. However, it seems likely that most workers form very specific plans about the timing of retirement, particularly as it becomes more imminent. It is difficult to understand why an individual would report that he intends to retire at age 63, if in fact he plans to do so at age 65.

It is, of course, possible that the negative results in table 11.2 all stem from a failure to identify appropriate instruments. I therefore present one final set of estimates in table 11.3. Here, I have regressed the forecast error (RET − ERET) on the full complement of informational variables. This procedure yields consistent estimates even if ERET is measured with error (unfortunately, it precludes us from testing the theory by examining the coefficient of ERET). If the mean value hypothesis is correct, then one can determine the kinds of information that individuals either ignore or use improperly by examining the coefficient estimates. Note first that the coefficients of the expectational variables are not significantly different from zero. This finding validates

**Table 11.3 Forecast Error Regression, 1971**

| Variable | Coefficient | Variable | Coefficient |
|---|---|---|---|
| Intercept | −11.3 | $ED/10^3$ | −0.59 |
| | (2.7) | | (7.02) |
| $ESS/10^5$ | 6.02 | $SPED/10^3$ | −5.37 |
| | (8.70) | | (8.47) |
| DSS | −0.197 | $W/10^7$ | 9.06 |
| | (0.215) | | (7.39) |
| $EOI/10^5$ | 1.60 | GH | −0.219 |
| | (6.59) | | (0.112) |
| DOI | −0.121 | PH | 0.082 |
| | (0.149) | | (0.174) |
| $CSS/10^5$ | −1.60 | $KIDS/10^2$ | 0.06 |
| | (6.59) | | (2.65) |
| AGE | 0.195 | COMPRET | −0.847 |
| | (0.044) | | (0.125) |
| $SPAGE/10^3$ | −6.04 | MOVE | 0.301 |
| | (7.62) | | (0.182) |
| MAR | 0.653 | $R^2$ | 0.051 |
| | (0.484) | | |
| DIV | −0.147 | Observations | 1919 |
| | (0.303) | | |
| WID | 0.306 | | |
| | (0.242) | | |

the use of these variables as instruments, and strengthens the conclusion that my negative results are not attributable to measurement error. Variables appearing with statistically significant coefficients include AGE, GH, and COMPRET. The last of these is particularly interesting, since it suggests that workers at jobs with mandatory retirement ages tend to believe that they will be able to continue working longer than they actually can. However, I caution that this conclusion is based upon a suspect empirical specification, in that my findings are generally unfavorable to the mean value hypothesis.

## 11.5 The Modal Value Hypothesis

I now turn to the possibility that respondents report their most likely dates of retirement, rather than mean dates. To investigate this hypothesis, I group observations by common values of ERET for each survey year and compute the modal realization for each group. Table 11.4 presents results for married men. This table contains 20 cells, identified by the survey year and value of ERET. In each cell, I report (in order) the modal value of RET minus ERET, the fraction of the group for which RET and ERET coincide, the fraction of the group for which RET is within one year of ERET, and the total number of observations.

The most striking aspect of table 11.4 is that the modal realization coincides with ERET in 16 out of 20 cells. In the four remaining cases, the mode differs from ERET by only a single year, and ERET is the second most common outcome, lagging the mode by a relatively small margin. Since ERET exceeds the mode in exactly half (two) of these cases, there is no indication of systematic bias.

One can also obtain some feeling for the accuracy of reported expectations by examining the second and third entries in each cell. I caution against placing too much emphasis on the fraction of respondents for whom RET and ERET coincide exactly. An individual who is 62 years old in 1969 and who reports that he intends to retire when 65 could plan to leave his job in either 1971, 1972, or 1973, depending upon his exact date of birth. Since it is impossible to identify the month during which an individual retires, I cannot adjust for this ambiguity. Accordingly, it is more appropriate to examine the fraction of individuals for which RET differs from ERET by at most one year. Note that as long as individuals do not intend to retire too far in the future, expectations are highly accurate; in all 16 cells for which ERET exceeds the survey year by four years or less, more than 60 percent of the respondents retired within one year of ERET.

As an individual approaches retirement, he presumably forms his expectation on the basis of more complete information. We would

**Table 11.4** **Expectations and Model Realizations for Married Men**

| | Survey Year | | | |
|---|---|---|---|---|
| ERET | 1969 | 1971 | 1973 | 1975 |
| 1969 | 1 | — | — | — |
| | 0.26 | | | |
| | 0.61 | | | |
| | 157 | | | |
| 1970 | 0 | — | — | — |
| | 0.39 | | | |
| | 0.65 | | | |
| | 311 | | | |
| 1971 | 0 | 0 | — | — |
| | 0.28 | 0.43 | | |
| | 0.67 | 0.75 | | |
| | 411 | 281 | | |
| 1972 | 0 | 0 | — | — |
| | 0.29 | 0.44 | | |
| | 0.57 | 0.74 | | |
| | 375 | 367 | | |
| 1973 | −1 | 0 | 0 | — |
| | 0.22 | 0.32 | 0.44 | |
| | 0.60 | 0.79 | 0.79 | |
| | 290 | 309 | 198 | |
| 1974 | 0 | 0 | 0 | — |
| | 0.26 | 0.32 | 0.47 | |
| | 0.50 | 0.60 | 0.80 | |
| | 240 | 241 | 225 | |
| 1975 | 1 | 0 | −1 | 0 |
| | 0.18 | 0.24 | 0.29 | 0.47 |
| | 0.53 | 0.64 | 0.80 | 0.85 |
| | 263 | 255 | 253 | 135 |
| 1976 | 0 | 0 | 0 | 0 |
| | 0.22 | 0.23 | 0.39 | 0.62 |
| | 0.38 | 0.42 | 0.66 | 0.84 |
| | 112 | 93 | 78 | 76 |

therefore expect the accuracy of his forecast to improve. It is possible to examine this prediction in two different ways. First, one can investigate the relationship between ERET and accuracy during any survey year by reading down columns. While accuracy does not decline monotonically with the expected date of retirement, there is a general tendency for it to fall. Second, one can examine the relationship between accuracy and the survey date for any given value of ERET by reading across rows. Note that in 23 of 24 possible pairwise comparisons (12 for fractions with RET = ERET, 12 for fractions with RET within one year of ERET), accuracy improves when the question is posed at a later date. In the one remaining case, it is simply unchanged. This finding provides striking confirmation for the view that information improves as individuals approach retirement.

An additional feature of table 11.4 merits comment. Let $T$ denote the survey year. Fix $t$, and consider individuals who expect to retire in year $T + t$. There is a strong tendency for the accuracy of expectations to rise with $T$ (to see this, read table 11.4 diagonally). The reason for this phenomenon is not immediately obvious. At first, one might suppose that, given $t$ (expected length of time until the event of interest), the date of reporting should not affect accuracy. However, one must bear in mind that average age is greater in later survey years. This causes significant compression of the retirement distribution, which leads in turn to greater accuracy. This observation underscores an important point: one should not assume that the shape of the conditional distribution is invariant with respect to either ERET or age. I will return to this point shortly.

Table 11.4 also sheds some light on the question of whether unanticipated changes in Social Security benefits during the early 1970s caused many workers to retire earlier than expected. Recall that by far the largest real benefit increase took place in 1972. If this change induced substantial early retirement, we would expect to see abnormal deviations from retirement plans during this period. There is little evidence of this in table 11.4. A substantial number of respondents in both 1969 and 1971 reported that they expected to retire after 1972. In 6 of the 8 relevant cells, the modal expectation still coincides with ERET. For those reporting ERET = 1975 in 1969, the modal realization was actually *after* 1975, not before. Only for those reporting ERET = 1973 in 1969 was the modal realization less than ERET, and indeed in this case 1972 was the most frequent date of retirement. Note, however, that 1969 forecasts for those with ERET = 1973 are only slightly less accurate than 1971 forecasts for those with ERET = 1975 (also 4 years in the future). Note also that 1971 forecasts for those with ERET = 1973 are actually more accurate than either 1973 forecasts for those with ERET = 1975, or 1969 forecasts for those with ERET = 1971 (both also 2 years in the future). Together, these observations suggest that changes in benefit levels did not induce substantial early retirement for individuals who had expected to stop working in 1973.

The substantial divergence of means and modes (tables 11.1 and 11.4) suggests that the conditional distributions of retirement dates may be highly skewed. This supposition is in fact correct. Figures 11.1 and 11.2 illustrate the distribution of retirement dates by ERET for 1969. One can see that when ERET is low, the conditional distribution is skewed to the right; as ERET rises, the skew shifts to the left. If reported expectations represent modes rather than means, this pattern is natural. Those expecting to retire very soon will, if surprised, generally retire later, and those expecting to retire late will, if surprised, generally retire sooner. This explains why the mean moves so much less than the mode, as noted in tables 11.1 and 11.4.

Failure to recognize this pattern can easily lead to misinterpretation of the data. Consider for example the study by Anderson, Burkhauser, and Quinn (1986). These authors examined the relationship between unexpected deviations from retirement plans and unexpected changes in Social Security benefits. They calculated the latter variable by comparing actual benefits available in the year of expected retirement to the level of benefits that would have been available had the 1969 statutes been adjusted for cost of living only. Through multinomial logit analysis, they found that respondents who experienced larger unexpected increases in Social Security benefits were much more prone to retire earlier than planned. Yet it now seems likely that this finding is merely an artifact of the data. Note that the authors' measure of unexpected benefit increases is primarily determined by ERET; the later the respondent expects to retire, the more the 1969 legislation will understate benefits available in the year of expected retirement. Furthermore, the pattern of skewness implies that higher values of ERET are associated with a greater frequency of unexpected early retirement. Combining these two observations leads one to expect a strong positive association between unexpected benefits and early retirement, even in the absence of a behavioral response. It is therefore conceivable that the finding is entirely spurious.

In fact, figures 11.1 and 11.2 provide only a very slight indication that the 1972 benefit changes may have induced some early retirement. In particular, the distributions for ERET = 1969, 1973, and 1975 exhibit somewhat higher frequencies for 1972 than one might ordinarily expect. However, the pattern is certainly far from overwhelming.

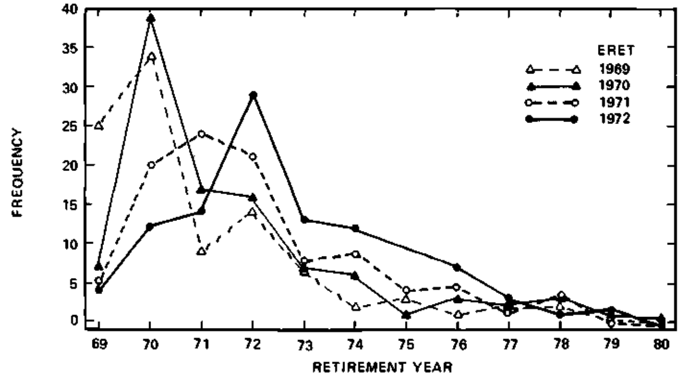


**Fig. 11.1** Distribution of RET by ERET, 1969, part 1

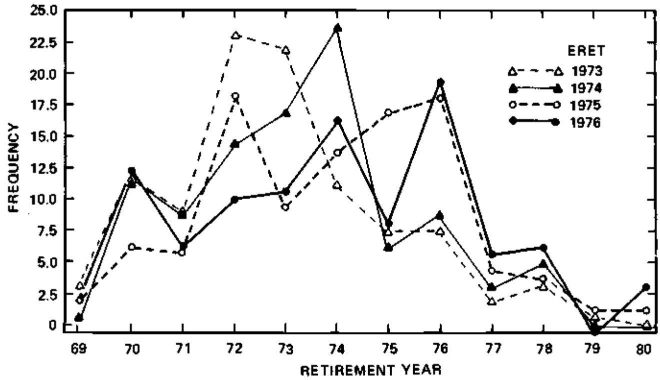


**Fig. 11.2** Distribution of RET by ERET, 1969, part 2

As a final step, I provide a comparison of expectations and realizations for various population subgroups, including married men, married women, single men, single women, widowers, widows, married men with high wealth, married men with low wealth, married men with high levels of educational attainment, and married men with low levels of educational attainment. I present results in tables 11.5 through 11.8, which correspond to each of four different survey years (1969 through 1975). Several consistent patterns emerge. First, married women form the least accurate expectations and are most likely to work longer than planned. Lower accuracy results in part from the fact that women tend to be younger and therefore further from retirement than their husbands. However, even if one compensates for this by, for example, comparing married men in 1969 to married women in 1975, the pattern is still evident. Second, there is a general tendency for single individuals, widows, and widowers to retire earlier than expected more frequently than married individuals. Third, in early survey waves the expectations of single women and widows were much less accurate than those of married men. However, in later waves this gap narrowed, and indeed the expectations of single women became more accurate than those of married men. Fourth, education appears to be inversely related to accuracy. Wealth is positively related to accuracy in early survey waves, but negatively related in later waves.

Overall, the evidence presented in this section is strongly consistent with the joint hypotheses that 1) when asked to report an expected date of retirement, an individual will describe the outcome that he or she considers most likely, and 2) the subjective distribution of retirement dates coincides with the objective distribution. Since this distribution is highly skewed, and since the skewness is related to the expected date of retirement, one cannot interpret the data as reflecting mean retirement dates. Finally, there is little or no evidence to support the

**Table 11.5 Expectations and Realizations by Subgroup, 1969**

| Subgroup | Number of Observations | Mean of ERET | Fraction with RET = ERET | Fraction with RET = ERET ± 1 | Fraction with RET < ERET − 1 | Fraction with RET > ERET + 1 |
|---|---|---|---|---|---|---|
| Married men | 2,240 | 72.5 | 0.259 | 0.570 | 0.232 | 0.198 |
| Married women | 482 | 74.8 | 0.154 | 0.361 | 0.247 | 0.392 |
| Single men | 73 | 71.7 | 0.301 | 0.562 | 0.315 | 0.123 |
| Single women | 114 | 72.0 | 0.246 | 0.526 | 0.281 | 0.193 |
| Widowers | 77 | 72.2 | 0.234 | 0.610 | 0.182 | 0.208 |
| Widows | 272 | 72.1 | 0.250 | 0.533 | 0.261 | 0.206 |
| High wealth: married men | 1,383 | 72.4 | 0.267 | 0.587 | 0.244 | 0.168 |
| Low weatlh: married men | 857 | 72.7 | 0.247 | 0.543 | 0.212 | 0.245 |
| High education: married men | 1,002 | 72.7 | 0.237 | 0.566 | 0.235 | 0.200 |
| Low education: married men | 1,238 | 72.3 | 0.278 | 0.574 | 0.230 | 0.196 |

**Table 11.6** **Expectations and Realizations by Subgroup, 1971**

| Subgroup | Number of Observations | Mean of ERET | Fraction with RET = ERET | Fraction with RET = ERET ± 1 | Fraction with RET < ERET − 1 | Fraction with RET > ERET + 1 |
|---|---|---|---|---|---|---|
| Married men | 1,619 | 73.3 | 0.334 | 0.672 | 0.170 | 0.158 |
| Married women | 639 | 75.4 | 0.178 | 0.426 | 0.224 | 0.351 |
| Single men | 55 | 73.1 | 0.400 | 0.673 | 0.218 | 0.109 |
| Single women | 91 | 73.5 | 0.396 | 0.747 | 0.121 | 0.132 |
| Widowers | 86 | 73.1 | 0.256 | 0.640 | 0.209 | 0.151 |
| Widows | 230 | 73.3 | 0.291 | 0.604 | 0.222 | 0.174 |
| High wealth: married men | 1,071 | 73.3 | 0.347 | 0.676 | 0.169 | 0.155 |
| Low wealth: married men | 548 | 73.4 | 0.307 | 0.664 | 0.173 | 0.162 |
| High education: married men | 773 | 73.5 | 0.339 | 0.656 | 0.172 | 0.172 |
| Low education: married men | 846 | 73.1 | 0.329 | 0.687 | 0.169 | 0.144 |

**Table 11.7 Expectations and Realizations by Subgroup, 1973**

| Subgroup | Number of Observations | Mean of ERET | Fraction with RET = ERET | Fraction with RET = ERET ± 1 | Fraction with RET < ERET − 1 | Fraction with RET > ERET + 1 |
|---|---|---|---|---|---|---|
| Married men | 853 | 74.8 | 0.353 | 0.720 | 0.158 | 0.122 |
| Married women | NA | NA | NA | NA | NA | NA |
| Single men | 32 | 75.1 | 0.344 | 0.719 | 0.156 | 0.125 |
| Single women | 51 | 74.5 | 0.431 | 0.804 | 0.176 | 0.020 |
| Widowers | 48 | 74.3 | 0.375 | 0.688 | 0.188 | 0.125 |
| Widows | 140 | 74.5 | 0.386 | 0.707 | 0.179 | 0.114 |
| High wealth: married men | 574 | 74.9 | 0.357 | 0.711 | 0.171 | 0.118 |
| Low wealth: married men | 279 | 74.7 | 0.344 | 0.738 | 0.133 | 0.129 |
| High education: married men | 423 | 75.1 | 0.336 | 0.671 | 0.184 | 0.144 |
| Low education: married men | 430 | 74.6 | 0.370 | 0.768 | 0.133 | 0.100 |

**Table 11.8 Expectations and Realizations by Subgroup, 1975**

| Subgroup | Number of Observations | Mean of ERET | Fraction with RET = ERET | Fraction with RET = ERET ± 1 | Fraction with RET < ERET − 1 | Fraction with RET > ERET + 1 |
|---|---|---|---|---|---|---|
| Married men | 297 | 76.3 | 41.4 | 76.8 | 10.8 | 12.5 |
| Married women | 478 | 78.6 | 19.7 | 48.1 | 20.1 | 31.8 |
| Single men | 6 | 77.2 | 16.7 | 50.0 | 16.7 | 33.3 |
| Single women | 16 | 76.0 | 50.0 | 87.5 | 12.5 | 0.0 |
| Widowers | 21 | 75.5 | 38.1 | 66.7 | 33.3 | 0.0 |
| Widows | 50 | 76.3 | 40.0 | 78.0 | 18.0 | 4.0 |
| High wealth: married men | 198 | 76.3 | 40.9 | 75.8 | 11.6 | 12.6 |
| Low wealth: married men | 99 | 76.2 | 42.4 | 78.8 | 9.1 | 12.1 |
| High education: married men | 168 | 76.5 | 38.1 | 72.0 | 13.1 | 14.9 |
| Low education: married men | 129 | 75.9 | 45.7 | 82.9 | 7.8 | 9.3 |

view that unanticipated benefit increases led many workers to retire unexpectedly during the early 1970s.

# Notes

1. It is worth noting that the data do not appear to be consistent with the hypothesis that individuals report the medians of objective distributions. In fact, the pattern of medians is quite similar to the pattern of means.

2. It is possible to obtain the standard deviation of the mean retirement date in each cell from the standard deviation of the retirement date and the number of observations.

3. This assumption may seem peculiar when attrition is due to death. If, however, one believes (as seems natural) that individuals report expected dates of retirement conditional upon surviving until retirement, then the assumption is appropriate, since one wishes to know what each individual would have done had he survived.

# References

Aiginger, Karl. 1979. Empirische Information zur Bildung von Erwartungen. *Ifo-Studien* 25:83–135.

Anderson, Kathryn H., Richard V. Burkhauser, and Joseph F. Quinn. 1986. Do retirement dreams come true? The effect of unanticipated events on retirement plans. *Industrial and Labor Relations Review* 39:518–26.

Bernheim, B. Douglas. 1988. Social Security benefits: An empirical study of expectations and realizations. In *Issues in contemporary retirement*, ed. E. Lazear and R. Ricardo-Campbell. Stanford, Calif.: Hoover Institution. Pp. 312–48.

Burtless, Gary. 1986. Social security, unanticipated benefit increases, and the timing of retirement. *Review of Economic Studies* 53 (October): 781–805.

Curtin, Richard T. 1982. Determinants of price expectations: Evidence from a panel study. In *International research on business cycle surveys*, ed. M. Laumer and M. Ziegler. Aldershot: Gower.

Gramlich, Edward M. 1983. Models of inflation expectations formation: A comparison of household and economist forecasts. *Journal of Money, Credit, and Banking* 11:155–73.

Hall, Arden, and Terry Johnson. 1980. The determinants of planned retirement age. *Industrial and Labor Relations Review* 33:240–55.

Huizinga, John. 1980. *Real wages, employment, and expectations*. Ph.D. diss., Massachusetts Institute of Technology.

Hurd, Michael. 1983. The effects of Social Security on retirement: Results and issues. State University of New York, Stony Brook. Mimeo.

Hurd, Michael, and Michael Boskin. 1981. The effect of Social Security on retirement in the early 1970s. National Bureau of Economic Research Working Paper no. 659.

Mitchell, Olivia S., and Gary S. Fields. 1982. The effects of pensions and earnings on retirement: A review essay. In *Research in labor economics*, ed. Ronald E. Ehrenberg. Greenwich, Conn.: JAI Press.

Papadia, Francesco. 1982. Rationality of inflationary expectations in the European Community's countries. European Economic Community Working Paper, Brussels.

## Comment Edward P. Lazear

I am a fan of Bernheim's work, and this paper has done nothing to change my view. It is a careful analysis of the relation of expected retirement age to actual retirement age. The main question that I want to raise is, "Why is this interesting?" Or, put otherwise, "What else would we like to know?" Bernheim motivates the analysis by stating that life-cycle theory is based on the premise that consumers think seriously and coherently about the distant future. But for most of life-cycle theory, what is relevant is opportunities rather than outcomes. Retirement is an outcome, a point chosen on an opportunity locus, and it is not clear what it means or why it is relevant. Bernheim did some work on Social Security that spoke directly to this issue.

To make the point, consider the work-leisure diagram in figure 11.3. Point A is the expected outcome, and points B and C are actual outcomes under two different scenarios. Both B and C correspond to lower retirement ages than A since more years of retirement leisure are taken. The retirement age could be lower than the expected retirement age for one of two reasons. First, as illustrated by point B, opportunities get worse and wage offers fall. The substitution effect induces the worker to take more leisure, but he is worse off than he was at A. Second, income from nonlabor sources may rise, as illustrated by point C. The income effect induces the worker to take more leisure, and he is better off than he was at A. Bernheim looks at Social Security payments, so he goes part of the way there. But pension buyouts are important, as Laurence Kotlikoff and David Wise have found in their data (ch. 10, in this volume). Additionally, spousal income may be a factor. Although both cases show up as an actual retirement date that is earlier than the expected one, and although both have implications for life-cycle models, the implications are very different. In one case, early retirement implies an unanticipated fall in standard of living. In the other, it implies a rise.

The following extensions are the most important:

First, forecasts of variables that are the determinants, rather than the outcome, of retirement are more interesting. Data are lacking, except perhaps for health status.

Second, it would be useful to find out what happens to those who fail to predict their ages of retirement correctly. Specifically, what happens to income, assets, housing, and food consumption? Hausman and

Edward P. Lazear is a Professor of Industrial Relations at the Graduate School of Business, University of Chicago, a Senior Fellow, Hoover Institution, and a Research Associate of the National Bureau of Economic Research.

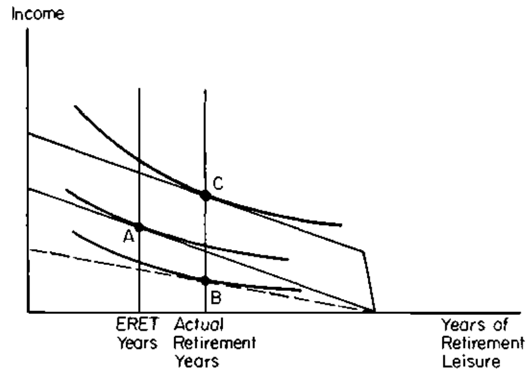


**Fig. 11.3** Work-leisure diagram

Paquette (1987) find that food consumption falls for those who suffer early retirement.

Third, the distinction between anticipated and unanticipated and between voluntary and compulsory retirement is significant. Voluntary might be defined as an improvement in one's leisure or alternative work opportunities. Compulsory might be defined as the current job situation getting worse. Is the group of those who retire earlier than anticipated dominated by compulsory or voluntary retirements? The answer has important implications for welfare.

Fourth, and related, married women forecast badly. They seem to work longer than expected. Are they widows who forecast their husband's life expectancy badly, or are they pleasantly surprised by the wonderful job offer that McDonald's made them to make french fries?

Fifth, and a more general way to put the point: When workers are wrong, does it matter? Is the change in rent associated with the unanticipated event large or small? For example, it could be that expected retirement deviates by a large amount from actual retirement because wages are close to alternatives so that a small change in either induces a large change in retirement ages. Under these circumstances, little consumer surplus is lost even when there are big differences between expected and actual retirement.

Sixth, are bad forecasters bad because they experience more unanticipated events than others or are they poorer at data processing? It is one thing to say that women forecast badly because they are faced with more uncertainty. It is another to say that they are worse at making decisions, given the same amount of uncertainty.

Here are some other minor points:

1. The analysis does not exploit the panel fully. Although it does look at whether predictions get better as the individual nears the

retirement age, it is useful to know how the expected age of retirement changes as one ages. Do all 55 year olds think that they will work to 70 and subsequently revise their forecasts downward, or do they think that they will work to 57 and revise it upward? More generally, what determines how ERET changes over time?

2. Are good Social Security forecasters also good retirement forecasters? This bears on the issue of whether bad forecasters are bad because of their data-processing abilities or because they face more uncertainty. If there is no consistency between ability to forecast Social Security and retirement, then perhaps difference in forecast ability reflects differences in the variance in opportunities, rather than in data processing.

This was a good paper. I learned something from it, and the results are credible. At worst, it does not go far enough.

## Reference

Hausman, Jerry A., and Lynn Paquette. 1987. Involuntary early retirement and consumption. In *Work, health and income among the elderly*, ed. Gary Burtless. Washington, D.C.: The Brookings Institution.