This PDF is a selection from an out-of-print volume from the National Bureau of Economic Research

Volume Title: Issues in the Economics of Aging

Volume Author/Editor: David A. Wise, editor

Volume Publisher: University of Chicago Press, 1990

Volume ISBN: 0-226-90297-8

Volume URL: http://www.nber.org/books/wise90-1

Conference Date: May 19-21, 1988

Publication Date: January 1990

Chapter Title: How Do the Elderly Form Expectations? An Analysis of Responses to New Information

Chapter Author: B. Douglas Bernheim

Chapter URL: http://www.nber.org/chapters/c7120

Chapter pages in book: (p. 259 - 286)

9 How Do the Elderly Form Expectations? An Analysis of Responses to New Information

B. Douglas Bernheim

A large fraction of the existing work on the economics of aging and the retirement period procedes on the basis of life-cycle assumptions, which hold that individuals form very rational and deliberate long-range plans. Implicit in these assumptions is the notion that individuals develop well-informed opinions about the economic factors that will affect their well-being in the future. Despite the existence of a small body of work on the accuracy of expectations concerning Social Security benefits (Bernheim 1988), the timing of retirement (Hall and Johnson 1980; Parnes and Nestel 1981; Anderson, Burkhauser, and Quinn 1986; Wolpin and Gönül 1987; and Bernheim 1989) and inflation (see Zarnowitz 1984 and the references contained therein), very little is actually known about the manner in which individuals incorporate new information into expectations.

The purpose of this paper is to examine the evolution of self-reported expectations about Social Security benefits during the preretirement period and to examine the responses of these expectations to the arrival of new information. The central questions are as follows. Do expectations evolve in the manner predicted by theory? What kind of information leads individuals to revise their expectations, and what is the nature of the responses? Are revisions "rational," in the sense that they closely resemble the effects of new information on objective measures of expected benefits? Since models of consumer decision making inevitably invoke a host of assumptions concerning expectations, these questions logically precede any analysis of behavior. I plan to study the relation between self-reported expectations and behavior in subsequent work.

B. Douglas Bernheim is the Harold Hines Jr. Distinguished Professor of Risk Management of the J. L. Kellogg Graduate School of Management, Northwestern University, and a research associate at the National Bureau of Economic Research.

Funding from the National Institute on Aging through its grant to the NBER is gratefully acknowledged.

The current investigation employs a data sample drawn from the Retirement History Survey (RHS). The longitudinal nature of this survey makes it possible to observe and compare expectations reported by the same households at different points in time and to relate observed changes to intervening events. My central conclusions are as follows.

First, a variety of simple tests appear to reject the most basic implications of the theory that forms the basis for this analysis. As in Bernheim (1988), I attribute these apparent failures to the fact that reported expectations are extremely noisy. When one corrects for the presence of reporting error through the appropriate use of instrumental variables, the resulting estimates are generally consistent with the theory. In particular, one cannot reject the hypotheses that expectations evolve as a random walk and that innovations in this process are unrelated to prior information.

Having concluded that the data support these basic implications, I use the theory to formulate an empirical specification that relates changes in expectations to the arrival of new information. Using this specification, I estimate responses of expectations to informational events and test for the rationality of these responses. The results are striking. Responses to new information during the period immediately preceding retirement appear to be highly rational. The bulk of information affects the evolution of expectations only through its effect on actual benefit levels computed from contemporaneous benefit formulas and earnings histories. Furthermore, the data support the view that individuals form accurate assessments of the ultimate effect of new information on actual benefits.

These results contrast sharply with findings based on analyses of expected benefit *levels* rather than *changes* in expected benefits. In Bernheim (1988), I found that certain variables—especially current statutory Social Security benefit entitlements—were highly correlated with subsequent forecast errors. This implies that individuals do not make complete use of all the information contained in these variables. Nevertheless, these same individuals are very good at processing information that arrives just prior to retirement. Specifically, while they are apparently incompletely informed about the level of benefits associated with contemporaneous benefit formulas, they revise expectations as if they understand how new information affects the benefits prescribed by these formulas *on the margin.* This result suggests that individuals formulate expectations about the retirement period much more carefully as retirement approaches and therefore corroborates some speculative conclusions based on more sketchy evidence that appeared in Bernheim (1988). At the same time, this finding supports the hypothesis that, because individuals appreciate the links between behavior and benefits at the margin, benefit formulas may have incentive effects. This hypothesis has formed the basis for many previous studies of the retirement decision (see Hurd 1983).

The remainder of this paper is organized as follows. Section 9.1 presents the basic model of expectations. I discuss the data in section 9.2. Section 9.3

contains tests of the model's central implications, and section 9.4 examines responses of expectations to new information. The paper closes with a brief conclusion.

9.1 A Model of Expectations

Suppose that, at each point in time t, an individual forms an expectation, X_t^e, about the value of a variable X that is realized at some point in the future. During period t, he has access to certain information, which I denote Ω_t. Throughout, I assume that the individual's memory is perfect, so that all information available at time t is also available in period $t + 1$. Formally, $\Omega_{t+1} = (\Omega_t, \omega_{t+1})$, where ω_{t+1} represents information that becomes available between periods t and $t + 1$.

In subsequent sections, I interpret X as Social Security benefits. When an individual reports expected Social Security benefits, there is, of course, some ambiguity as to what this means. While he may have in mind something like a mathematical expectation, it is also possible that his report reflects his view of the most likely outcome (i.e., the mode). As long as the distribution of X is approximately symmetric and single peaked, this ambiguity is probably of very little consequence. However, it is important to bear in mind that the mathematical interpretation that one places on a reported expectation becomes a joint hypothesis with any other proposition that one wishes to test. In particular, failure of tests for "rationality" (discussed below) could simply reflect misinterpretation of the reported data. With this qualification in mind, I henceforth focus on the hypothesis that individuals report expected values, that is,

$$X_t^e = E(X|\Omega_t), \tag{1}$$

where E is the expectations operator.

From equation (1), it follows that

$$E(X_{t+1}^e|\Omega_t) = E\,[E(X|\Omega_t, \omega_{t+1})|\Omega_t] = E(X|\Omega_t) = X_t^e. \tag{2}$$

This expression describes the stochastic evolution of expectations through time and is the basis for the conclusion that expectations should follow a random walk. In particular, (2) implies that

$$X_{t+1}^e = X_t^e + \eta_{t+1}, \tag{3}$$

where

$$E(\eta_{t+1}|\Omega_t) = 0. \tag{4}$$

Furthermore, η_{t+1} should be a function of new information received since period t, ω_{t+1}.

The analysis of this paper is based on the simple model described in equations (3) and (4). Using these as the basis for an empirical specification, I investigate the manner in which expectations respond to new information. The validity of my empirical results depends critically on the appropriateness of this underlying framework. It is therefore essential to test the framework as thoroughly as possible.

Fortunately, the model lends itself to a number of direct tests. Note that we can write

$$\text{(5)} \qquad \begin{aligned} \operatorname{var}(X_{t+1}^e) &= \operatorname{var}(X_t^e) + \operatorname{var}(\eta_{t+1}) \\ &= \operatorname{var}(X_t^e) + \operatorname{var}(X_{t+1}^e - X_t^e) > \operatorname{var}(X_t^e). \end{aligned}$$

Two implications follow directly from equation (5). First, the population variance of expectations reported at a particular point in time should be greater than the population variance of expectations reported at earlier points in time. Second, the difference between these population variances should be *exactly* the variance of their differences. In Bernheim (1987), I studied the first of these implications and found the data somewhat supportive. However, since the focus of that study was a comparison of expectations and realizations (rather than a comparison of expectations at different points in time), I did not consider the second implication.

Equations (3) and (4) also suggest a regression format that facilitates further testing of the underlying model. Suppose in particular that we use ordinary least squares to estimate an equation of the form

$$\text{(6)} \qquad X_{t+1,i}^e = \alpha + \beta X_{t,i}^e + \Omega_{t,i}\gamma + \epsilon_{t,i}\,,$$

where i indexes individuals. Theory implies that we should obtain $\alpha = \gamma = 0$ and $\beta = 1$. Furthermore, our estimate of σ_ϵ^2 measures σ_η^2. This test is quite demanding, in that the underlying hypothesis includes the assertion that, in forming his expectation, the individual actually uses—and uses efficiently—all information observed by the econometrician. I therefore refer to it as a test of "strong" rationality. One can also conduct a weaker, less demanding test by omitting the informational variables and simply regressing X_{t+1}^e on X_t^e. Theory implies that the intercept and slope coefficients should be zero and one, respectively. This test allows for the possibility that individuals do not form expectations on the basis of all available information. However, the underlying hypothesis retains the key feature that expectations evolve as a random walk, responding only to new information.

If the tests of the underlying model prove favorable, then one can use the model of expectations embodied in equations (3) and (4) to measure responses of expectations to new information. Since η_{t+1} is related exclusively to *new* information, I write it as a function of *surprises:*

$$\eta_{t+1} = \psi[\omega_{t+1} - E(\omega_{t+1}|\Omega_t)].$$

Substitution into (3) yields an expression for adjustments in expectations:

$$X^e_{t+1} - X^e_t = \psi[\omega_{t+1} - E(\omega_{t+1}|\Omega_t)]. \tag{7}$$

This in turn suggests a regression of changes in expectations on variables that contain new information received after period t. The coefficients in this regression will reflect the magnitude of responses to particular types of information. Implementing this strategy is somewhat problematic, in that all variables have both expected and unexpected components and therefore measure blends new and old information. I take up specific estimation issues in section 9.4.

9.2 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 (58–63 years old in 1969) for a period of ten years, beginning in 1969. Each household was surveyed once every two years (1969, 1971, 1973, 1975, 1977, and 1979). Although the initial wave included more than 11,000 households, there was substantial attrition over successive waves.

In 1969, 1971, and 1973, respondents reported the level of Social Security benefits that they expected to receive on retirement. In subsequent sections, the variables ESS71 and ESS73 (expected Social Security in 1971 and 1973, respectively) reflect answers to these questions, adjusted to an annual basis. Inspection of the data for 1969 revealed a low response rate (due in part to survey skip patterns) as well as a high frequency of nonsensical values. I have therefore confined attention to responses given in 1971 and 1973. In what follows, the variable CESS measures the *change* in expectations between these two years (i.e., CESS = ESS73 − ESS71).

Unfortunately, interpretation of expected benefits is somewhat problematic, in that the treatment of inflation is ambiguous. Certainly, the survey instrument does not specify whether the individual is to report a real or a nominal figure. Throughout, I simply assume that respondents report expected benefits in current (i.e., survey year) dollars. This seems the most natural choice since respondents would otherwise have had to forecast future inflation rates before formulating an answer to the question. To the extent that my assumption is incorrect, the scale of expectations may vary somewhat between 1971 and 1973.

Tests of the strong rationality hypothesis, as well as some of the other exercises conducted in section 9.4 of this paper, require the collection of informational variables that are candidates for inclusion in Ω_t. In this paper, I employ essentially the same informational variables as in Bernheim (1988). I group these into three distinct categories.

The first category contains variables that measure other reported expectations. These are natural candidates for inclusion in Ω_t since they necessarily reflect information that the individual has used to generate forecasts. If any of these variables appear with significant coefficients in estimates of equation (6), it would indicate that, at a minimum, individuals use different kinds of information to form different expectations. Definitions of specific variables follow:

ERET71: Expected date of retirement reported in 1971.
ERET69: Expected date of retirement reported in 1969.
EOI71: Expected retirement income other than Social Security, reported in 1971.
EOI69: Expected retirement income other than Social Security, reported in 1969.

Data on expectations are, of course, incomplete—many individuals who report expected Social Security benefits do not, for example, report an expected date of retirement. Accordingly, I also use dummy variables, DRET71 and DRET69, which equal one if the individual reports the associated expectation and zero otherwise. In the final sample (described below), all individuals responded to questions about retirement income other than Social Security, so no companion dummies for the EOI71 and EOI69 variables were required.

The second category includes various demographic variables and other household characteristics that might be useful in predicting future Social Security benefits. The list of variables includes:

AGE: The respondent's age.
SAGE: The respondent's wife's age.
ED: The respondent's education (measured in number of years).
SED: The respondent's wife's education.
W: The household's net wealth (including financial assets, businesses, and real property).
HGOOD: A dummy variable, indicating whether the respondent reports his health as being better than average for his age (1 = better, 0 = other).
HBAD: A dummy variable, indicating whether the respondent reports his health as being worse than average for his age (1 = worse, 0 = other).
KIDS: Number of living children.
COMPRET: A dummy variable, indicating whether the respondent's employer maintains a compulsory retirement age (1 = yes, 0 = no).
MOVE: A dummy variable, indicating whether the respondent has moved within the past two years (1 = has moved, 0 = has not moved).

The third and final category consists of a single variable, which is the individual's current Social Security entitlement, CSS71, defined as the level of benefits he would receive under current law if he retired immediately. CSS71 is, theoretically, part of each individual's information set in 1971, in that it depends only on his own past earnings history and on current law (which is public information). Special treatment of CSS71 is warranted in light of my earlier findings (Bernheim 1988), which indicated that individuals fail to use much of the information contained in this variable and furthermore that, quantitatively, this is by far the most important source of unused information.

Since this study focuses on the responses of expectations to new information, it is also essential to compile a list of variables that are candidates for inclusion in ω_{t+1}. Each of the following variables describes some aspect of a change in an individual's status between 1971 and 1973 and could conceivably be related to the ultimate realization of Social Security benefits:

HBET: A dummy variable, indicating whether the self-reported index of health status improved (1 = improvement, 0 = other).

HWOR: A dummy variable, indicating whether the self-reported index of health status deteriorated (1 = deterioration, 0 = other).

WIDM: A dummy variable, indicating whether the wife died between 1971 and 1973 (1 = wife died, 0 = other).

WIDW: A dummy variable, indicating whether the husband died between 1971 and 1973 (1 = husband died, 0 = other).

LJOB: A dummy variable, indicating whether the respondent was employed in 1971 but not in 1973 (1 = lost job, 0 = other).

GJOB: A dummy variable, indicating whether the respondent was employed in 1973 but not in 1971 (1 = obtained job, 0 = other).

CJOB: A dummy variable, indicating that the respondent was employed in different jobs in 1971 and 1973 (1 = different jobs, 0 = other).

NMOVE: A dummy variable, indicated whether the respondent moved between 1971 and 1973 (1 = moved, 0 = other).

CW: The change in the respondent's wealth between 1971 and 1973.

CCSS: The change in the respondent's statutory Social Security entitlement between 1971 and 1973.

Finally, while the focus of this analysis is on changes in expectations (rather than on the accuracy of expectations *per se*), some of the exercises in section 9.4 require measures of ultimate realizations. I calculate each realization by applying the benefit formula in effect at the individual's date of retirement to earnings histories from matching administrative records provided by the Social Security Administration. For details, I refer the reader to Bernheim (1988).

The basic sample population for this analysis consisted of RHS respondents who in 1971 were married and not yet receiving Social Security benefits. Individuals who failed to report expectations about Social Security benefits in 1971, as well as a few who reported nonsensical values (in excess of $20,000 per year), were dropped. In order to compare expectations across years, I restricted attention to respondents who still had not begun to receive Social Security benefits in 1973 and who reported an expectation in that year as well. I dropped a small number of observations for which key variables (marital status in 1973, health status in 1973, spouse's age, number of children, and compulsory retirement) were either missing or nonsensical. The resulting sample contained one individual who failed to report an expectation about retirement income other than Social Security in either 1969 or 1971—rather than create a dummy variable like DRET71, I simply dropped this observation. This left a total of 370 observations.

Since the sample used here is a rather small fraction of the total survey population, one naturally wonders whether it is very representative. In particular, the majority of individuals fail to report expectations. Nonreporting might itself reflect a failure to think seriously about the retirement process. If so, statistical analysis based on the fraction of the sample that reports expectations may be very misleading.

Fortunately, nonreporting appears to be fairly random and is perhaps more commonly attributable to fatigue resulting from the length of the survey instrument or to the styles of different interviewers. If nonreporting reflected a failure to think seriously about and plan for the retirement period, then one would expect nonreporting of expected benefits and nonreporting of expected retirement dates to be highly correlated. In fact, this is not the case. Of those married men who reported expected Social Security benefits in 1971, 42 percent also reported an expected retirement date. For those who did not report expected benefits in 1971, the figure was only slightly lower (40 percent). Of those who reported expected benefits in 1973, 34 percent also reported an expected retirement date. For those who did not report expected benefits in 1973, the figure was slightly higher (36 percent). In addition, there is only a mild correlation between reporting of expected benefits in 1971 and reporting of expected benefits in 1973. Forty-five percent of married males reported expected benefits in 1971, as did 39 percent in 1973. Of those who reported expected benefits in 1971, only 49 percent also reported this expectation in 1973.

One might also argue that those who reported expected benefits in both 1971 and 1973 could be atypical. Some insight into this issue can be gained from considering a few summary statistics. Table 9.1 contains means and standard deviations for the variables that measure expectations about Social Security benefit levels (ESS71 and ESS73), the change in expectations (CESS), and the actual realization (SS). Note that the average expected benefit rose just over 2 percent between 1971 and 1973. In 1971, expectations were about 10 percent

Table 9.1 **Summary Statistics on Expectations**

Variable	Mean	Standard deviation
ESS71	2,307	881
ESS73	2,362	1,164
CESS	55	1,229
SS	2,550	1,003

lower than realizations, while in 1973 they were about 8 percent lower. All these numbers (including the standard deviations) coincide very closely to summary statistics presented in Bernheim (1988). Those earlier calculations were based on much larger samples owing to the fact that it was not necessary to restrict attention to respondents who reported expected benefits *both* in 1971 *and* in 1973 (in that paper, the object was to compare expectations to realizations rather than to subsequent expectations). The similarity of these summary statistics suggests that the smaller sample is representative.

Before passing on to analysis of the data, it is important to discuss two potential problems. The first concerns sample selection biases. Many of the criteria for dropping observations are based on characteristics that were observed in 1971. In principle, such factors are part of Ω_{71}, the respondent's information set in 1971, and, according to theory, are therefore unrelated to η_{73}. Sample selection of this sort is therefore not likely to produce systematic biases. Other selection criteria are based on characteristics that are observed after 1971. In principle, these could be systematically related to new information and hence to η_{73}. In Bernheim (1988), I argued that some of these (e.g., attrition due to death) are not likely to create significant problems. Unfortunately, owing to the nature of the current excercise, I have had to impose more demanding requirements on data availability during the period after an expectation is reported (most important, the individual must report an expectation in 1973 as well as in 1971). This enhances significantly the probability that one or more of the selection criteria are in fact problematic. I have therefore given some explicit attention to these issues in the econometric implementation.

The second and perhaps more serious problem concerns nonindependence of realizations. Tests such as those described in section 9.1 are most commonly conducted with time-series data on the same individual or set of individuals, so that, under the null hypothesis, independence of the error terms is guaranteed. When one instead relies primarily on cross-sectional data from a short panel such as the RHS, theory does not rule out systematic correlation of error terms across observations. Correlation could arise for a variety of reasons.

The most important potential source of correlation is a macro event that affects a significant fraction (perhaps all) of the sample simultaneously. Suppose, for example, that subsequent to the date at which X_t^e is recorded, Congress unexpectedly raises benefits by 20 percent. Assuming that individuals process this information, one would presumably discover that on average η_{t+1} is significantly positive. Such an event did in fact occur in September 1972. However, this was for the most part an across-the-board increase in benefit levels. As a result, it probably affected little more than the scale of expectations. To put it another way, one would not be surprised to find $\beta > 1$ in estimates of equation (6), and one should not construe this as contrary to theory. Indeed, through estimates of β, one can hope to discern the extent to which this change was either anticipated *ex ante* or ignored *ex post*. Finally, one would still expect to find $\alpha = \gamma = 0$ if the theory is accurate. The data would fail to satisfy these restrictions only if elements of Ω_{71} were related to the probability of processing information about the new law (or processing it correctly) or to the nature of behavioral responses to the law. I tend to discount both possibilities. In particular, the results in Bernheim (1988) suggest that the 1972 legislation was largely anticipated, and the summary statistics in table 9.1 show little evidence of an upward surge in expectations after 1972. Furthermore, the analyses of Burtless (1986) and Bernheim (1989) suggest that the effect of the 1972 legislation on the timing of retirement was small.

9.3 Tests of the Model

In this section, I test various implications of the model presented in section 9.1 using the data described in section 9.2. This nature of these tests is very similar to those in Bernheim (1988), except that in my earlier work I focused on the relation between realizations and expectations rather than that between expectations at different points in time. Many findings from my earlier study are relevant to, and corroborated by the results of, this section. Most important, the previous study found that survey responses to questions about expected benefits are quite noisy and that failure to deal with this problem leads to apparent rejection of the theory. However, when the noise is treated through an appropriate instrumental variables technique, the results are highly favorable to the hypothesis of weak rationality and indeed indicate that individuals are quite good at forming expectations based on the subset of information that they do use. These issues reappear in the current context and must be dealt with explicitly.

9.3.1 Tests of Weak Rationality

Section 9.1 describes a theory of information processing. That theory does not necessarily assume or imply that individuals use all the information that is in principle available to them. Fortunately, even in the absence of any prior knowledge about what kinds of information individuals do and do not use to form expectations, the theory still has some testable implications. As I have already discussed, there are several natural tests based on equation (5), and these certainly do not require knowledge of Ω_t. In addition, since X_t^e is (trivially) part of the information set used in forming expectations at time t, expectations always evolve according to equation (3), where in place of (4) we substitute

$$E(\eta_{t+1}|X_t^e) = 0,$$

(i.e., they always follow a random walk). Thus, another minimalistic test would be based on ordinary least squares estimation of the equation

$$X_{t+1}^e = \alpha + \beta X_t^e + \eta_{t+1} . \tag{8}$$

Regardless of what Ω_t contains, theory implies that $\alpha = 0$ and $\beta = 1$. This section is devoted to the implementation of these tests.

I begin with tests based on equation (5). The summary statistics in table 9.1 are certainly consistent with the prediction that $\text{var}(X_{t+1}^e) > \text{var}(X_t^e)$, and indeed this corroborates the finding of Bernheim (1988). However, on the basis of these statistics, it is also evident that support for the theory is superficial at best. In particular, the difference between the variances *cannot* equal the variance of the differences (i.e., var[CESS]) since the latter by itself exceeds var(ESS73). Indeed, the standard deviation of ESS73 would have to be about 30 percent larger than its actual value in order to satisfy the equality in (5).

One can make this same point through estimation of equation (8). Results for ordinary least squares are contained in column 1 of table 9.2. Note that the intercept is quantitatively large and statistically significant, while the slope is less than one-half and estimated with great precision. On the basis of these estimates, one would be inclined to conclude that the data resoundingly reject even the simplest implications of our central hypothesis.

Fortunately, this negative conclusion is premature. As emphasized in Bernheim (1988), much evidence indicates that expectations about Social Security benefits are reported with a great deal of noise. This may at first seem peculiar. With a variable like wealth or income, noise may arise from imprecise measurement on the part of respondents. In contrast, an individual creates his own expectations and therefore cannot have any problem measuring

Table 9.2 **Tests of Weak Rationality**

	Equation				
Variable	1	2	3	4	5
Technique	OLS	IV	IV	IV-Heckit	IV-Heckit
Intercept	1,429	−559	−307	−93.1	−213
	(176)	(1,287)	(661)	(879)	(685)
ESS71	.400	1.27	1.16	1.37	1.22
	(.0791)	(.557)	(.285)	(.400)	(.373)
MILLS				−897	−314
				(757)	(546)

them. There are, however, other sources of noise. Some individuals may tend to exaggerate, reporting a higher number than they believe, while others may be prone to understate their assets. Alternatively, individuals may use relatively precise figures when formulating financial plans but provide only "ballpark" figures to interviewers. Respondents might also think in terms of replacement rates (i.e., the percentage of preretirement income provided by Social Security) rather than absolute levels and may err in the process of converting one to the other. Finally, some noise is undoubtably attributable to recording and coding errors.

The analysis of my previous paper established that a standard errors-in-variables specification, combined with the basic theory of expectations outlined above, explained the relation between expectations and realizations rather well. It is therefore quite possible that reporting error also accounts for the apparent failure of the theory in the current context.

Unfortunately, one cannot in the absence of additional information adjust the tests based on equation (5) for the presence of reporting error. Nevertheless, one can "back out" the variance of the measurement error that would make the observed variances consistent with theory. This is accomplished as follows.

Suppose that for each τ we observe $\tilde{X}^e_\tau$, which is related to X^e_τ as follows:

$$\tilde{X}^e_\tau = X^e_\tau + \mu_\tau ,$$

where μ_τ is uncorrelated with X^e_τ. Suppose further that the μ_τ are independently and identically distributed, with variance σ^2_μ. Then equation (5) implies that

$$\text{var}(\tilde{X}^e_{t+1}) - \sigma^2_\mu = \text{var}(\tilde{X}^e_t) + \text{var}(\tilde{X}^e_{t+1} - \tilde{X}^e_t) - 3\sigma^2_\mu .$$

From this expression, it follows that

$$\sigma^2_\mu = \frac{\text{var}(\tilde{X}^e_t) + \text{var}(\tilde{X}^e_{t+1} - \tilde{X}^e_t) - \text{var}(\tilde{X}^e_{t+1})}{2} .$$

Substitution of the summary statistics from table 9.1 into this formula reveals that $\sigma_\mu = 682$, so that approximately 60 percent of the variance in ESS71 and 35 percent of the variance in ESS73 is attributable to measurement error.

While the preceding calculation assumes the existence of reporting error, one can actually test this hypothesis through estimation of equation (8). The standard prescription for reporting error is to employ instrumental variables. One requires that the instrument is uncorrelated with both η_{t+1} and μ_t but correlated with X^e_t. Accordingly, valid instruments must be related to information that the individual actually uses to construct X^e_t. Thus, one necessarily tests the basic theory and the measurement error hypothesis jointly with the assumption that individuals actually use certain information (i.e., that contained in the instruments) in a manner consistent with theory.

The second column in table 9.2 contains estimates of (8) for which I have instrumented ESS71 with measures of other expectations (i.e., the first group of variables discussed in sec. 9.2 as candidates for inclusion in Ω_t). The use of these variables as instruments is based on the plausible assumption that individuals' expectations are internally consistent, in the sense that all expectations are based on the same information. The results in Bernheim (1988) lend strong support to this assumption. Note that the estimated coefficients change dramatically. The intercept is now negative and statistically insignificant, while the slope coefficient rises to 1.27 and is statistically indistinguishable from unity.

The third column in table 9.2 contains estimates of (8) for which I have instrumented ESS71 with various socioeconomic and demographic variables (i.e., the second group of variables discussed in sec. 9.2 as candidates for inclusion in Ω_t). The use of these variables as instruments is supported by the findings in Bernheim (1988)—while individuals do not appear to use all this information efficiently, the extent of the departure from theory is not of much quantitative importance. Once again, the estimated coefficients change dramatically. The intercept becomes negative and statistically insignificant, while the slope coefficient rises to 1.16 and is statistically indistinguishable from unity.

For both sets of estimates, one cannot reject the hypothesis that $\alpha = 0$ and $\beta = 1$ at reasonable levels of confidence. Of course, this is in large part due to the fact that standard errors are enormous. By itself, this evidence is only weakly supportive of the underlying hypotheses. It becomes far more persuasive in the context of my earlier results. In regressions of realizations on expectations (see Bernheim 1988), precisely the same pattern emerged—simple regressions produced large positive intercepts and slope coefficients of roughly .5, while instrumental variables techniques drove the intercepts toward zero and generated slope coefficients of about 1.1. Furthermore, since the earlier study made use of much larger samples, the precision of these estimates was substantially greater. The fact that the predicted pattern arises in two different estimation contexts lends strong support to the underlying joint hypotheses.

It is also possible to "back out" estimates of σ_μ from the IV results. Standard calculations reveal that the bias in the OLS estimate of the slope parameter is proportional to the noise-to-signal ratio. Furthermore, the IV estimates are consistent. Using these facts, it is easy to show that

$$\hat{\sigma}_\mu^2 = \hat{\sigma}_{\tilde{x}}^2(1 - \beta_{\text{ols}}/\beta_{\text{iv}})$$

yields a consistent estimate of σ_μ^2, where $\hat{\sigma}_{\tilde{x}}^2$ is the population variance of $\tilde{X}_t^e$, and β_{ols} and β_{iv} are, respectively, the OLS and IV estimates of β. The preceding paragraphs describe two sets of IV results. For the first set, the implied value of $\hat{\sigma}_{\mu}$ is 728, while for the second it is 712. Since the estimated β's are quite close to unity, these values are not far from the figure derived from equation (5) (i.e., 682). Moreover, one can undertake a similar exercise for the regressions of realizations on expectations contained in Bernheim (1988). The implied variance for measurement error for 1971 is 660, which is in the same ballpark. The striking similarity of estimates obtained from two distinct empirical exercises again lends support to the joint hypotheses outlined above.

In section 9.2, I mentioned that this analysis suffers from potential sample selection problems. To assess the importance of these factors, I introduced a statistical correction based on the procedure outlined by Heckman (1976). First, I created a larger data sample containing the original sample plus all the observations that were excluded on the basis of characteristics observed after 1971. Next, I estimated a probit relation that explained inclusion in the original sample as a function of the instrument list and used these estimates to form inverse Mill's ratios. I then augmented equation (8) with the inverse Mill's ratio term and estimated it with two-stage least squares, using both the original instrument list and the inverse Mill's ratio as instruments. This procedure treats both the endogeneity of ESS71 and the sample selection problem simultaneously and yields consistent estimates.

Results for the two instrument lists discussed above appear in columns 4 and 5 of table 9.2. While the slope coefficients rise slightly, this change is dwarfed by the original standard errors. In addition, the Mill's ratios do not appear to enter significantly (note, however, that I have not adjusted the standard errors for the fact that these terms are estimated rather than observed). Overall, the sample selection correction appears to make very little quantitative or qualitative difference. Indeed, none of the estimates in this paper were significantly affected by the introduction of similar corrections. In subsequent sections, I have conserved space by presenting only uncorrected OLS and IV estimates. Results based on sample selection corrections are available on request.

In summary, the data are consistent with the hypothesis of weak rationality. This fact is obscured by the presence of significant reporting error, which biases simple regression estimates and leads to apparent rejections of the theory. Unfortunately, estimates that correct for the presence of measurement error are imprecise, so that the associated tests have little power. However, taken in conjunction with previous work, this analysis validates the use of weak rationality as a maintained hypothesis in subsequent sections.

9.3.2 Tests of Strong Rationality

In my previous study of expectations and realizations (Bernheim 1988), I found that, while the data were consistent with the hypothesis of weak rationality, they were highly inconsistent with strong rationality. In particular, individuals appeared to ignore much of the information contained in current statutory entitlements and to a lesser extent failed to make complete use of several socioeconomic variables.

In the current context, tests of strong rationality have a much different flavor. To understand these differences, consider equation (6). If we replace X^e_{t+1} with X (so that the equation explains realizations rather than later expectations), then any failure to process information contained in Ω_t should show up as nonzero components in the coefficient vector γ. However, as the equation stands, elements of γ will be nonzero only if either (i) individuals are slow to adjust expectations, and incorporate certain aspects of Ω_t into their forecasts sometime after period t and before period $t + 1$, or (ii) individuals ignore elements of Ω_t that are useful in predicting events that these individuals *will* subsequently incorporate into their forecasts. Failure to reject the hypothesis that $\gamma = 0$ does not, in the current context, imply that individuals process all information correctly. Most obviously, if individuals *never* adjust their expectations, then we will certainly estimate $\gamma = 0$, despite the fact that expectations are not informationally efficient. Thus, the tests of strong rationality have power against a much narrower range of alternatives in the current context than in my earlier paper.

I implement these tests through estimation of equation (6). In light of my conclusions concerning the presence of reporting error, it is hardly surprising that OLS estimates of (6) are highly at variance with the theory. I therefore omit these results and turn directly to procedures that correct for this problem.

There are two alternative methods of dealing with measurement error. First, one can impose the constraint that $\beta = 1$, thereby moving $\tilde{X}^e_t$ to the left-hand side of the equation. The term μ_t then becomes part of the standard regression error; while it renders the estimates less precise, it does not affect consistency. One can then test the hypotheses that $\alpha = \gamma = 0$. Second, one can estimate (6) with instrumental variables. It is then possible to test all the relevant constraints (including $\beta = 1$). The drawback of this approach is that, as in the previous section, in order to identify instruments one must maintain the hypothesis that individuals actually use certain information.

Table 9.3 contains the results of the procedures outlined in the preceding paragraph. Estimates in the first column are generated by regressing the change in expectations (CESS) on the full list of informational variables. Note that none of the corresponding coefficients is significant at the 95 percent level of confidence. Even CSS71, which played such a large role in my earlier analysis of expectations and realizations, appears to explain very little of the change in expectations. In fact, the F-statistic for the hypothesis that $\gamma = 0$ is .834, and the F-statistic for the joint hypotheses that $\alpha = \gamma = 0$ is .829, so that it is impossible to reject strong rationality at any standard level of confidence.

Failure to reject might, of course, be attributable to imprecision. It is therefore appropriate to consider the magnitudes of point estimates. Certain

Table 9.3 Tests of Strong Rationality

Variable	Equation 1	Equation 2	Equation 3
Dependent variable	CESS	ESS73	ESS73
Technique	OLS	IV	IV
Intercept	−3,593	−4,427	10.2
	(2,913)	(6,352)	(554)
ESS71		1.24	.966
		(.688)	(.286)
ERET71	40.1		51.4
	(46.8)		(50.3)
DRET71	−2,801		−3,519
	(3,446)		(3,708)
EOI71/100	−.184		−.230
	(.173)		(.335)
ERET69	8.89		1.51
	(39.9)		(39.0)
DRET69	−649		−130
	(2,920)		(2,848)
EOI69/100	−1.92		−1.68
	(2.89)		(3.10)
CSS71	−.102	−.162	−.0284
	(.0744)	(.202)	(.0892)
AGE	49.5	63.0	
	(50.8)	(108)	
SAGE	13.8	9.74	
	(14.3)	(27.1)	
ED	−3.99	−5.15	
	(12.1)	(12.7)	
SED	1.14	−2.84	
	(11.5)	(16.4)	
$W/10^4$	.481	−1.05	
	(8.41)	(10.2)	
HGOOD	−156	−175	
	(141)	(157)	
HBAD	−122	−158	
	(223)	(251)	
KIDS	5.18	3.16	
	(39.0)	(38.2)	
COMPRET	329	408	
	(179)	(200)	
MOVE	−203	−219	
	(223)	(259)	

coefficients stand out as very large relative to the mean value of expected benefits. The most notable among these are DRET71 and DRET69. The reason for this is simply that the variables ERET71 and ERET69 have also been included in the regression. Since the mean value of ERET71 is around 74, the product of this variable with its coefficient is typically around 2,900. The corresponding dummy variable simply takes out the mean of this product so that the fitted value of CESS is not substantially different for those who do and do not report ERET71. Since the t-statistic for the coefficient of ERET71 is small, the standard error for the coefficient of DRET71 must be enormous.

Other variables with quantitatively significant coefficients are HGOOD, HBAD, COMPRET, and MOVE. Of these, only the coefficient of COMPRET approaches statistical significance. Nevertheless, it is somewhat disturbing that the standard deviations of these coefficients are so large. For example, although the point estimates indicate that a recent move is associated with roughly an 8 percent decline in expected benefits during the subsequent period, we are unable to determine with any reasonable confidence whether this association is the result of chance.

The second column of table 9.2 contains IV estimates, where the instrument list consists of other reported expectations (i.e., the first set of variables listed in sec. 9.2 as candidates for inclusion in Ω_t). The coefficient of ESS71 is only slightly changed from the corresponding regression in table 9.2. Of the various informational variables, only COMPRET appears with a significant coefficient. Of course, with a large number of informational variables, it is hardly surprising that one should appear with a coefficient that is significant at the 95 percent level of confidence. A formal test of the hypothesis that none of the informational variables matters ($\gamma = 0$) reveals that this hypothesis cannot be rejected. Similarly, the data fail to reject the full implications of strong rationality—$\alpha = \gamma = 0$ and $\beta = 0$—at the 95 percent level of confidence.

These conclusions follow with even greater force from estimates based on the use of socioeconomic and demographic variables (i.e., the second set of variables listed in sec. 9.2 as candidates for inclusion in Ω_t) as instruments. The associated results appear in the third column of table 9.3. Note that the intercept is nearly zero, that the estimate of β differs only slightly from unity, and that none of the informational variables appears with either a statistically significant or a quantitatively important coefficient (recall my earlier comments concerning the interpretation of the coefficient for DRET71). Not surprisingly, one cannot reject the hypothesis of strong rationality on the basis of these estimates.

Taken together, these results bear out the strongest implications of the theory outlined in section 9.1. One should, however, be cautious in interpreting these results. In this regard, it is worth reiterating some of the opening remarks for this subsection. This evidence suggests that we can rule out the possibilities that (i) individuals incorporate certain information into their expectations only after a lag and (ii) information that individuals fail to use is highly correlated with subsequent events that they do incorporate into their expectations. The evidence does not allow us to conclude that individuals make efficient use of all available information, and indeed the results of Bernheim (1988) suggest the contrary.

9.4 Responses to New Information

The analysis of section 9.3 lends support to the theoretical model of expectations outlined in section 9.1. Unfortunately, it does not tell us very much about the manner in which individuals process new information. For example, this evidence does not rule out the possibility that individuals form expectations at some early date and thereafter cling stubbornly to their original forecasts, ignoring all new information. The current section is therefore devoted to an analysis of the manner in which new information affects the evolution of expectations.

On the basis of the simple summary statistics in table 9.2, it seems apparent that some adjustment of expectations occurs. For one thing, the variance of CESS is very large. Of course, this could be partly attributable to the fact that both ESS71 and ESS73 contain measurement error—indeed, the observed variance of CESS could in principle be entirely spurious. If the variance of measurement error remains constant over time, then the variance of CESS simply equals the variance of the true change in expectations plus two times the variance of the measurement error, σ_{μ}^2. In section 9.3.1, I presented several different estimates of σ_{μ}, all of which clustered around 700. Combining this figure with the observed standard error of CESS, it is possible to recover the variance of the true change in expectations. Specifically, I calculate the standard error of the true change to be 728. Thus, individuals appear to have adjusted their expectations significantly between 1971 and 1973. One can illustrate this same point simply by comparing the variances of ESS71 and ESS73—unless measurement error increased dramatically between these years, the rise in variance must reflect the processing of new information.

The observations raise two important questions. First, what kind of information leads individuals to revise their expectations, and what is the nature of the response? Second, do individuals process new information "rationally," in the sense that the adjustment of observed expectations closely resembles an adjustment to some objective expectation of the realized value? The next two subsections are devoted to analyses of these questions.

9.4.1 Measurement of Responses

The starting point for this analysis is equation (7), which relates changes in expectations to unanticipated events. To the extent that such events determine subsequent earnings, applicable statutes, or the timing of retirement, they may also have large effects on eventual benefits. Estimation of equation (7) requires some notion of what the function ψ looks like as well as some technique for distinguishing between anticipated and unanticipated events. Lacking any prior information about the form of ψ, I simply estimate a linear approximation. In addition, I try out three different procedures for measuring unanticipated events.

It is natural to begin with the simple assumption that expectations are largely myopic, so that any change in status is unanticipated. This motivates a regression of CESS on the set of variables listed in section 9.1 as candidates for inclusion in ω_{t+1}. Since my earlier study (Bernheim 1988) suggested that individuals ignore much of the information contained in current statutory entitlements, I begin with a regression that omits CCSS (the change in current entitlement) from this list. The results appear in column 1 of table 9.4.

Only one of the variables in this regression—WIDW—appears with a coefficient that is significant at the conventional 95 percent confidence level. However, many of the other coefficients have t-statistics in the neighborhood of 1.5. It is therefore not surprising that the F-statistic for the hypothesis that all these coefficients equal zero is 2.09, which is significant at the 95 percent level of confidence. This joint hypothesis test indicates that some of the change in expectations observed between 1971 and 1973 is a response to the information contained in these variables.

It is also clear that the lack of statistical significance for a number of individual coefficients reflects imprecision rather than small point estimates. Several of the dummy variables have coefficients in the neighborhood of 600, which indicates that the event changes expectations by about 25 percent of its

Table 9.4 **Estimates of Responses to New Information**

	Equation			
Variable	1	2	3	4
Intercept	52.7	−163	55.1	55.1
	(81.4)	(116)	(60.2)	(61.0)
HBET	−633	−639	−38.2	−280
	(379)	(352)	(275)	(384)
HWOR	−604	−692	−322	−316
	(693)	(697)	(646)	(686)
WIDM	−659	−591	−313	−234
	(414)	(408)	(384)	(412)
WIDW	−804	−531	−181	−11.7
	(306)	(304)	(379)	(372)
LJOB	190	199	−46.7	−47.2
	(145)	(143)	(141)	(147)
GJOB	−602	−480	−444	−298
	(388)	(402)	(439)	(491)
CJOB	636	605	581	363
	(377)	(354)	(387)	(415)
NMOVE	239	264	191	253
	(170)	(169)	(166)	(183)
CW/100	−.442	−.450	−.307	−.307
	(.345)	(.337)	(.416)	(.383)
CCSS		.177	.441	.546
		(.0694)	(.0844)	(.114)

mean value. Nevertheless, standard errors are simply too large to say with confidence that the specific event (as opposed to events collectively) has an effect on expectations. Unfortunately, several coefficients also have counter-intuitive signs. Specifically, finding a job depresses expected benefits, while losing a job raises them.

It is particularly interesting to compare these results with the second column of table 9.4, which differs from the first only in that I have added CCSS (the change in statutory entitlement). Note first that the coefficient of this variable is statistically significant, which indicates that individuals do to some extent process information that affects their benefit levels through the benefit formulas. Furthermore, the addition of CCSS renders all other coefficients individually insignificant. Indeed, the F-statistic for the hypothesis that all these other coefficients equal zero is 1.71, which is significant at the 90 percent confidence level, but *not* at the 95 percent level. Closer inspection reveals that the introduction of CCSS renders the other coefficients jointly insignificant by reducing the estimated effects of several key variables (especially GJOB and WIDW) rather than by reducing the precision of these coefficients.

These results raise the interesting possibility that events affect expectations only through their effects on actual benefit calculations. This would entail a very high degree of rationality with respect to the processing of information received on the margin—certainly a much greater degree of rationality than was apparent in my analysis of the levels of expectations (see Bernheim 1988). Much of the following analysis is designed to investigate this possibility in greater detail.

The problem with the preceding set of estimates is, of course, that much of the observed changes in status may have been anticipated. This is especially important for the CCSS variable, in that statutory entitlements (what an individual would obtain on immediate retirement) rise steeply during the period immediately prior to retirement. Thus, much of the change in CCSS may have been anticipated. This would tend to bias the coefficient of CCSS toward zero, thereby overstating the extent to which other events affected expectations through channels other than the benefit formulas.

The next logical step is therefore to reestimate this specification using a more elaborate model for distinguishing between anticipated and unanticipated events. The object is to measure the component of an event that is unanticipated, given whatever method individuals actually use to forecast these events. Since we do know that individuals use information contained in ESS71, one possibility is to forecast (through regressions) the informational events on the basis of ESS71 and to use the residuals as measures of the unanticipated components.

Results based on this procedure appear in the third column of table 9.4. The list of independent variables should now be interpreted as measures of unanticipated changes, constructed as described above. Note that the coefficient of CCSS rises dramatically to .441 and its t-statistic now exceeds 5. In addition, the absolute value of *every other coefficient* declines, in some cases very significantly, and none of these other coefficients is even close to being statistically significant. Jointly, the significance of these other coefficients is no longer even marginal—the F-statistic for the hypothesis that they all equal zero is .65, which does not permit rejection at any meaningful level of confidence.

Even with this second procedure, measures of unanticipated events may still contain anticipated components. I therefore implement a third procedure in which the informational events are regressed on the full array of variables listed in section 9.2 as candidates for inclusion in Ω_t, as well as on ESS71. I then use the residuals from these variables as measures of the unanticipated changes. The justification for this procedure is that it is better to overexplain rather than underexplain the changes in status between 1971 and 1973. If one uses more information than do the respondents, then one's prediction will better than theirs, and the residual will then certainly reflect only unanticipated changes in status. Since the respondents' forecasts are presumed to be inferior, part of the predicted change will also be unanticipated. Fortunately, the nature of regression analysis is such that these other unanticipated components must be orthogonal to the residuals, and consequently the omission of these components will not bias the coefficients in a regression of CCSS on the residuals.

The results of this procedure appear in column 4 of table 9.4. The coefficient of CCSS again increases significantly to .546, and it remains highly significant. The absolute values of the coefficient estimates continue to decline significantly for WIDM, WIDW, GJOB, and CJOB. In fact, for WIDW, the coefficient is reduced practically to zero. In contrast, the coefficients for HBET and NMOVE rise somewhat. The statistical significance of these other individual coefficients continues to be low, and one cannot reject the joint hypothesis that they are all zero at any reasonable level of confidence.

Note that the second and third procedures described above implicitly treat the increase in average benefits between 1971 and 1973 as anticipated. Thus, the relative importance of CCSS does not simply reflect the fact that most individuals were aware of the benefit increase and adjusted their expectations accordingly. Rather, these results suggest that cross-sectional variation in unanticipated changes in statutory entitlements is the most important factor explaining cross-sectional variation in changes of expected benefits.

Two qualifications are in order. First, for the second and third procedures I have not adjusted the standard errors for the fact that the residuals are estimated rather than observed. It is in principle possible to obtain correct standard errors, as well as more efficient estimates, by estimating the entire system simultaneously through the use of seemingly-unrelated-regression (SUR) techniques. Unfortunately, computational requirements for SUR estimation of the full system exceeded the capacity of the available computer facilities. Second, the power of the tests discussed above is questionable in light of the fact that the standard errors of many coefficients are, from an economic point of view, extremely large.

Nevertheless, the general pattern of results, and especially the progression of coefficients through the second, third, and fourth columns in table 9.4, lends significant support to a remarkable conclusion: despite the fact that individuals do not appear to use all information contained in their statutory entitlements, the bulk of new, marginal information is incorporated into expectations through its effect on statutory entitlements. Although individuals do not appear to be well informed about the level of benefits, they appear to have a very good sense for how the benefit formulas operate on the margin.

The remaining question is whether these responses to new information are rational, in the sense that they closely resemble adjustments to an objective measure of expected benefits. Even if individuals incorporate new information as if they evaluate its effect on statutory entitlements, it is still possible that they do not fully exploit this information or that they misperceive the relation between entitlements and ultimate benefits. These issues are the subjects of the next subsection.

9.4.2 Evaluation of Response Quality

In order to test the rationality of responses to new information, it is necessary to add some additional structure to the basic model. I will suppose that the objective expectation concerning the realization of X is given by a linear function of information:

$$E(X|\Omega_t) = \Omega_t\zeta_1 .$$

When new information arrives, the objective expectation adjusts in response to unanticipated shocks. In particular, I suppose that

$$E(X|\Omega_{t+1}) = \Omega_t\zeta_1 + [\omega_{t+1} - E(\omega_{t+1}|\Omega_t)]\zeta_2 . \tag{9}$$

I now allow for the possibility that reported expectations differ from objective expectations. Suppose in particular that subjective expectations are given not by equation (9) but rather by

$$X^e_{t+1} = \Omega_t\theta_1 + [\omega_{t+1} - E(\omega_{t+1}|\Omega_t)]\theta_2 . \tag{10}$$

Then, combining (9) and (10), and using the fact that

$$X = E(X|\Omega_{t+1}) + \nu,$$

where ν is uncorrelated with Ω_{t+1}, we have

$$\begin{aligned} X - X^e_{t+1} = {} & \Omega_t(\zeta_1 - \theta_1) \\ & + [\omega_{t+1} - E(\omega_{t+1}|\Omega_t)](\zeta_2 - \theta_2) + \nu. \end{aligned} \tag{11}$$

The empirical analysis in Bernheim (1988) established that individuals do not process all available information in a fully rational manner (i.e., $\zeta_1 - \theta_1 \neq 0$). In this paper, I focus on the processing of new information (i.e., on the value of $\zeta_2 - \theta_2$).

To estimate the value of $\zeta_2 - \theta_2$, I regress the forecast error from expectations reported in 1973 on the 1971 information set as well as on measures of unanticipated events that occurred in the intervening period. I present estimates of equation (11) based on the three distinct methods of measuring unanticipated events discussed in the preceding subsection. It should be noted that the use of the first two procedures does not conform strictly to the theory outlined above.

Results appear in table 9.5. In order to conserve space, I have omitted coefficients for all the Ω_t variables and concentrate exclusively on the effects of new information. It is worth mentioning that the pattern of coefficients for the Ω_t variables was very similar to that obtained in my previous study. Most important, CSS71 entered with a positive, economically significant, and statistically significant coefficient, indicating that individuals fail to process all the information contained in statutory entitlements.

The first thing to notice about table 9.5 is that the results differ very little across the three procedures. There is a particularly striking similarity between the second and the third set of estimates. This should not be surprising—were

Table 9.5 **Estimates of Response Quality**

	Equation		
Variable	1	2	3
HBET	−28.0	33.5	34.1
	(688)	(643)	(655)
HWOR	473	493	496
	(334)	(347)	(342)
WIDM	9.95	54.8	50.5
	(232)	(246)	(243)
WIDW	−174	−101	−121
	(250)	(270)	(264)
LJOB	−205	−229	−228
	(139)	(137)	(137)
GJOB	325	357	348
	(373)	(378)	(376)
CJOB	−346	−352	−356
	(442)	(442)	(442)
NMOVE	−178	−186	−186
	(191)	(191)	(191)
CW/1000	6.68	6.91	6.87
	(3.64)	(3.62)	(3.63)
CCSS	.0141	.0654	.0532
	(.103)	(.122)	(.118)

it not for the presence of ESS71 in the first stage regressions, the independent variables would be related by a linear transformation, and the estimated coefficients for the new information variables would in fact be identical across procedures. The second and third sets of estimates differ only because the first-stage estimates for the coefficients of ESS71 differ.

Note that none of the variables in table 9.5 appears with a statistically significant coefficient in any regression (although the change in wealth variable, CW, does have t-statistics ranging from 1.8 to 1.9). In each case, one cannot reject the hypothesis that $\zeta_2 - \theta_2 = 0$ at any reasonable level of confidence. The data therefore support the view that individuals rationally process new information.

It is worth emphasizing that the CCSS variable has a statistically insignificant coefficient in each of these equations and that the point estimates of this coefficient are small in economic terms. Although individuals do not appear to use all information contained in statutory entitlements, they do seem to act rationally toward new information that changes statutory entitlements on the margin.

9.5 Conclusions

In this paper, I have outlined and tested a simple theory that describes the evolution of expectations concerning Social Security benefits during the preretirement period. While the raw data do not appear to support the empirically testable implications of this theory, the evidence indicates that this failure is attributable to the presence of measurement error. After correcting for the presence of this error, I find that expectations do appear to evolve as a random walk and that the innovations in this process are unrelated to previously available information.

After concluding that the data support the theory, I estimate responses of expectations to the arrival of new information and test for the rationality of these responses. The results here are striking. Although individuals do not form expectations on the basis of all available information, and in particular ignore much of the information contained in concurrent statutory entitlements to Social Security benefits, responses to new information during the period immediately preceding retirement appear to be highly rational. The bulk of information affects the evolution of expectations only through its effect on actual benefit calculations. Furthermore, the data support the view that individuals form accurate assessments of the ultimate effect of new information on actual benefits. These findings corroborate more speculative results from Bernheim (1988), which suggested that individuals formulate expectations about the retirement period much more carefully as retirement approaches.

References

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, 312–45. Palo Alto: Hoover Institution.

Bernheim, B. Douglas. 1989. The timing of retirement: A comparison of expectations and realizations. In *The economics of aging*, ed. D. Wise. Chicago: University of Chicago Press.

Burtless, Gary. 1986. Social Security, unanticipated benefit increases, and the timing of retirement. *Review of Economic Studies* 53:781–806.

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

Heckman, James J. 1976. The common structure of statistical models of truncation, sample selection, and limited dependent variables and a sample estimator for such models. *Annals of Economic and Social Measurement* 5:475–92.

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

Parnes, H. S., and G. Nestel. 1981. The retirement experience. In *Work and retirement*, ed. H. S. Parnes. Cambridge, Mass.: MIT Press.

Wolpin, Kenneth I., and Fürsun Gönül. 1987. On the use of expectations data in micro surveys: The case of retirement. Ohio State University. Mimeo.

Zarnowitz, Victor. 1984. Business cycle analysis and expectational survey data. NBER Working Paper no. 1378. Cambridge, Mass.: National Bureau of Economic Research.

Comment Sherwin Rosen

This paper continues Bernheim's imaginative use of rational modeling and modern conditional expectations apparatus to analyze Social Security benefit expectations of persons close to retirement. The orthogonality and related restrictions on the survey expectation data studied here are of interest, but forging the linkages between survey expectation responses (or lack thereof) and actual behavior will make the work even more important. Expectations of agents are important insofar as they affect savings, labor force participation, and other economic decisions of the elderly, but we do not yet know if there is any relation of that kind in these data.

The orthogonality restrictions that are sought here in some sense celebrate what old-fashioned empirical investigators would have called "poor" results.

Sherwin Rosen is the Edwin A. and Betty L. Bergman Professor of Economics at the University of Chicago and a research associate of the National Bureau of Economic Research.

Indeed, lack of correlation among variables in the retirement history sample is not exactly unknown, as some of the few surviving readers of my own work with these data can attest. The extent of measurement error that Bernheim finds in his own data gives one pause about the quality of the entire data set, so it remains to be seen whether Bernheim's results are found because the implicit behavioral restrictions of the theory really apply or because the data are not very good. Barring replication in a new survey, examining the behavioral linkages with real behavior is the only way that the research can most meaningfully be assessed. Perhaps the next paper in this sequence will be devoted to that important task.

Let me turn now to a few remarks about the work on its own terms. The data are incomplete because many people do not take the trouble to report expectational data at all. Bernheim's checks on the biases caused by this go about as far as possible given the data available to him. Nonetheless, in extrapolating these results to the population at large, it is well to keep in mind that people who are hooked up to a retirement information network such as AARP or a private pension system administered by a large company or who are just more interested in retiring in the near future would have had easy access to much of the information requested in the survey questions, but others would not. These are after all very specific questions about a legal entitlement that was changing very often over the survey period, and in very confusing ways as well, such as double indexing. Social Security does not send out financial statements unless specifically requested to do so, and most people must make substantial efforts to get the necessary information. Perhaps it is not too surprising that many people do not report an expectation given the costs of it.

One of Bernheim's most remarkable results is that the rational expectations work better for changes in benefits than for levels of benefits. The meaning of this is not entirely clear. Given the many legislated changes in the Social Security law during the sample period, it is not surprising that most people would be confused and uncertain about their benefit amounts at retirement. Benefits are calculated from a complicated table mapping earnings histories into a monthly benefit amount, and the table changed many times in these years. Saying that people have expectations of benefits means that they implicitly have expectations concerning the table parameters that take earnings histories to retirement payments.

If the expectations process is thought of in these terms, the only thing that could make much difference to the benefit calculation over time is changes in these parameters. After all, covered earnings histories do not change very much in a one- or two-year period. Most of the action in changes in monthly benefit amounts is year-to-year changes in the statutes, but virtually none of the information variables used in the statistical work relate directly to the benefit calculation or even indirectly in terms of congressional actions concerning it. This might account for why that kind of information shows no effects in the regression, and to that extent the result may not be so remarkable. Whatever that may be, it is difficult to see how these individuals could be confused about benefit levels but at the same time be much more rational in reacting to new information given the political environment of that time. Surely knowledge of benefit levels is as important to the behavioral effects of Social Security as are changes in benefits.

This Page Intentionally Left Blank