

NBER WORKING PAPER SERIES

RETIREMENT CONSUMPTION: INSIGHTS FROM A SURVEY

John Ameriks
Andrew Caplin
John Leahy

Working Paper 8735
<http://www.nber.org/papers/w8735>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2002

We thank Daron Acemoglu, Robert Barsky, Douglas Bernheim, Norman Boone, Peter Diamond, Harold Evensky, Douglas Fore, Jake Jacobson, Thomas Juster, Miles Kimball, David Laibson, Kevin Lang, Donghoon Lee, George Loewenstein, Mimi Lord, Annamaria Lusardi, Anna Madamba, Antonio Rangel, Alice Rivlin, Andrew Schotter, Matthew Shapiro, Robert Solow, George Szybillo, Richard Thaler, Tom Tyler, and Mark Warshawsky for their valuable comments. We gratefully acknowledge financial support for our survey provided by the TIAA-CREF Institute. All opinions expressed herein are those of the authors alone, and not necessarily those of TIAA-CREF or any of its employees. Caplin thanks the C.V. Starr Center at New York University and Leahy thanks the NSF for financial support. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

© 2002 by John Ameriks, Andrew Caplin and John Leahy. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Retirement Consumption: Insights from a Survey
John Ameriks, Andrew Caplin and John Leahy
NBER Working Paper No. 8735
January 2002
JEL No. E2, D1

ABSTRACT

Prior research has established that consumption falls significantly at retirement. What is not known is the extent to which this fall is anticipated during the working years. Do working households expect such a large fall in consumption upon retirement, or are they taken by surprise? Using data from a new survey, we show that many working households do expect a considerable fall in consumption when they retire. In fact, those who are already retired report significantly smaller falls in consumption than are expected by those who are still working. We show that participation in the stock market plays a dominant role in explaining the gap between expectations and outcomes, indicating that much of the gap is a result of unexpected stock market appreciation. The survey produces new insights into the high level of uncertainty in the period leading up to retirement, and the surprises that may lie in store when households actually retire.

John Ameriks
TIAA-CREF Institute
730 Third Avenue (730/24/01)
New York, NY 10017-3206
Tel: 212-916-4693
Fax: 212-916-6849
Email: jameriks@tiaa-cref.org
<http://www.tiaa-crefinstitute.org>

Andrew Caplin
Department of Economics
New York University
269 Mercer Street
New York, NY 10003
and NBER
Tel: 212-998-8950
Fax: 212-995-3932
Email: andrew.caplin@nyu.edu
<http://www.econ.nyu.edu/user/caplina/>

John Leahy
Department of Economics
Boston University
270 Bay State Road
Boston, MA 02215
and NBER
Tel: 617-353-6675
Fax: 617-353-4449
Email: jleahy@bu.edu
<http://econ.bu.edu/leahy/>

1 Introduction

The facts are not in dispute. Many households consume a great deal less after retirement than they did before (Hamermesh [1984], Mariger [1987]). Yet the interpretation is hotly disputed. Is this finding consistent with standard theories of saving? Banks, Blundell, and Tanner [1998] and Bernheim, Skinner, and Weinberg [1997] argue that it reflects a significant failure of forward-looking behavior: “many retirees...take stock of their finances only to discover that their resources are insufficient to maintain their accustomed standard of living.” In contrast, proponents of the classical life cycle model argue that retirement is merely a relatively inexpensive life phase, and that there is no failure of foresight.

Understanding retirement consumption is critical in tracking the twin processes of wealth creation and capital accumulation. It is also vital to the current debate on the adequacy of the U.S. savings rate. Yet research is hampered by the difficulty of measuring expectations. It is relatively easy to confirm that actual consumption falls in retirement, yet far harder to show that the drop in consumption is a deviation from prior expectations. The standard way to pin down expectations is to use data on actual consumption, together with strong identifying assumptions. The critical identifying assumption in Banks, *et. al* [1998] is that expected expenditure patterns among retirees are similar to actual patterns observed among workers who have been laid off. The key assumptions in Bernheim *et. al* [1997] concern which broad categories of realized expenditure are most strongly work-related. Unfortunately, it is difficult to test the validity of these identifying assumptions.

Given the ambiguity in the behavioral data, direct survey data on expectations may be of value. The systematic use of survey measures of expected consumption was proposed long ago by Juster [1966] in connection with demand for indivisible commodities. Yet there has been very little macroeconomic use of this form of expectational data.¹ Our own research on this subject represents a much-belated effort to adapt Juster’s ideas to the topic of aggregate consumption. We use data on expected and realized consumption at retirement from two recent surveys of some 2,000 TIAA-CREF participants. These data allow us to derive new insights on the determinants of retirement consumption. In addition, they open up some new questions concerning households’ financial preparedness

¹Barsky, Juster, Kimball, and Shapiro [1997] present a pioneering application of survey techniques to another important set of macroeconomic questions involving wealth accumulation.

for retirement.

We introduce the survey and the sample in section 2 of the paper. Section 3 presents our specific questions concerning the change in consumption at retirement. In section 4 we use our data to analyze the expected change in consumption at retirement among working households. We confirm that retirement is indeed seen as a relatively inexpensive life phase, as stressed by proponents of the life cycle model. In our sample, the typical household reports an expected fall in the order of 10%. For a household with a pattern of education, income, and wealth more in line with the general population, our results suggest an expected fall in the 20% range.

In section 5 we analyze how outcomes compare with these expectations. According to the wealth shock hypothesis, one might expect actual consumption to fall below prior expectations as retirees come face-to-face with their disappointing prospects. We find the opposite. Rather than having lower than expected consumption due to negative wealth shocks, many retired households appear to be consuming more than they had expected. For the retired households in our sample, the reported drop in consumption at retirement is significantly smaller than the expected drop, at around 2%.

In section 6 we explore the extent to which the high levels of retirement consumption are the result of stock market appreciation. The huge increase in equity values during the 1990s is an obvious candidate for a wealth shock that might have raised consumption among retired stock holders. By combining our expectational data with measures of stock ownership, we are able to confirm that most of the gap between expectation and outcome can indeed be accounted for in this way. The finding of a large effect of portfolio shares on reported consumption sheds light on another important macroeconomic issue: the marginal propensity to consume out of increases in equity values.

While retirement does not produce a negative wealth shock, our survey suggests that households may nevertheless be poorly prepared for retirement in other respects. We show in section 7 that many households enter retirement profoundly uncertain about how their change in work status will impact their consumption needs. When this uncertainty is resolved, we find that more households are surprised by how high are expenditure needs in retirement rather than by how low they are.

Does the systematic underestimation of retirement needs reflect an unpredictable shock (say unpredictably high health expenditures), or does it represent a bias in expectations? We find some intriguing evidence supporting the bias theory. In comparison with those whose survey responses indicated that they had done little retirement planning,

those who reported that they had spent a great deal of time planning were less biased toward underestimation of needs. This last result points out how important it will be to analyze the impact of planning activities on retirement preparation, as in Lusardi [1999] and Ameriks, Caplin, and Leahy [2002].

2 The Sample and the Survey

2.1 The SPF and the FAB

The data used in this paper are drawn primarily from two surveys sent to a sample of TIAA-CREF participants: the Survey of Participant Finances (SPF) conducted in January 2000, and the Survey of Financial Attitudes and Behavior (FAB) conducted in January 2001. The SPF was designed to examine in detail the type and the amount of financial assets owned by a sample of TIAA-CREF participants. The FAB explored these participants' financial preferences, expectations, and attitudes.

The sample for both surveys are members of the "TIAA-CREF Research Panel." In 1993, 60,000 TIAA-CREF participants were randomly selected as prospective members of this panel. They were mailed a letter requesting their participation in research-related surveys along with an initial survey containing questions regarding demographic and other characteristics. A total of 9,847 responses were received. Respondents were sent additional surveys in 1995, 1997, 1998, and 1999 prior to the SPF (2000) and the FAB (2001). At the end of each survey, the Research Panel was defined as the set of households who had responded to at least one prior survey, and who were still TIAA-CREF participants. The number of households in this panel changed over time according to the survey response rates and the rate of attrition among TIAA-CREF participants. Since response rates were in the 12%-21% range, the sample had to be replenished several times in order to maintain a universe of adequate size.

At the end of 1999, the Research Panel comprised some 9,234 households. The SPF was mailed to this group in January 2000. In total, 2,835 households responded to that survey, for an overall response rate of 30.7%. The universe for the FAB comprised 2,687 of the 2,835 households who had responded to the SPF; the difference being accounted for by changes of address, death, requests to be removed from the panel, etc. A response rate on the order of 30% or less would have been very damaging to our project, and considerable thought was given to issues of survey design in order to maximize response.

In particular, we followed procedures suggested by Dillman [1978] to boost the response rate. The end result was a remarkably high response rate. Fully 2,064 of the 2,687 households responded to the survey (76.8%).

2.2 The Sample: Demographic and Financial Characteristics

In order to compare retirement expectations with retirement outcomes, we divide our sample into two groups based on labor force participation. To preserve a clear distinction between retired and nonretired households in multi-person households, we define the household to be working only if no adult member is retired, and to be retired only if all adults are retired. Of the 2,064 respondents who filled out the FAB, 1,074 were in working households, while 735 were in retired households. The remaining 255 respondents were either in partially retired households, or provided insufficient data to be clearly categorized, and were therefore left out of the analysis.

Table 1 shows the basic demographic characteristics of sample households. We tabulate answers to questions on the respondent's: gender; marital status (married, never married, previously married); number of dependent children; and age. In addition to age, we tabulate the expected number of years to retirement, which is derived as the difference between the current age, and the best estimate to a question concerning the expected age at retirement. We also tabulate educational and occupational characteristics. The "other" employment category corresponds to secretarial, maintenance, and other support positions.

It is clear that our sample is extremely well-educated. Unlike those with Ph.D's, the group with high school education alone is not large enough to warrant a separate category: the lowest level of education we tabulate is college or below. In terms of employment, more than 1 in 3 are teaching faculty, while the majority of the others have management or professional positions. The gender, marital status, educational, and occupational characteristics of the retired households in our sample are very similar to those of the working households.

Table 1
Demographic Characteristics of 2001 Survey Respondents

Characteristic	Working Households		Retired Households		Total	
	(n)	(%)	(n)	(%)	(n)	(%)
Gender						
Female	480	44.7	334	45.4	814	45.0
Male	594	55.3	401	54.6	995	55.0
Marital Status						
Curr. married	690	64.6	465	63.4	1,155	64.1
Prev. married	179	16.8	182	24.8	361	20.0
Never married	199	18.6	87	11.9	286	15.9
Education						
College or below	281	26.2	247	33.6	528	29.2
Masters or Prof.	411	38.3	243	33.1	654	36.2
Ph.D.	382	35.6	245	33.3	627	34.7
Occupation						
Teaching faculty	382	35.6	325	44.5	707	39.2
Mgmt., Sen. Admn.	216	20.1	116	15.9	332	18.4
Other Tech./Prof.	267	24.9	120	16.4	387	21.5
Other	207	19.3	169	23.2	376	20.9
Age						
Below 35	112	10.5	0	0.0	112	6.2
35-44	209	19.6	0	0.0	209	11.6
45-54	382	35.8	10	1.4	392	21.8
55-64	272	25.5	117	16.0	389	21.6
65-74	84	7.9	388	52.9	472	26.2
75+	9	0.8	218	29.7	227	12.6
Years to/from retirement						
0-5	191	20.3	195	35.5	386	25.9
6-10	180	19.1	173	31.5	353	23.7
11-15	165	17.5	145	26.4	310	20.8
16-20	126	13.4	28	5.1	154	10.3
21+	281	29.8	8	1.5	289	19.4

Source: Authors' tabulations of 2001 FAB survey data.

Table 2 presents additional qualitative data on our sample, including information on home ownership, retiree health and long-term care insurance coverage, and defined benefit pension plan coverage. The table shows that a very high proportion of the sample own their homes rather than rent them. The majority of working owners have positive balances remaining on their mortgages, while the majority of retired owners have no remaining mortgage. With respect to health care, roughly 50% of households have at

least one member of the household with employer-provided health insurance extending into retirement, but very few have coverage for long term care. About a quarter of the respondents in working households, and a third of those in retired households, indicated that they participate in a defined benefit pension plan (of course, this pension plan is in addition to any TIAA-CREF and/or other defined contribution pension plans).

Table 2
Homes, Mortgages, Insurance, and DB Plans

Characteristic	Working		Retired		Total	
	(n)	(%)	(n)	(%)	(n)	(%)
Home ownership						
Rents	144	14.0	58	8.4	202	11.8
Owns	885	86.0	631	91.6	1,516	88.2
Mortgage status (owns)						
Has mortgage	726	82.0	244	38.7	970	64.0
Retiree health ins.						
Has retiree HI	538	51.2	400	55.9	938	53.1
LTC insurance						
Res. has LTC insurance	167	15.7	232	32.1	399	22.4
Sps. has LTC insurance	82	12.0	135	29.3	217	19.0
DB plans						
Res. has DB plan	253	24.4	221	32.6	474	27.7
Sps. has DB plan	177	17.1	132	19.5	309	18.0

Source: Authors' tabulations of 2000 and 2001 survey data.

Table 3 summarizes households' economic characteristics. Data on earnings is from questions on the FAB in which we asked households to provide estimates of their overall taxable income from employment in 1999. Most of the asset and debt information is drawn from the SPF. We record not only the total level of wealth, but also the division between retirement assets and non-retirement assets. Within the non-retirement assets, we separate out real estate wealth, which includes both owner-occupied and investment assets. With regard to debt, we distinguish between mortgage debt, and all other forms of debt, including credit card and educational debts. Note that Table 3 uses all of the data available on an item-by-item basis, so that sample sizes differ across rows of the table. Complete wealth data (i.e., a complete response to each and every question on the survey) was provided by only 55% of those who are working, and 40% of those who are retired.

Table 3
Wealth Aggregates for Surveyed Households

	Percentiles			
	Mean	10th	50th	90th
Sample and measure	(\$000)	(\$000)	(\$000)	(\$000)
Working households				
Earnings	87	30	73	150
Total assets	809	83	436	1,821
Total real estate assets	239	0	160	500
Total financial assets	587	38	273	1,453
Total ret. fin. assets	216	3	44	523
Total non-ret. fin. assets	406	21	192	1,035
Total debt	85	0	60	198
Total mortgage debt	74	0	50	180
Total personal debt	9	0	0	21
Total networth	737	48	368	1,761
Retired households				
Earnings	13	0	0	50
Total assets	1,092	241	812	2,213
Total real estate assets	276	30	180	500
Total financial assets	824	126	568	1,796
Total ret. fin. assets	357	6	149	892
Total non-ret. fin. assets	465	47	310	1,145
Total debt	40	0	0	92
Total mortgage debt	35	0	0	80
Total personal debt	4	0	0	5
Total networth	1,066	216	792	2,244

Source: Authors' tabulation of 2000 and 2001 survey data.

Notes: Earnings is the reported level of household taxable income from employment in the calendar year 1999. "Total financial assets" is the sum of all retirement account balances, mutual funds (except real estate mutual funds), directly held stocks, directly held bonds, checking accounts, savings accounts, and CDs. "Total assets" is total financial assets plus the value of homes and other real estate. "Total net worth" is total assets minus mortgage debt, outstanding educational loans, outstanding personal loans, and credit card balances. All aggregates exclude the value of real estate mutual funds, whole life insurance policies, trusts, and educational savings accounts (Education IRAs and 529 plans). All wealth data are as of December 31, 1999.

In Table 3 we use net worth as our overall summary statistic for the household's balance sheet situation. As one would expect, retired households have generally accumulated more wealth than working households. Going beyond this, the most striking features of the wealth data are the generally high means and the relatively low variability. We discuss these features in more detail when we compare our sample with a representative sample in the next section.

With respect to earnings, the table shows that mean household income from employment among working households is \$87,000, while the median is \$73,000. However given that our fundamental interest is in the lifetime budget set, we are more interested in the present value of future household income from employment rather than income in any given year. Fortunately, our sample contains many data items of potential value in computing this present value. We have estimates of the remaining years before the respondent retires. In addition, we asked households not only for their actual household income from employment in the years 1998, 1999, and 2000, but also for their estimated future incomes in 2001, 2005, and 2010. Our procedure for using these data to arrive at a final estimate of the present value of future employment income is detailed in Appendix 1.

2.3 A Comparison with the SCF

In what ways does our sample differ from a random sample of U.S. households? Table 4 records financial characteristics of households in the Survey of Consumer Finances, which is a representative sample of U.S. households.

Our sample clearly possesses far greater wealth than the representative group of households in the in the SCF. Net worth among both our working and retired households is some 2.5-3 times higher than in the SCF. In addition, there appears to be far greater homogeneity in our sample than in the SCF. As one reflection of this, median net worth in our sample is some 5-6 times as high in the SCF. A particularly striking difference between our sample and the SCF concerns the distribution of financial assets. In our sample, mean financial assets are \$587,000, with a median of \$273,000. In the SCF, the mean is smaller by a factor of 5 at \$124,000, while the median is smaller by a factor of 18, at \$15,000.

In addition to differences in net worth, there are important differences in composition. Average debt levels are generally lower in our sample than in the SCF. The difference is accounted for by the fact that there are a significant number of households in the SCF with very high levels of personal debt, while in our sample the level of personal debt is almost universally low.

The fact that our sample of working households have generally high levels of non-retirement financial assets and low debt levels implies that very few are liquidity constrained in the sense of being “hand-to-mouth” consumers. In standard treatments, households with levels of non-retirement financial assets below \$10,000 are treated as

liquidity constrained. If we use this criterion, well under 10% of our sample are liquidity constrained.

Table 4
Wealth Aggregates in 1998 Survey of Consumer Finances

Sample and measure	Mean	Percentiles		
		10th	50th	90th
	(\$000)	(\$000)	(\$000)	(\$000)
Working households				
Income	56	9	37	97
Total assets	329	3	110	566
Total real estate assets	33	0	0	42
Total financial assets	124	0	15	230
Total ret. fin. assets	42	0	1	100
Total non-ret. fin. assets	82	0	6	120
Total debt	59	0	22	142
Total mortgage debt	47	0	0	125
Total personal debt	12	0	4	29
Total networth	270	0	58	455
Retired households				
Income	37	7	22	70
Total assets	383	6	155	678
Total real estate assets	46	0	0	110
Total financial assets	198	0	39	347
Total ret. fin. assets	35	0	0	75
Total non-ret. fin. assets	163	0	30	294
Total debt	18	0	0	50
Total mortgage debt	13	0	0	38
Total personal debt	6	0	0	8
Total networth	364	5	147	656

Source: Authors' tabulation of 1998 SCF data.
Note: Income includes non-labor income.

3 Spending in Retirement: The Question

3.1 The Questions

Questions 16a and 16b were the fundamental questions on retirement consumption:

- 16a: For many households, overall spending changes dramatically upon retirement. Please indicate below (what) your experience has been (if you are retired), or what your expectations are (if not retired)

- My household had (or expects to have) no change in spending at retirement
 - My household has spent (or will spend) more after retirement than before (also answer 16b)
 - My household has spent (or will spend) less after retirement than before (also answer 16b).
- 16b: About how much more or less (as a percentage of your annual pre-retirement spending)?

We asked for three distinct estimates in response to 16b: a low estimate, a closest estimate, and a high estimate. The precise meaning of these three estimates was clarified earlier in the survey:

- Because of the uncertainty that is inherent in any question about your financial future, the following questions ask you to provide three estimates:
 - a **LOW** estimate, where you are 90% sure that the correct answer is **above** this number
 - your **CLOSEST** estimate of the most accurate answer
 - a **HIGH** estimate, where you are 90% sure the correct answer is **below** this number.

There are two entirely different approaches one can take when posing questions on the level of consumption. Our approach was to ask the question in a manner that accords with the “common-sense” notion of consumption as directly tied to expenditures. The obvious alternative involves making a strong effort to convey to the respondent the precise economic meaning of consumption expenditures, so that they make appropriate adjustments. It was our belief that the former approach would be easier for households to understand, especially given our request for low, closest, and high quantitative estimates of the change at retirement. This may have contributed to the high response rates. Questions 16a and 16b were answered by the vast majority of respondents. 90% of the sample answered question 16a, and of those who were eligible, 90% answered the closest estimate part of question 16b. Even the high and low estimates were provided by almost 70% of those eligible to answer.

The potential drawback of our common-sense approach is that it exposes us to potential ambiguity. One important question is whether or not respondents include debt repayments (especially mortgage payments) as part of their overall spending. Fortunately, we have a great deal of data on households' levels of debt which we can use to assess the quantitative importance of this ambiguity in determining the answers to our questions. As we will see in the next section, the results are surprisingly clean. It appears that many households correctly interpreted question 16a as referring to changes in expenses associated with the event of retirement itself, discounting any change in debt payments as irrelevant. More generally, we find that the impact of potential ambiguities on the quantitative answers to question 16 to be rather low.

3.2 Imputations and Corrections

The quantitative answers to question 16b require interpretation. One issue is imputation: what numerical value should be filled in for those who report no expected change in consumption at retirement, and for those who report an expected change but do not fill in a numerical value? A second issue is the relatively large number of households who report expected falls in consumption of either 80% or 90%. Not only do such huge falls in consumption seem a priori unlikely, but there were very few households in the 51%-69% range, suggesting even more strongly that there is something artificial in the second mode of the distribution.

Through the body of the paper we use simple and natural procedures to adjust the data. For those who report no expected change in consumption at retirement, we impute the numerical answer of a zero expected change. For those who report an expected increase but provide no numerical estimate, we use the conditional average among those giving the same qualitative answer. Finally, we assume that all households who fill in an expected fall in consumption of more than 50% misread the question, filling in the more familiar "replacement rate" (the ratio of retirement consumption to pre-retirement consumption) rather than answering the question about the change per se.

In Appendix 2 we detail a more intricate procedure for going from the survey answers to the underlying distribution of answers to question 16b. This approach involves using a maximum likelihood procedure based on the assumption of lognormality to correct for all reporting errors, including a possible bias associated with the rounding process. Given our identifying assumptions, we can then estimate the distribution of the "correct"

answers to the survey. Fortunately, there is very little quantitative difference between the simple procedures used in the paper and the more intricate procedure described in the Appendix. In particular the complete maximum likelihood procedure strongly suggests that all of the data lying at 50 and above are the result of reporting error.

3.3 What Window around Retirement?

The retirement expectations of very young households may be driven by entirely different forces than those of older households. In fact, our data suggest that younger households expect significantly smaller falls in consumption at retirement than do the remaining households. One possible explanation may be that younger households are naive (in the sense of O’Donoghue and Rabin [1999]) about their ability to save for their old age. A second possible explanation involves liquidity constraints. When these are currently binding, households may look forward to retirement as a time of relatively high consumption in which these constraints have been removed. Whatever the explanation, it is not clear that the expectations of very young households should have any weight in our analysis.

Just as we need to take care in our specification of the working sample, so we need to be aware that households who live for unusually long periods following retirement may answer question 16 in a qualitatively different manner than do recent retirees. Again, it is not clear that the answers of the super-aged are relevant, since our fundamental interest concerns a change associated with the event of retirement itself, rather than the aging process per se. To limit the impact of youth and old age in distorting our results, we set a 20 year pre-retirement window and a 10 year post-retirement window for inclusion in the regression samples. In addition we exclude from the retired sample those who have been retired for less than one year, on the grounds that they may have insufficient experience of the retired style of life to have adequately adjusted expectations.

4 Expectations of Working Households

4.1 Basic Data

Table 5 presents simple summary statistics of the answers to questions 16a and the “best estimate” part of question 16b among working households. We present data both for the entire working population, and for the “regression sample” that satisfies the

two conditions (no more than 20 years to retirement and adequate data) required to be included in the regression analysis. In these tabulations, and throughout the rest of the paper, we have converted the reported percentages into corresponding implied changes in the natural logarithm of consumption (this produces a symmetric treatment of consumption increases and decreases, and is more consistent with measures used in other studies).

Table 5
Spending Expectations among the Working

	Fraction	Mean	Std. Dev.
Sample and expectation	(%)	($\Delta \ln c$)	($\Delta \ln c$)
All respondents (n=1,005)			
Expects same	35.2	0.000	0.000
Expects lower	54.7	-0.255	0.139
Expects higher	10.0	0.199	0.121
Regression sample (n=313)			
Expects same	35.7	0.000	0.000
Expects lower	56.8	-0.255	0.134
Expects higher	7.5	0.194	0.122

Source: Authors' calculations based on 2001 survey data.

Note: Respondents were asked to provide data on the expected percentage change in spending at retirement; we have converted these responses to implied changes in the natural logarithm of consumption.

It is clear that the majority of households expect consumption in retirement to fall. In the regression sample, more than 55% expect a fall in consumption, while less than 8% expect an increase. A second striking fact is that among those who expect consumption to fall, the average expected fall is in the order of 25%. More generally there appears to be massive individual heterogeneity. Even though only 8% of sample households expect an increase in consumption at retirement, the mean among these is an expected increase of close to 20%.

4.2 Regression Results

Many economic and demographic forces may influence the expected level of retirement in consumption. Table 6 summarizes our basic regression results for the level of expected retirement consumption among sample households. The dependent variable in this regression is the expected change in consumption (expressed as a natural logarithm). The

right hand side variables include demographic dummies based on the respondent's gender and marital status, as well as other variables tracking the number of dependent children, and the respondent's age. We include also indicators of education and occupation, as well as home ownership and health insurance status. The financial variables included are net worth, the present value of future income from employment, and indicators concerning defined benefit pension coverage.

Table 6
OLS Regression of Expected Spending Change at Retirement, for Working Households

RHS Variable	Coeff.	Std. Err.	Pr > t
ln(Net worth)	0.0555**	0.0119	0.000
ln(Fut. earnings)	-0.0055	0.0135	0.684
Male	-0.0036	0.0220	0.869
Prev. married	0.0316	0.0301	0.294
Nev. married	0.0531	0.0371	0.153
Num. kids	-0.0054	0.0118	0.651
Age	-0.0031	0.0021	0.140
Coll. or below	0.0546*	0.0287	0.058
Prof. degree	0.0031	0.0234	0.895
Faculty	0.0396*	0.0227	0.082
Owns home	0.0101	0.0393	0.797
R. DB plan	-0.0420*	0.0245	0.088
S. DB plan	0.0009	0.0290	0.976
Ret. HI ins.	-0.0208	0.0203	0.308
R. LTC ins.	0.0038	0.0383	0.921
S. LTC ins.	0.0088	0.0484	0.856
Constant	-0.3066*	0.1716	0.075

Source: Authors' tabulation of 2001 survey data.

Notes: The dependent variable is our measure of expected change in the log of consumption, as described in the text. There were 331 observations used in this regression. The R^2 was 0.1078, and $F(16, 314)$ was 2.37.

The most striking finding is the strongly significant positive coefficient on net worth. The coefficient suggests that a doubling of net worth (an increase of one log point) is associated with a 5-6% increase in expected consumption at retirement. The importance of net worth is especially striking given the apparent irrelevance of future income from employment. How can these two variables, which should have similar impacts in terms of shifting the lifetime budget set, have such different implications for retirement consumption? The most natural explanation would seem to be based on differences in the

desire to save for retirement. If households differ in their preference for consumption in retirement relative to consumption in the working years, then those who desire a higher relative level of retirement consumption will be accumulating more wealth in order to pay for it. But unless this taste parameter is correlated with the overall net present value of future earnings, there is no reason to expect a relationship between the value of human capital and expected retirement consumption.

Aside from wealth, there are no other variables that are significant at the 5% level. Lack of continuing employer-sponsored health insurance, lack of long term care insurance, and ownership of a defined benefit pension might all be expected to raise expected spending in retirement, yet these effects are apparently not statistically significant. The lack of impact of home ownership on expected retirement consumption is also striking, since our question may have been interpreted differently by home owners than by renters. To a first approximation, for a renter whose expenditure includes payments for housing services, the answer should represent the change in *total* consumption, housing included. In contrast owners' answers should measure the expected change in *non-housing* consumption, since there is no payment that corresponds to the implicit flow of consumption derived from the house. Apparently the expected fall in consumption at retirement is very similar, regardless of the definition used.

4.3 Is Net Worth Enough?

Do all assets and debts enter equally into the determination of retirement consumption? One obvious question concerns whether assets and debts should be netted out, or instead treated as having potentially different impacts on retirement expectations. A second question is whether or not retirement and non-retirement assets should be treated symmetrically. A third is whether residential real estate represents an entirely separate asset category due to its illiquidity. Finally, on the debt side of the balance sheet, it is natural to wonder whether or not mortgage debt has the same relationship with expected retirement consumption as do other forms of personal debt, such as credit card debt.

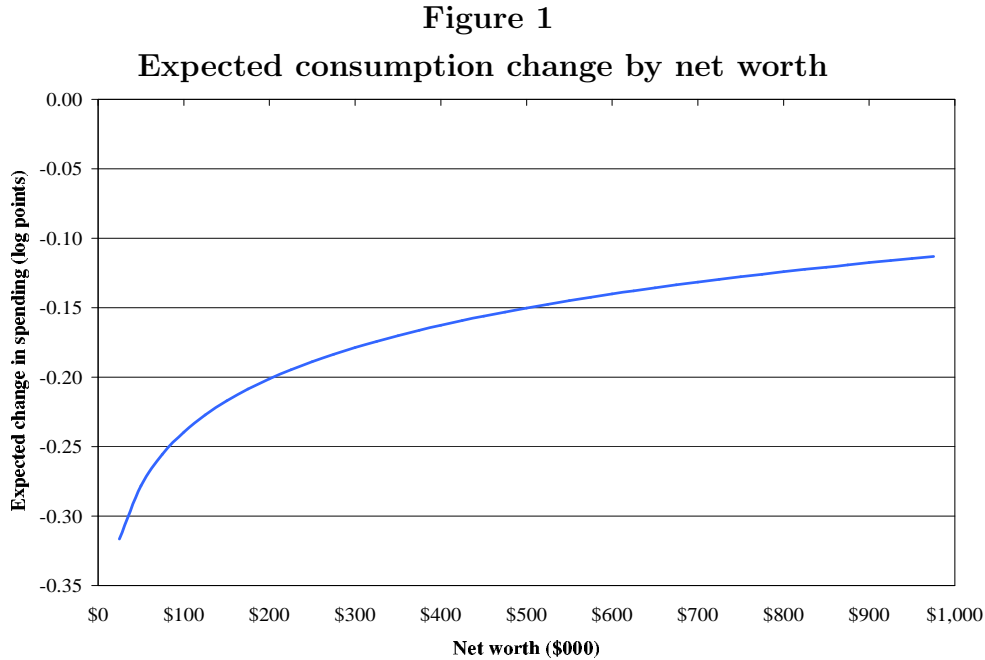
To check whether or not our data suggest a separate role for any of these decompositions of net worth, we have run a series of regressions in which we include measures of the share of net worth comprised of each subcategory of financial instrument (shares were constructed for retirement assets, real estate assets, mortgage debt and personal debt). We estimated specifications including each subcategory by itself (in addition to

the other variables listed in Table 6), and in combination with all the others. We found that *none* of the subcategories had statistical significance at the 10% level or less, either singly or jointly. Apparently the information contained in the net worth figure conveys all that is needed to understand the level of expected retirement consumption.

The apparent independence of expected retirement consumption from the composition of net worth is striking. In part, it may be related to special characteristics of the households in our sample. For example the irrelevance of the retirement asset share may be connected to the relatively small number of households in the sample for whom liquidity constraints are binding. In similar manner, the irrelevance of the composition of debts may be a result of our sample households having relatively low levels of credit card debt, which may in turn be associated with their having relatively high levels of self control. The irrelevance of the share in real estate assets may be related to financial sophistication: it suggests that many households in the sample view their homes at least in part as assets that are potentially available for consumption in retirement. Finally, the irrelevance of aggregate debt levels is suggestive of an entirely different form of sophistication: it suggests that most respondents interpreted question 16a just as we had intended, treating as irrelevant changes in debt service payments not specifically related to retirement.

4.4 Expectations at Retirement

We are particularly interested in using our regression analysis to characterize the expected level of retirement consumption for a household on the verge of retirement. We consider a married couple with no dependent children, with a male respondent of age 65 who is a non-faculty member with a masters of professional degree. The household owns their own home, has no employer-provided health insurance in retirement, no defined benefit plan, and no long term care insurance. Since the household is at the point of retirement, they have zero future income from employment. In figure 1 we plot the expected change in consumption at retirement implied by our regression as a function of the level of wealth at retirement.



One point to note is that this expectation is somewhat lower than the raw numbers describing our sample average. The raw data produce an average expected fall in the order of 10%, yet figure 1 indicates that a typical household on the verge of retirement with net worth of \$500,000 expects a fall of some 15%. The two primary factors accounting for the expectation of an additional 5% fall in consumption are that figure 1 is drawn for a non-faculty member who is considerably older than average working household (there is a small negative coefficient on age).

Figure 1 goes some way to reconciling our finding with the earlier literature in which the fall in retirement consumption has been far larger than in our sample. At a net worth of \$200,000, figure 1 suggests that the expected fall in retirement consumption would be in the order of 20%, which is close to the numbers in the prior literature.

5 Outcomes among the retired

5.1 Basic Data

Table 7 presents simple summary statistics of the answers to questions 16a and to the best estimate part of 16b among retired households. As with the working population, we present data both for the entire retired population, and for the regression sample.

Table 7
Spending experience among the retired

Sample and experience	Fraction (%)	Mean ($\Delta \ln c$)	Std. Dev. ($\Delta \ln c$)
All respondents (n=661)			
Experienced same	43.9	0.000	0.000
Experienced lower	36.2	-0.257	0.135
Experienced higher	20.0	0.196	0.101
Regression sample (n=140)			
Experienced same	47.1	0.000	0.000
Experienced lower	30.7	-0.233	0.142
Experienced higher	22.1	0.175	0.123

Source: Authors' calculations based on 2001 survey data.

The main finding is immediately obvious: recently retired households have generally had smaller falls in consumption than is anticipated by those who are still working. While more than 55% of working households in our sample expect consumption to fall, only 30% of the retired reported that they had experienced such a fall. On the other side of this, while less than 8% of working households in our sample expect consumption to rise, more than 20% of the retired households experienced such a rise.

5.2 Regression Analysis

Table 8 summarizes our basic regression results for the determinants of actual retirement consumption for our sample of retired households. The regression is identical to that for working households in section 4.2 above, with the exception that we exclude income from employment, the number of dependent children, and home ownership, since they show insufficient variation. As with the working sample, net worth is the only variable that is significant at the 5% level. No other variables are significant even at the 10% level.

Table 8
Regression Results for Retired Households

Variable	Coeff.	Std. Err.	Pr > t
ln(Net worth)	0.0562**	0.0227	0.015
Male	-0.0130	0.0360	0.718
Prev. married	0.0649	0.0447	0.148
Nev. married	0.0248	0.0557	0.657
Age	0.0006	0.0031	0.850
Coll. or below	-0.0364	0.0422	0.390
Prof. degree	0.0474	0.0430	0.272
Faculty	0.0014	0.0387	0.971
R. DB plan	0.0019	0.0369	0.960
S. DB plan	0.0251	0.0429	0.560
Ret. HI ins.	0.0065	0.0332	0.845
R. LTC ins.	-0.0135	0.0491	0.784
S. LTC ins.	0.0537	0.0640	0.403
Constant	-0.4767*	0.2656	0.075

Source: Authors' tabulation of 2001 survey data.

Notes: The dependent variable is our measure of the realized change in the log of consumption, as described in the text. There were 140 observations used in this regression. The R^2 was 0.1300, and $F(13, 126)$ was 1.45.

5.3 The Gap

Comparing regression results it seems clear that recent retirees experience a smaller fall in consumption than expected by working households who are close to retirement. To provide simple statistical confirmation, we run a pooled regression in which we add a dummy for whether or not the household is retired, and allow the coefficient on wealth to depend on whether or not the household is retired. We exclude household income, the number of children, and homeownership status from this regression, since they are absent from the regression for retired households, and irrelevant to the regression for working households. Table 9 summarizes regression results.

Table 9
Regression Results for All Households

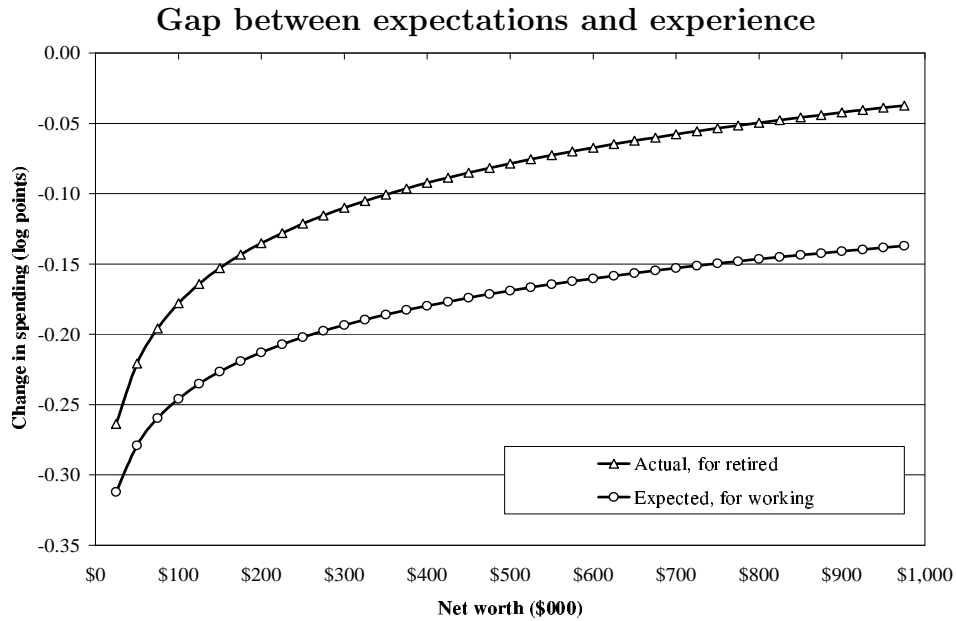
Variable	Coeff.	Std. Err.	Pr > t
Retired	0.0035	0.1513	0.981
ln(Net worth)	0.0478**	0.0105	0.000
ln(Net worth)*r	0.0140	0.0221	0.527
Male	-0.0025	0.0183	0.893
Prev. married	0.0421*	0.0241	0.081
Nev. married	0.0476*	0.0286	0.097
Age	-0.0017	0.0013	0.209
Coll. or below	0.0274	0.0230	0.234
Prof. degree	0.0189	0.0199	0.343
Faculty	0.0359*	0.0193	0.063
R. DB plan	-0.0294	0.0202	0.145
S. DB plan	0.0153	0.0237	0.520
Ret. HI ins.	-0.0119	0.0170	0.484
R. LTC ins.	0.0000	0.0294	1.000
S. LTC ins.	0.0370	0.0375	0.325
Constant	-0.3724**	0.0834	0.000

Source: Authors' tabulation of 2001 survey data.

Notes: The dependent variable is our measure of the expected (for working households) or realized (for retired households) change in the log of consumption, as described in the text. There were 471 observations used in this regression. The R^2 was 0.1419, and $F(15,455)$ was 5.02. A test of the joint significance of the retired dummy and the net worth / retired interaction variable has an $F(2,455)$ of 7.38.

Note that both the retirement dummy and the retirement wealth dummy are positive, and the test for their joint significance has an F-statistic of 7.38. The economic importance of the retirement effect is illustrated in figure 2. We plot the model-fitted expected and actual retirement consumption for a married couple with no dependent children, with a male respondent of age 65 who is a non-faculty member with a masters of professional degree, and with no employer-provided health insurance in retirement, no defined benefit plan, and no long term care insurance. As the curve shows, at a net worth level of \$500,000, the expected level of retirement consumption falls a full 10% below the realized level for this same household.

Figure 2



6 Wealth Shocks

6.1 Equity Shares

The potential for increases in equity values during the 1990's to explain the high level of consumption in our retired sample is highlighted by the fact that equity shares are so large. Our data on household portfolio shares comes from two primary sources. In addition to the SPF, we have gone directly to TIAA-CREF data to extract a precise value for the TIAA-CREF holdings of all respondents as of the survey date of this wealth survey. We replace the survey data with this precise measure in arriving at our estimate of the proportion of equities in the household's entire portfolio of financial assets. Table 10 characterizes these equity holdings for the retired and the working households in the regression samples.

Table 10
Stockownership among Sample Members

Fraction of net worth in stocks	Working (%)	Retired (%)
Less than 10%	6.7	19.1
11%-25%	10.4	11.1
26%-50%	31.8	35.7
51%-75%	36.1	31.0
76%-90%	11.7	3.2
91% or more	3.3	0.0

Source: Authors' tabulation of 2001 survey data.

Equity holdings are high in our sample in comparison to the population as a whole; more than two thirds of both our working and retired samples have at least 25% of their household's net worth invested in stocks. Very few working households have no equity; however, 19.1% of the retired respondents indicate that less than 10% of their net worth is held in the form of stocks. The median working household has around 40% of financial wealth in equity, while the median retired household has somewhat less.

6.2 Adding Equities

We use a simple regression to confirm that equities appear to play a large part in explaining the high realized level of retirement consumption. Table 11 summarizes results of a pooled regression in which we add the equity shares as a fraction of net worth to our pooled regression of the last section. Given our existing findings concerning the irrelevance of the composition of net worth for working households, we expect a similar irrelevance to hold for the share of equities. However we also add a dummy to capture the possible increased impact of equities in retirement. Here we expect an asymmetry. A household with a high share of equities would likely have been impacted by a far larger wealth shock than would a household with relatively low equity levels. Hence we expect a positive coefficient on the equity share in retirement.

Table 11
Regression Results for All Households,
Controlling for Stock Market Exposure

Variable	Coeff.	Std. Err.	Pr > t
Retired	0.0278	0.1633	0.865
ln(Net worth)	0.0440**	0.0111	0.000
ln(Net worth)*r	-0.0017	0.0250	0.945
Equity share	0.0001	0.0004	0.866
(Equity share)*r	0.0022**	0.0009	0.011
Male	-0.0066	0.0192	0.731
Prev. married	0.0415*	0.0248	0.096
Nev. married	0.0473	0.0303	0.119
Age	-0.0016	0.0014	0.257
Coll. or below	0.0220	0.0239	0.356
Prof. degree	0.0143	0.0208	0.494
Faculty	0.0458**	0.0203	0.024
R. DB plan	-0.0402*	0.0215	0.062
S. DB plan	0.0214	0.0251	0.394
Ret. HI ins.	-0.0101	0.0181	0.576
R. LTC ins.	0.0059	0.0302	0.845
S. LTC ins.	0.0329	0.0391	0.401
Constant	-0.3538**	0.0897	0.000

Source: Authors' tabulation of 2001 survey data.

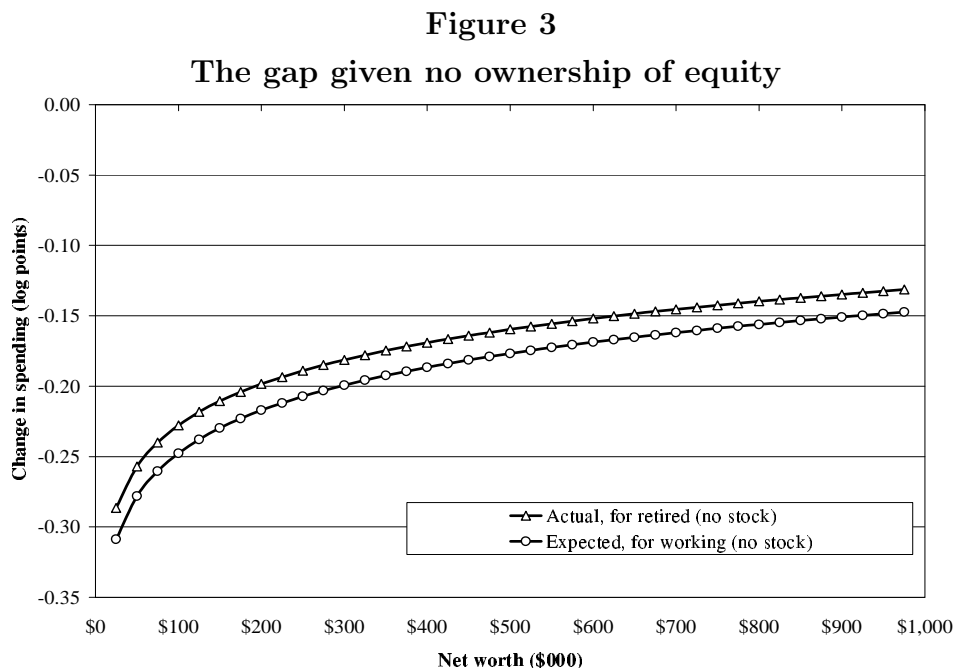
Notes: The dependent variable is our measure of the expected (for working households) or realized (for retired households) change in the log of consumption, as described in the text. There were 425 observations used in this regression. The R^2 was 0.1645, and $F(17, 407)$ was 4.71. A test of the joint significance of the retired dummy and the net worth/retired interaction variable has an $F(2, 407)$ of 0.09.

As hypothesized, we find that the equity effect is entirely irrelevant for working households, yet has a significant positive impact on actual consumption among retired households. Once the equity effect is removed, there appears to be no significant impact of retirement on the level of consumption. In fact the test for the joint significance of the retirement dummy and the retirement wealth coefficient has an F-statistic of 0.09.

6.3 Equities and the Gap

Figure 3 indicates that equity effects play a very large role in explaining the gap between expectations and outcomes. The figure shows our standard 65 year-old household of the last section where we condition not only on wealth, but also on the share of wealth in equity. The particular comparison drawn is for a household with no equities. As the

figure shows, a household with \$500,000 in wealth but no equity expects consumption to fall some 18% in retirement, while the realized outcome is closer to 16%, an insignificant difference.



7 Planning and Surprise

The above findings suggest that households' expectations of retirement consumption are not systematically in error. Yet it would be a mistake to conclude that households were therefore entirely well-prepared for their retirements. In this section we consider two separate indications of lack of preparedness. We begin by analyzing the subjective uncertainty of households in the pre-retirement period, showing the large amount of uncertainty with which retirement is approached. We follow up by exploring the ex post impressions of retired households concerning the types of cost surprises that they encountered in retirement, and show that the nature of this surprise is influenced by planning behavior.

7.1 Subjective Uncertainty

Question 16b was designed to shed light not only on expected consumption at retirement, but also on the degree of subjective uncertainty. Table 12 produces basic facts about the range of uncertainty, which we define as the absolute difference between the answers to

the high and low estimates in response to question 16b. In terms of the subjective prior, the range of uncertainty is defined as the central eight deciles of the distribution. In addition to producing unconditional averages, table 12 characterizes the impact of years to retirement on the range of uncertainty.

Table 12
Subjective Uncertainty in Expected Changes in Retirement Spending for Working Households, Reported Range from 10th to 90th Percentile

	Count	Mean	S.D.
Category	(n)	(%)	(%)
All Working	533	27.3	22.6
Years to retirement			
0-5	89	28.3	24.2
6-10	95	27.3	27.5
11-15	98	28.1	21.5
16-20	71	28.6	23.7
21+	153	26.8	20.0

Source: Authors' tabulation of 2001 survey data.

The range of uncertainty is very large, amounting to over 25% of pre-retirement consumption. It is also striking that the number of years to retirement has no systematic impact on the range of uncertainty. One might have expected less uncertainty for those who are closer to retirement, due both to reduced wealth uncertainty and to a possibly increased understanding of what is in store for them in retirement.

One other factor that one might expect to impact the range of uncertainty is the wealth portfolio, with less uncertainty for those who hold less risky portfolios. A simple regression of the range of uncertainty on portfolio shares for those close to retirement reveals that there is indeed such a relationship. However, the best estimate of the residual uncertainty for a household with no risky assets remains well above 20%.

High as they are, there is reason to believe that the numbers provided above may understate the true level of uncertainty. The question as posed suggests that all three answers (high, best, and low) should lie on the same side of zero. This makes it a very hard question to answer for a household expecting a 10% fall, but believing that a 5% or more rise has probability 1/10. The household might reasonably feel that the correct answer is 0%. They may even interpret the question to be asking about the top decile of their subjective distribution conditional on consumption falling, in which case they

should provide a strictly positive answer (“if our consumption at retirement falls, there is a 90% chance that the fall will be larger than 3%”). In either case, the bias operates to artificially shrink the stated range.

7.2 Surprise

One implication of high pre-retirement uncertainty is that precautionary savings may be important not only in the early period of life when liquidity constraints are important, but also later in life before retirement uncertainty is resolved. For present purposes, the more crucial issue raised by the high level of prior uncertainty is the question of what happens when the uncertainty is resolved. Question 17 was designed to focus directly on the qualitative issue of whether or not there were expense surprises in retirement:

- 17a. If you are currently retired, how do your spending needs in retirement compare with those you expected before you retired:
 - About what you expected
 - Higher than you expected
 - Lower than you expected

Table 13 presents the basic data. The dominant feature is the large number who report not having been surprised. The second feature is the apparent, if small, bias in the direction of surprise. More households report being surprised by how high were their expenses in retirement rather than by how low they were.

Table 13
Actual spending in retirement relative to
prior expectations, among retired households

	All responses (%, n=694)	Regression sample (%, n=604)
Spending level		
As expected	74.9	74.8
Higher than expected	15.3	15.2
Lower than expected	9.8	9.9

Source: Authors' tabulation of 2001 survey data.

7.3 Planning and Surprise

Are mistakes in predicting retirement consumption due to some common shock (e.g. cost shocks for health expenditures), or are they due to poor planning? We are in a good position to address this issue, since many questions on the survey were focused on the extent of household planning activities. The most straightforward summary measure of planning activity is a question that is entirely qualitative. We asked respondents to indicate which of six statements (disagree strongly, disagree, disagree somewhat, agree somewhat, agree, and agree strongly) best characterized their attitude toward the assertion in question:

- Question 1a: I have spent a great deal of time developing a financial plan.

Table 14 shows that self-reported planning activity has a strong and predictable impact on the answer to question 17a.

Table 14
Planning and Reported Spending Surprise

Planning response	(n)	Response to Question 17a, Needs Relative to Prior Expectations		
		At (%)	Higher (%)	Lower (%)
Disagree (1-2)	100	67.0	24.0	9.0
Equivocate (3-4)	328	76.8	14.9	8.2
Agree (5-6)	279	76.0	12.5	11.5

Source: Authors' tabulation of 2001 survey data.

The striking finding concerns the asymmetry. For those who have planned the most, the asymmetry is entirely absent. For those at intermediate levels of planning, the asymmetry matches its level in the broader population. For those who have done the least planning, the asymmetry is significantly more marked than it is in the broader population.

Does the relationship between planning and surprise survive in a multivariate context? Table 15 summarizes results of an ordered probit analysis of the answers to question 17a on planning activity (question 1a), insurance status, and demographic controls. The only significant variable is planning activity, which is significant at the 3% level. Note that the regression we report is for all retired households: if we restrict attention to those 1-10 years from retirement, the coefficient on planning increases slightly, but the reduction in sample size reduces significance to the 7% level.

Table 15
Ordered probit regression results,
for all retired households

Variable	Coeff.	Std. Err.	Pr > t
Planning variable	0.0886**	0.0407	0.029
Male	0.0886	0.1186	0.455
Prev. married	-0.0206	0.1465	0.888
Nev. married	0.1396	0.1866	0.455
Age	-0.0008	0.0083	0.927
Coll. or below	0.0964	0.1351	0.476
Prof. degree	-0.0581	0.1281	0.650
Faculty	0.1295	0.1188	0.276
R. DB plan	0.0082	0.1179	0.945
S. DB plan	-0.1204	0.1440	0.403
Ret. HI ins.	-0.0272	0.1042	0.794
R. LTC ins.	0.1265	0.1513	0.403
S. LTC ins.	-0.0972	0.1916	0.612

Source: Authors' tabulation of 2001 survey data.

Notes: The dependent variable is an ordered measure of the direction of the "spending surprise" reported by the individual (higher than expected, as expected, lower than expected). There were 604 observations in the regression, the Likelihood Ratio ($\chi^2(13)$) for the regression was 9.49; with a p-value of 0.735.

Why is low planning associated with a systematic underestimation of spending needs in retirement? The most intriguing possibility connects planning with a less downward-biased estimate view of retirement consumption. It may be easier to recognize which of one's current expenses will be less pressing than it is to visualize entirely new categories of expense associated with the retired life style. Do planning activities help to reduce this asymmetry, and force households to recognize that some expenses will actually increase? Clearly, this is an area in which more research is needed.

8 Concluding Remarks

We present some new insights on retirement consumption that are derived from survey data on some 2,000 TIAA-CREF participants. We show that retirement is seen as a relatively inexpensive life phase, as stressed by proponents of the life cycle model. We further find that many retired households appear to be consuming more than they had expected. We confirm that almost all of the gap between expectation and outcome can be accounted for by stock market appreciation.

Even though retirement does not produce a negative wealth shock, our survey suggests that households in the survey may be poorly prepared for retirement in other respects. Many households enter retirement profoundly uncertain about how their change in work status will impact their consumption needs. When this uncertainty is resolved, we find that more households are surprised by how high are expenditure needs in retirement rather than by how low they are. We find suggestive evidence that the systematic underestimation of retirement may reflect a bias in expectations.

While our sample is far from representative, it is our hope that the main methodological lessons, as well as several of the substantive findings, will apply more generally. In particular, our results suggest that survey measures of expected consumption may be of great value in macroeconomic research.

Appendix 1

Our goal is to estimate the present value of all remaining household income from employment from year 2000 on until retirement. We proceed as follows. For each household, we first compute how many years there are remaining until the expected date of retirement. This defines the number of years of (non-zero) employment income that we are going to estimate for that household. Some of this data is derived directly from the survey. For each of the future years for which data was requested (2001, 2005, and 2010), we enter the household's best estimate (if supplied) as a point estimate of the expectation of income for the year in question. We then fill in all of the intermediate years using simple interpolation assuming a constant rate of (geometric) growth in unreported years.

At the end of the interpolation stage, all households have data for some set of contiguous years. We then develop a fairly standard growth regression to predict the increase in income in a given year as a function of the prior year's growth in income, the prior year's level of income, age, age squared, age cubed, occupation, level of education, and gender. The regression is fitted to our interpolated income data for 2001-2010, as well as actual income data for 1998-2000, where available. The regression is summarized in table A1.

Table A1
Earnings growth regression estimates for working households


Variable	Coeff.	Std. Err.	Pr > t
$\Delta \ln Y(t - 1)$	0.1336**	0.0323	0.000
$\ln Y_{t-1}$	-0.0155**	0.0043	0.000
Age	-0.0190*	0.0103	0.065
Age ²	0.0004*	0.0002	0.064
Age ³	0.0000**	0.0000	0.044
Management	0.0095*	0.0054	0.081
Professional	0.0092**	0.0041	0.025
Other	-0.0020	0.0057	0.724
Coll. or below	-0.0026	0.0045	0.565
Prof. degree	0.0070	0.0045	0.124
Male	0.0111**	0.0035	0.002
Constant	0.3970**	0.1645	0.016

Source: Authors' tabulation of 2001 survey data.

Notes: The dependent variable is the change in income from period $t - 1$ to t (which would be denoted $\Delta \ln Y(t)$). There were 7,577 observations of 1,019 individuals; the standard errors reported above have been corrected for clustering based on individuals. The R^2 of the regression was 0.054; and the regression has an $F(11, 1018)$ of 11.76.

To illustrate the implications of this regression, Figure A1 shows the age-earnings profile that would be obtained for a male, Ph.D. faculty member, age 40 in 2000, whose income in 1998 was \$75,000, and whose income growth in 1999 was 0.03 log points over 1998.

Figure A1
Projected age-earnings profile



We use our earnings growth equation to extend all income data to the final retirement date. With this we have estimates of complete income profiles for all sample households from 2000 until retirement. All that remains is to pick an appropriate rate of interest to discount the stream to the present. We use the simple assumption of a constant interest rate of 3% per annum in the measure of employment income in the paper. We have performed some sensitivity analyses in which this interest rate is allowed to vary from 2% p.a to 7% p.a. None of the regressions in the paper are significantly impacted by these changes in the discount rate.

Appendix 2

To estimate the underlying distribution of responses to question 16b, we begin by noting that the propensity to answer question 16b varies with the response to question 16a. All of those who answered that they expected no change were automatically assigned an answer of zero to the “best estimate” part of question 16b, while only 90% of those expecting a change answered this question. To neutralize these response rate differentials, we assume that conditional on their state (retired/non-retired) and their answer to question 16a, the non-respondents are distributed in the same manner as the respondents. This amounts to reweighting the answers to question 16b by the inverse of the conditional response rate.

Figure A2
Rewighted Responses for Working Households





Figure A3
Reweighted Responses for Retired Households



Figures A2 and A3 show the reweighted distributions of answers to question 16b for non-retirees and retirees respectively. Two features stand out. First, answers are generally rounded. Most of those who gave positive or negative values gave answers that are rounded to the nearest five, but some may have rounded to the nearest ten. More importantly, there appears to be a fairly generous interpretation of “no change”. The spike at zero is quite large and seems to draw mass from a fairly wide range about the origin. In the case of the non-retirees, respondents are more likely to answer ± 20 than ± 5 , ± 10 , or ± 15 . Second, there is a hump in the distribution below -50% centered at -80% . These declines are implausibly large. It appears that these agents provided the replacement ratio rather than the percent change in consumption.

We estimate the underlying distribution of responses by maximum likelihood, allowing for the behavior noted in the previous paragraph. If x is the underlying change in consumption, we assume that $x + 100$ is log-normally distributed. We allow the mean and the variance of this distribution to depend on whether or not the respondent is retired. Let μ_N and σ_N^2 denote the mean and variance of the non-retired and μ_R and σ_R^2 the mean and variance of the retired. To handle the rounding to the nearest 5, we round all x 's to the nearest 5. Leaving the unrounded numbers has very little effect on the estimated coefficients. To handle the rounding to zero, we assume that respondent's with $x \in [-15, 15]$ round to zero with probability $1 - p(x)$. We restrict p to be symmetric about zero and to be the same for both retired and non-retired. To handle the observations below

-50, we assume that conditional on reporting that consumption will fall, respondents give the correct x with probability q and mistakenly report $1 - x$ with probability $1 - q$. The probability q is assumed to be the same for both retired and non-retired and to be independent of $x < 0$.

The results of the estimation are as follows

Table A2
Maximum Likelihood Parameter Estimates for
Distribution of Expected Change in Spending at Retirement

Parameter	Value.
μ_N	4.4777
σ_N^2	0.0403
μ_R	4.5471
σ_R^2	0.0404
q	0.8407
$p(5)$	0.0863
$p(10)$	0.6922
$p(15)$	0.5192

When we use this distribution to create a simulated sample, the result is generally very close to our actual sample. In particular, the simulated sample fits the hump below -50% very well. This confirms that almost all observations below -50% are due to reporting error. One place that the simulation does not fit so well is that it underestimates the number of actual respondents expecting a 50% fall in consumption. Presumably this reflects a second focal point that could be better captured in a yet more intricate model of rounding error.

9 Bibliography

References

- [1] Ameriks, J., A. Caplin, and J. Leahy [2002], “Wealth Accumulation and the Propensity to Plan”, in progress.
- [2] Banks, J., R. Blundell, and S. Tanner [1998], “Is There a Retirement Savings Puzzle?”, *American Economic Review*, vol. 88, 769-788.

- [3] Barsky, R., M. Kimball, T. Juster, and M. Shapiro, [1997] “Preference Parameters and Behavioral Heterogeneity: An Experimental Approach Using the Health and Retirement Survey”, *Quarterly Journal of Economics*, vol. 62, 537-579.
- [4] Bernheim, D., J. Skinner, and S. Weinberg [1997], “What Accounts for the Variation in Retirement Wealth Among U.S. Households”, Working Paper 6227, National Bureau of Economic Research, Cambridge, Mass.
- [5] Dillman, D. [1978], *Mail and Telephone Surveys: the Total Design Method*, Wiley and Sons, New York.
- [6] Hamermesh, D. [1984], “Consumption During Retirement: The Missing Link in the Life Cycle”, *Review of Economics and Statistics*, vol.66, 1-7.
- [7] Juster, T. [1966], “Consumer Buying Intentions and Purchase Probability: An Experiment in Survey Design”, *Journal of the American Statistical Association*, vol. 61, 658-696.
- [8] Lusardi, A. [1999], “Information, Expectations, and Savings”, ch. 3 in *Behavioral Dimensions of Retirement Economics*, H. Aaron ed., Brookings Institution/ Russell Sage Foundation, New York.
- [9] Mariger [1987], R., “A Life-Cycle Consumption Model with Liquidity Constraints: Theory and Empirical Results”, *Econometrica*, vol. 55, 533-557.
- [10] O’Donoghue, T., and M. Rabin [1999], “Doing it Now or Later”, *American Economic Review*, vol. 89, 103-124.