ARE LIFE INSURANCE HOLDINGS RELATED TO FINANCIAL VULNERABILITIES?

B. DOUGLAS BERNHEIM, KATHERINE GRACE CARMAN, JAGADEESH GOKHALE, and LAURENCE J. KOTLIKOFF*

Using the 1995 Survey of Consumer Finances and an elaborate life-cycle model, we quantify the potential financial impact of each individual's death on his or her survivors and measure the degree to which life insurance moderates these consequences. Life insurance is essentially uncorrelated with financial vulnerability at every stage of the life cycle. As a result, the impact of insurance among at-risk households is modest, and substantial uninsured vulnerabilities are widespread, particularly among younger couples. We also identify a systematic gender bias: For any given level of financial vulnerability, couples provide significantly more protection for wives than for husbands. (*JEL* D10, G22)

I. INTRODUCTION

The purpose of this article is to examine life insurance holdings and financial vulnerabilities among couples over the life cycle. Two separate concerns motivate our analysis. First, there are reasons to suspect that life insurance coverage is poorly correlated with underlying financial vulnerabilities. A well-known insurance industry adage holds that life insurance is sold and not bought. Alternatively, households may purchase long-term policies relatively early in life and subsequently fail to adjust coverage appropriately because of inertia or other psychological considerations. Second, households that purchase little or no life insurance may leave either or both spouses at risk of serious financial consequences.

With respect to the first concern, the available evidence is limited. Using a sample of older workers drawn from the 1992 wave of the Health and Retirement Study (HRS), Bernheim et al. (2003) documented a startling mismatch between life insurance holdings and underlying vulnerabilities (defined as the decline in living standard that a survivor would experience in the absence of insurance). In particular, they found virtually no correlation between these two variables. However, because the HRS focuses on individuals approaching retirement, this finding did not permit them to distinguish between the two hypotheses mentioned in the previous paragraph.

The second concern has received considerably more attention. One branch of the literature examines the experience of widows. Holden et al. (1986) and Hurd and Wise (1989) documented sharp declines in living standards and increases in poverty rates (from 9% to 35%) among women whose husbands actually passed away. A second branch of the literature projects the consequences of widowhood for married individuals. Analyzing data gathered during the 1960s from households in middle age through early retirement,

*We are grateful to the National Institute of Aging provided financial support. Economic Security Planner (ESPlanner) is used in this study with the permission of Economic Security Planning, Inc.

Bernheim: Professor, Department of Economics, Stanford University, Stanford, CA 94305-6072 and National Bureau of Economic Research. Phone 650-725-8732, Fax 650-725-570, E-mail bernheim@stanford.edu

Carman: Postdoctoral Fellow, Center for Basic Research in the Social Sciences, Harvard University, Cambridge, MA 02138. Phone 617-495-5365.

Gokhale: Senior Economic Advisor, Federal Reserve Bank of Cleveland, P.O. Box 6387, Cleveland, OH 44101-1387. Phone 216-579-2970, Fax 216-579-3050, E-mail jgokhale@clev.frb.org

Kotlikoff: Professor, Boston University, Department of Economics, 270 Bay State Rd., Boston, MA 02215. Phone 617-353-4002, Fax 617-353-4001, E-mail kotlikof@bu.edu

ABBREVIATIONS

HRS: Health and Retirement Study
OECD: Organisation for Economic Co-operation and Development
OLS: Ordinary Least Squares
SCF: Survey of Consumer Finances

Economic Inquiry
(ISSN 0095-2583)
Vol. 41, No. 4, October 2003, 531–554

DOI: 10.1093/ei/cbg026

Auerbach and Kotlikoff (1987; 1991a; 1991b) found that roughly one-third of wives and secondary earners would have seen their living standards decline by 25% or more had their spouses died. Bernheim et al. (2003) adopted a similar approach but used more recent, high-quality data, more accurate estimates of survival–contingent income streams, and a more elaborate life-cycle simulation model. They documented significant underlying vulnerabilities among older working couples; ignoring life insurance, one-third of secondary earners would have experienced significant (20% or greater) declines in living standards had their spouses died in 1992, whereas one-fifth would have experienced severe declines (40% or greater). More important, only one in three households in these at-risk populations held sufficient life insurance to avert significant or severe financial consequences. Though primary earners (typically husbands) were less frequently at risk for significant or severe consequences, they generally had much less protection (insurance on the secondary earner's life) for any given level of risk exposure. Notably, none of the studies mentioned examined uninsured exposures among younger households. This is a significant omission, because younger couples generally have more contingent income to protect and hence greater underlying vulnerabilities.

In the current study, we reexamine both sets of issues using data obtained from the 1995 wave of the Survey of Consumer Finances (SCF). Relative to the HRS (studied by Bernheim et al., 2003), the primary advantage of the SCF is that it includes adult respondents of all ages. Consequently, we extend the existing literature by examining financial vulnerabilities and life insurance coverage at all stages of the life cycle. This permits us to provide the first available evidence concerning uninsured financial vulnerabilities among young households, as well as to distinguish between the two alternative explanations for the absence of a correlation between financial vulnerabilities and insurance holdings mentioned—that insurance bears little relation to needs at the time of purchase and that households purchase long-term contracts with initially appropriate insurance coverage but fail to adjust this coverage through time as their circumstances change.

Because the SCF data are in some ways less comprehensive than the HRS data, this extension comes at a cost. Most notably, matching Social Security earnings histories are not available for the SCF sample. In addition, certain measurement problems (e.g., the likelihood that a nonworking widow or widower will enter the labor force or the potential for remarriage) are more severe for younger households than for older households. Nevertheless, because the SCF does contain information on many important factors that determine financial vulnerabilities, it is reasonable to infer that measured and actual financial vulnerabilities are correlated. If households respond to greater financial vulnerabilities by purchasing more insurance, one would still expect to observe a correlation between measured vulnerabilities and insurance holdings.

In practice, we find that aside from scale effects associated with the level of income, the correlation between financial vulnerabilities and life insurance holdings is essentially zero *throughout the entire life cycle*. This finding is consistent with the hypothesis that life insurance bears little relation to needs at the time of purchase; however, it tends to refute the hypothesis that households purchase long-term contracts with initially appropriate insurance coverage but fail to adjust this coverage through time as their circumstances change.

We also find that uninsured vulnerabilities are considerably greater among younger couples than among older couples. Nearly two-thirds of secondary earners between the ages of 22 and 39 have significant uninsured financial vulnerabilities (projected reductions in living standards exceeding 20%), and nearly one-third have severe uninsured vulnerabilities (projected reductions exceeding 40%). Moreover, for this age group, only one in five households with at-risk secondary earners (those who would experience significant or severe declines in living standard on the death of their spouse, ignoring insurance) held sufficient life insurance to avert significant or severe financial consequences. For the entire sample, in the absence of life insurance, 56% of secondary earners and 6% of primary earners would have experienced significant or severe declines in living standard on the death of a spouse. Actual life insurance holdings reduce these figures to, respectively, 42% and 5%. Thus, the overall impact of life insurance holdings on financial vulnerabilities among at-risk households is modest. Roughly two-thirds of poverty among surviving women and more than one-third of

poverty among surviving men resulted from a failure to ensure survivors of an undiminished living standard through insurance.

Whereas our analysis of the correlation between vulnerabilities and insurance holdings only requires us to assume that there is a reasonably strong correlation between actual and measured vulnerabilities, our findings concerning the prevalence of financial vulnerabilities rely more heavily on the accuracy of computed vulnerabilities. In light of the limitations of the SCF and the associated measurement problems, we therefore examine the robustness of the findings summarized in the preceding paragraphs with respect to a variety of alternative assumptions and imputations.

Finally, we find evidence of a significant gender bias in life insurance holdings. Specifically, for any given level of financial vulnerabilities, couples provide significantly more protection for wives than for husbands. For example, couples with severely at-risk wives on average hold $166,628 of life insurance, whereas couples with severely at-risk husbands on average hold only 15% of this amount ($24,827). Life insurance reduced the average impact of a husband's death on the living standard of a severely at-risk wife from a 65.5% decline to a 47.6% decline. In contrast, life insurance reduced the average impact of a wife's death on the living standard of a severely at-risk husband by a much smaller amount, from a 68.4% decline to a 64.1% decline. It is unlikely that these results are attributable to measurement problems. If anything, our analysis understates the true gender differential by overstating the vulnerabilities of nonworking spouses relative to working spouses.[1]

The article proceeds as follows. Section II discusses methods, section III describes the data, section IV presents basic results, section V provides sensitivity analysis, and section VI concludes.

II. A STRATEGY FOR MEASURING FINANCIAL VULNERABILITIES

Following Bernheim et al. (2003), we adopt a concrete and easily understood yardstick for quantifying financial vulnerabilities: the percentage decline in an individual's sustainable living standard that would result from a spouse's death.[2] The use of this yardstick permits us to make apples-to-apples comparisons of vulnerabilities across households and investigate correlations between vulnerabilities and insurance coverage. We also compare actual life insurance holdings to a natural benchmark, defined as the level of coverage required to ensure survivors of no change in their sustainable living standard. It is worth emphasizing that we do not regard this benchmark as a definitive standard of adequacy or rationality. Rational decision makers may elect to purchase either higher or lower levels of insurance. However, when combined with other evidence on household objectives, comparisons with the benchmark potentially shed light on the adequacy of life insurance coverage.

Concepts

We clarify our strategy for measuring financial vulnerabilities through an example. Imagine that a husband and wife each live for at most two years (equivalently, they are within two years of maximum life span). Both are alive initially, but either may die before the second year. The household's well-being depends on consumption in each year and survival contingency. As will be discussed further, we allow for the possibility that some ongoing expenditures are either exogenous or determined early in life by "sticky" choices. We refer to these expenditures as fixed consumption and to residual spending as "variable consumption." Let y_1 denote initial assets plus first period earnings net of fixed consumption, and let y_{2s} denote second period earnings net of fixed consumption in state $s = W, H, B$, where the state identifies survivors (wife, W, husband, H, or both, B). The couple divides first-period resources between variable consumption, c_1, saving, A, and insurance premiums, $p_i L_i$, $i = H, W$, where L_i represents the second-period payment to i if his or her spouse dies, and p_i denotes the associated price per dollar of coverage. Assets A earn the rate of return r. The couple faces the following constraints: $c_1 = y_1 - A - p_W L_W - p_H L_H$, $c_{2B} = y_{2B} + A(1+r)$,

1. Our base-case calculations assume that an individual does not increase hours worked in the labor force after the death of his or her spouse.

2. To calculate this decline, we make use of an elaborate life-cycle model. The model is embodied in financial planning software, Economic Security Planner (ESPlanner). Economic Security Planning, Inc., provides free copies of the software for academic research. For additional information, consult www.ESPlanner.com

and $c_{2i} = y_{2i} + A(1+r) + L_i$ for $i = W, H$, where c_{2i} denotes second period variable consumption in state i (for the moment, we ignore nonnegativity restrictions on life insurance and assets). Defining $p_B \equiv (1+r)^{-1} - p_w - p_H$, these equations imply

$$(1) \qquad c_1 + p_B c_B + p_W c_W + p_H c_H = y_1 + p_B y_B + p_W y_W + p_H y_H \equiv Y.$$

We equate living standard with per capita variable consumption adjusted for family composition. To determine each individual's living standard when both are alive, we divide variable consumption by a factor 2^{α}. We assume that $0 < \alpha \leq 1$; the second inequality reflects economies of scale associated with shared living expenses. To maintain a living standard that is constant across time and states of nature (in other words, one that is undiminished if and when either spouse dies), the couple must spend $2^{\alpha}C$ dollars in every period and state where both are alive for every C dollars in any state where only one survives. From (1), it is apparent that the household's highest sustainable living standard is

$$(2) \qquad c^* = Y/(2^{\alpha}[1 + p_B] + [p_W + p_H]).$$

The couple can guarantee that spouse j's death will not diminish i's living standard from its highest sustainable level, c^*, by purchasing a life insurance policy with face value $L_i^* = (c^* - y_{2i}) + (y_{2B} - 2^{\alpha}c^*)$.[3]

We measure underlying financial vulnerabilities by comparing an individual's highest sustainable living standard, c^*, with $c_i^n \equiv y_{2i} + A(1+r)$, which represents the living standard he or she would enjoy if widowed, ignoring life insurance. We define the variable $IMPACT$ (*ignoring insurance*) as $([c_i^n/c^*] - 1) \times 100$, $i = W, H$. This is a measure of the percent by which the survivor's living standard would (with no insurance protection) fall short of or exceed the couple's highest sustainable living standard. Similarly, we measure uninsured financial vulnerabilities by comparing c^* with $c_i^a \equiv y_{2i} + A(1+r) + L_i^a$, which represents the living standard that the individual would actually enjoy if widowed, based on actual life insurance coverage, L_i^a. We define the variable $IMPACT$ (*actual*) as $([c_i^a/c^*] - 1) \times 100$. This is a measure of the percent by which the survivor's living standard would (given actual levels of coverage) fall short of or exceed the couple's highest sustainable living standard. The $IMPACT$ variables are based on a concrete and easily understood yardstick for quantifying the consequences of a spouse's death.[4] We also compare actual household life insurance holdings, $L_H^a + L_W^a$, with the benchmark level, $L_H^* + L_W^*$.

For the preceding example, we implicitly assumed that individuals could borrow at the rate r and issue survival-contingent claims at the prices p_H and p_W. As a practical matter, households encounter liquidity constraints. They are also typically unable or at least very reluctant to purchase negative quantities of life insurance (buy annuities).[5] In solving for each household's highest sustainable living standard, we take these restrictions into account, smoothing consumption to the greatest extent possible.[6]

When the life insurance constraint binds, the benchmark living standard for a survivor, c_i^* (where $i = H$ *or* W), may be greater than the benchmark living standard for the couple while both spouses are still alive, c_B^*. This observation raises the following practical issue: When calculating $IMPACT$, should we set $c^* = c_i^*$, or $c^* = c_B^*$? Were we to use c_B^*, *actual* $IMPACT$ would be positive not only for households that depart from the benchmark by purchasing additional insurance ($L_i^a > L_i^*$) but also for constrained households that conform to the benchmark by purchasing no insurance ($L_i^a = L_i^* = 0$).

3. In the special case in which the household has Leontief preferences (defined over per capita adjusted expenditures), this is also the utility maximizing outcome.

4. Note that when actual life insurance is below the benchmark, the intact couple saves on insurance premiums, so its actual consumption exceeds c^*. Hence, the $IMPACT$ variables understate the change in living standard that an individual experiences on a spouse's death. However, because life insurance premiums typically account for a small fraction of expenditures, the degree of understatement is small.

5. A nonnegativity constraint for life insurance purchases is equivalent to the restriction that life annuities are not available for purchase at the margin. For further discussion, see Yaari (1965), Kotlikoff and Spivak (1981), and Bernheim (1987).

6. Formally, one can think of the outcome that we identify as the limit of the solutions to a series of utility maximization problems in which the intertemporal elasticity of substitution approaches zero. In the limit (the Leontief case), the household is actually indifferent with respect to the distribution of consumption across any years in which its living standard exceeds the minimum level.

In contrast, the use of c_i^* implies that *actual IMPACT* is positive when $L_i^a > L_i^*$ and zero when $L_i^a = L_i^* = 0$. Because we wish to use *actual IMPACT* as a measure of the extent to which a household deviates from the consumption-smoothed benchmark, we select c_i^* rather than c_B^*. As a result, the value of *IMPACT* ignoring insurance is always nonpositive (even though, absent insurance, the survivor's material living standard might actually increase on his or her spouse's death), and it equals zero whenever the corresponding benchmark insurance level, L_i^*, is zero.

Implementation

We actually evaluate each household's financial vulnerabilities using a more elaborate and realistic life-cycle model. As mentioned previously, the model is embedded in a financial planning software program, Economic Security Planner (ESPlanner). Although a complete description of the model would be prohibitively lengthy, it is important to summarize some key features.

For our base-case calculations, we assume that each individual lives to a maximum age of 95. We include children as members of the household through age 18. We represent household scale economies as follows: an expenditure of $(N+\beta K)^{\alpha}C$, when there are N adults and K children in the household provides the same standard of living for each household member as does an expenditure of C when there is only one adult in the household (this generalizes the adjustment factor used in our simple illustration). The coefficient β is a child–adult equivalency factor; we set it equal to 0.5.[7] The exponent α captures economies of scale in shared living. We set it equal to 0.678, which implies that a two-adult household must spend 1.6 times as much as a one-adult household to achieve the same living standard.[8]

Insurance needs depend on differences in survival-contingent income streams. Consequently, a careful and thorough treatment of the Social Security system is essential. In calculating benefits for retirement, survivors, parents, children, spouses, and dependent children, the model accounts for eligibility rules, early retirement reductions, delayed retirement credits, benefit recomputation, the legislated phased increase in the normal retirement age, the earnings test, restrictions on maximum family benefits, the wage indexation of average indexed monthly earnings, and the price indexation of benefits once received.

Various characteristics of the tax system, such as rate structure and the treatment of married couples, can alter insurance needs by influencing the distribution of after-tax income across the various survival contingencies. Consequently, a careful treatment of taxation is also critical. The model calculates federal and state income and payroll taxes for each year in each survival contingency.[9] It incorporates a wide range of provisions, including federal deductions and exemptions, the decision to itemize deductions, the taxation of Social Security benefits, the earned income tax credit, the child tax credit, the phase-out at higher income levels of itemized deductions, and the indexation of tax brackets to the consumer price index. In computing federal deductions, it determines whether the sum of state income taxes, mortgage interest payments, and property taxes is large enough to justify itemization. Contributions to tax-favored retirement savings accounts are excluded from taxable income, and withdrawals are included. Though the model determines total saving simultaneously with life insurance, tax-favored saving is specified exogenously.

Choices concerning housing may also affect life insurance needs. Unlike many other expenditures, housing outlays are not easily smoothed. It is difficult to scale mortgage, property tax, and insurance payments up and down with other expenditures. Cost and inconvenience discourage many households from moving or refinancing mortgages; others form psychological attachments to their homes and resist changing residences prior to death (Venti and Wise, 2001). Moreover, few households

7. Our child–adult equivalency factor is that used by the Organisation for Economic Co-operation and Development (OECD) (see Ringen, 1991). Nelson's (1992) work suggests a smaller value, but she considers total household expenditures, whereas our child–adult equivalency factor applies only to nonhousing consumption expenditure; for our base-case results, we treat housing expenditure as inflexible. It appears from Nelson's work that a higher equivalency factor is appropriate for nonhousing expenditures.

8. The OECD uses a value of 0.7 for α (see Ringen, 1991). Williams et al. (1998) consider values of 0.5 for both α and β.

9. The SCF does not contain data on the state of residency. Therefore, we assume all households reside in Massachusetts.

access the equity in their homes through refinancing or reverse annuity mortgages (Caplin, 2001). Consequently, for our base-case calculations we treat housing as fixed consumption. In effect, we assume that couples and survivors remain in the same home until death and die with home equity intact. Formally, we subtract housing expenses from income off the top, itemizing mortgage interest and property taxes as deductions for federal income tax purposes when it is optimal to do so, prior to smoothing variable consumption.

Several potentially important factors are omitted from our analysis. We do not model uncertainty concerning future income and nondiscretionary expenses (e.g., medical care). Because small groups of individuals can share risks to some extent, the adverse effect of uncertainty on living standard is probably greater for widows and widowers than for couples. For this reason, our analysis tends to understate insurance needs. We also neglect the possibility that an individual might remarry after a spouse's death. The extent to which remarriage mitigates the financial consequences of a spouse's death depends on one's view of the marriage market (see Lundberg, 1999, for a discussion). Although relatively few elderly individuals remarry after the death of a spouse (see Bernheim et al., 2003), remarriage is more common among younger households.

Table 1 summarizes some illustrative life insurance calculations.[10] We begin with a couple consisting of a 40-year-old man earning $45,000 per year and a 38-year-old woman earning $25,000 per year. The man intends to retire at age 64, the woman at age 63. They have two children, ages 5 and 7. The net value of their nonhousing assets is $50,000; in addition, they own a $150,000 home and have an unpaid mortgage balance of $120,000. They expect their real earnings to grow at the rate of 1% per year until retirement. They also expect to earn a real after-tax return of 3% on their nonhousing investments. According to our model, this couple must purchase $285,922 in term insurance on the husband's life, and no insurance on the wife's life, to ensure each potential survivor of an undiminished living standard. The remainder of the table illustrates the sensitivity of insurance needs to changes in various household characteristics and economic parameters. As one would expect, benchmark life insurance falls with age and with the earnings of the insured spouse; it rises with the addition of a child, with an increase in the rate of earnings growth, and with a reduction in the real interest rate (these last two changes increase the present discounted value of future human capital). There is one surprising finding: benchmark insurance also rises with the removal of one child. This result is attributable to the associated changes in Social Security survivor benefits, which are quantitatively important for the hypothesized family. For an otherwise identical family with high income (for which Social Security survivor benefits are less important), benchmark insurance rises monotonically with the number of children.

TABLE 1
Sample Life Insurance Benchmarks

	Insurance Benchmark for Husband	Insurance Benchmark for Wife
Base case	285,922	0
+Age (50,48)	127,455	0
–Age (30,28)	607,446	0
+Husband's earnings ($60K)	382,239	0
–Husband's earnings ($30K)	187,457	32,581
+Wife's earnings ($40K)	225,430	52,606
–Wife's earnings ($20)	282,246	0
+Child (age 9)	298,016	0
–Child (only the 5-year-old)	295,732	0
+Earnings growth (2%)	327,984	0
–Earnings growth (0%)	240,957	0
+Real interest rate (5%)	210,364	0
–Real interest rate (1%)	409,454	0

Notes: Assumptions for base case: age of husband: 40, age of wife: 38, husband's employee earnings: $45,000, wife's employee earnings: $25,000, husband's retirement age: 64, wife's retirement age: 63, number of children: 2, age of children: 5 and 7, nonhousing net wealth: $50,000, primary home value: $150,000, mortgage balance: $120,000, earnings growth: 1%, real interest rate: 3%.

Interpretation

We do not regard any single measure of financial vulnerability as ideal. Though elaborate, the life-cycle model used in our analysis is still an abstraction, and we have imperfect information concerning the economic circumstances of each household (see the appendix). Nor do we regard the benchmark level of life insurance as an objective standard of adequacy

10. For additional examples and for comparisons with recommendations generated by Quicken Financial Planner, see Gokhale et al. (2001).

or rationality. Optimal insurance coverage depends on a variety of considerations, including (but not limited to) the manner in which marginal utilities vary across survival states, the weights that households attach to the well-being of each family member, degrees of risk aversion, and load factors (more generally, the degree to which the industry departs from actuarially fair pricing). Consequently, it is possible to rationalize a wide range of behaviors.

Nevertheless, the absence of a significant correlation between life insurance and financial vulnerabilities (measured by benchmark life insurance) would be difficult to reconcile with theories of rational financial behavior. Even if a household places less weight on the well-being of a particular spouse, and even if it must pay actuarially unfair rates, it should still obtain greater insurance protection when the spouse in question is exposed to more severe financial consequences. To explain the absence of a correlation, one would need to believe either that our measure of benchmark life insurance is largely unrelated to underlying vulnerabilities, or that marginal utilities vary in a way that just offsets the differences in measured vulnerabilities. Both possibilities strike us as improbable.

Evidence of widespread and substantial uninsured vulnerabilities (as measured by *actual IMPACT* or by the divergence of actual insurance from the benchmark) are also more difficult to rationalize than it might at first appear. Most potential explanations presuppose that households deliberately choose different living standards for survivors. Yet this premise is inconsistent with preliminary findings from a financial planning case study at Boston University involving nearly 500 subjects. Each of these individuals constructed a comprehensive financial plan using the same financial planning software employed in the current study. Participants hoped to benefit from these sessions and therefore had strong incentives to provide accurate information. Though the software permits users to specify different living standards for intact couples and each potential survivor, the vast majority of subjects selected the same living standard for each contingency.[11] Though it is perfectly rational for individuals to have other objectives, it is irrational for individuals with these objectives to purchase coverage that diverges significantly from our benchmark (assuming, of course, that the benchmark is derived from a model that correctly depicts all important aspects of the household's opportunity set).

11. Even with risk aversion, such choices are reasonable if load factors are low. For evidence on load factors in the context of life annuities, see Mitchell et al. (1999).

III. DATA

The 1995 wave of the SCF was fielded between June and December 1995. It surveyed over 4,000 households (2,874 married couples, 1,425 single individuals), with oversampling of the wealthy. The data cover demographics, income, wealth, debt and credit, pensions, attitudes about financial matters, the nature of various transactions with various types of financial institutions, housing, real estate, businesses, vehicles, health and life insurance, current and past employment, current Social Security benefits, inheritances, charitable contributions, education, and retirement plans. The architects of the SCF data files imputed missing information, supplying five "implicates" for each household.[12] We use the first implicate.

Our final sample consists of 1,033 couples. We excluded couples for the following reasons: (1) a spouse was self-employed or owned and actively managed a business (67% of excluded observations); (2) a spouse was temporarily unemployed (11%); (3) neither spouse had regular earnings as an employee (54%); (4) labor earnings were defined in terms of a unit other than time worked, for example by the piece (0.7%); (5) mortgage information was inconsistent (7.4%); (6) property taxes were greater than 5% of the value of the home (2.6%); (7) a spouse was over the age of 85 (1.6%); or (8) the couple's reported income and other economic resources were insufficient to support its reported fixed expenditures (3.3%) (note that some households fell into more than one category).

Accurate measurement of life insurance coverage is, of course, particularly critical for our analysis. Fortunately, the SCF data match up reasonably well with other sources of information concerning this variable. In Table 2, we make some comparisons between statistics on life insurance coverage (including all individual and group policies) drawn from the SCF

12. Kennickell (1991) provides a description of the imputation procedure.

TABLE 2
Validation of the SCF Life Insurance Data

	Percent Covered (All Life Insurance)		Mean Coverage (All Life Insurance)	
	LIMRA	SCF	LIMRA	SCF
All households	78	85	132,304	163,973
Ages 18–24	47	51	92,222	96,438
Ages 25–34	75	80	149,476	159,916
Ages 35–44	85	87	202,150	203,759
Ages 45–54	83	92	159,569	200,058
Ages 55–64	83	86	96,567	105,441
65 and older	77	87	27,156	39,692

Note: Life Insurance Marketing Research Association (LIMRA) figures are based on a 1992 survey (see American Council of Life Insurance, 1994), with mean coverage adjusted to 1995 dollars.

and from a survey fielded by the Life Insurance Marketing Research Organization, an authoritative industry source. Generally, the figures are quite close. Certainly, there is no indication that the SCF understates life insurance coverage. It is, of course, possible that households underreported certain forms of insurance (e.g., employer-provided policies) in both surveys. As an additional check on the validity of the data, we computed the aggregate amount of in-force life insurance implied by the SCF survey responses and compared this with total in-force life insurance reported by the industry (obtained from the American Council on Life Insurance, 1999). Because the latter figure is derived directly from company records, it is presumably reliable. The SCF survey data accounts for roughly 81% of aggregate in-force life insurance ($9.52 trillion out of $11.70 trillion). Because some life insurance policies are owned by companies, trusts, and foreign individuals rather than by U.S. households, the SCF figure appears to be in the right ballpark.

An important limitation of the SCF is that it contains information only on the total amount of life insurance held by each household, and not on the division of this insurance between spouses. In contrast to Bernheim et al. (2003) (who treated an individual spouse as the unit of analysis throughout), we therefore focus in much of our analysis on household aggregates. However, for some purposes, we also impute the fraction of insurance held to protect each spouse. The imputation is based on a regression equation explaining the fraction of a couple's total life insurance held on the life of the husband as a function of the age of each spouse, the husband's earnings, the husband's share of the couple's total nonasset income, family size, and the husband's share of the couple's total benchmark life insurance, estimated using data from the 1992 HRS.

IV. RESULTS

Analysis of Household Insurance Holdings

We begin with an analysis of household insurance holdings. For the results in this section, it is not necessary to impute the fraction of a household's insurance held to protect each spouse.

Table 3 provides a variety of summary statistics, including some simple information on insurance coverage. According to figures in panel A, the average benchmark level of life insurance for a household exceeded actual life insurance holdings by a wide margin (more than two to one for the mean, and nearly three to one for the median). One might be inclined to conclude that on average, life insurance addresses between one-third to one-half of underlying exposures. However, this conclusion is premature. Underlying financial exposures vary dramatically across households. From an inspection of averages (Table 3), one cannot determine whether the distribution of insurance holdings matches up with the distribution of exposures. As we will see, the averages mask a startling mismatch between vulnerabilities and insurance holdings.

Table 4 provides a first glimpse at the relation between benchmark insurance and actual insurance across households. In panel A, we

TABLE 3
Descriptive Statistics

	Mean	Median
Panel A: Household-Level Variables		
Nonhousing net wealth	1,641,724	45,250
Primary home ownership	0.76	1
Primary home value	160,551	85,000
Household nonasset income	125,779	52,300
Total household life insurance, actual	206,022	71,000
Total household life insurance, benchmark	443,444	205,058
Number of children	0.97	1

	Husband		Wife		Primary Earner		Secondary Earner	
	Mean	Median	Mean	Median	Mean	Median	Mean	Median
Panel B: Individual-Level Variables								
Age	44	43	42	41	43	43	42	41
Nonwhite	0.172	0						
Gender	0	0	1	1	0.220	0	0.780	1
College degree	0.404	0	0.375	0	0.422	0	0.356	0
Pension coverage	0.716	1	0.538	1	0.750	1	0.504	1
Nonasset income	85,292	34,200	18,606	14,160	87,916	36,000	15,983	12,480
Imputed life ins.	166,215	54,391	39,806	14,685	160,539	47,597	45,483	15,598
Benchmark life ins.	405,916	149,535	37,528	0	415,229	168,224	28,214	0
Δ liv. std., ignore ins. (%)	−5.66	0.00	−28.6	−21.5	−3.77	0.00	−30.5	−24.6
Δ liv. std., imputed (%)	−2.40	0.66	−16.3	−7.75	−0.16	0.87	−18.5	−13.3

Notes: Imputed and benchmark life insurance refer to insurance on the life of the individual listed at the top of the column. Changes in living standard (Δ liv. std.) for the spouse listed at the top of each column depend on insurance on the life of the other spouse.

group households into four categories based on the benchmark insurance level (aggregated over spouses). As indicated in the second column, similar fractions of the data sample fall into each of these groups. The figures in the third column suggest that there is practically no relationship between benchmark insurance and the likelihood of holding some insurance. In this sense, the decision to obtain insurance appears to be uncorrelated with underlying vulnerabilities. The fourth and fifth columns permit comparisons between the medians of benchmark insurance and actual insurance (in each case, household aggregates) for each of the four groups. There is perhaps a small correlation across groups between insurance purchases and the benchmark. A similar conclusion follows from a comparison of means (columns six and seven), though the average level of actual insurance for those with no underlying vulnerabilities (benchmark insurance of zero) is anomalously large.

Though there is some indication that benchmark insurance is mildly correlated with actual insurance, the proper interpretation of this correlation is unclear. From panel A of Table 4, it is apparent that differences in the median and mean levels of actual insurance across groups correspond reasonably closely to differences in median household income (second to last column). Consider the alternative hypothesis that households with more income mechanically acquire more insurance without much deliberate consideration of vulnerabilities and needs. It is reasonable to entertain this possibility because many employers automatically provide life insurance equal to some multiple of salary, and because life insurance agents frequently base recommendations on simple rules of thumb involving multiples of earnings. This hypothesis would account for the positive correlation between benchmark and actual insurance levels among those for whom the benchmark is positive. Moreover, it would

TABLE 4
Characteristics of Households with Different Levels of Benchmark Insurance

Benchmark Household Insurance	Fraction of Households (%)	Percent Insured	Median Benchmark Insurance	Median Actual Insurance	Mean Benchmark Insurance	Mean Actual Insurance	Median Household Earnings	Median Average Age of Couple
Panel A: Insurance Levels (Household Aggregates)								
$0	21.2	83.1	0	68,000	0	215,174	65,000	56
$1–199,999	28.4	81.2	97,972	39,000	99,073	106,499	35,360	46
$200,000–449,999	25.8	82.8	309,668	100,000	313,821	174,565	50,000	38
$450,000 or more	24.6	83.5	674,480	100,000	1,359,286	346,001	62,220	31

Ratio of Benchmark Household Insurance to Household Earnings	Fraction of Households (%)	Percent Insured	Median Ratio of Benchmark Insurance to Household Earnings	Median Ratio of Actual Insurance to Household Earnings	Mean Ratio of Benchmark Insurance to Household Earnings	Mean Ratio of Actual Insurance to Household Earnings	Median Household Earnings	Median Average Age of Couple
Panel B: Ratios of Insurance to Earnings (Household Aggregates)								
0	21.2	83.1	0	0.96	0	2.53	65,000	56
0 to 3.99	26.8	87.0	2.07	1.37	2.08	2.09	54,000	47
4 to 7.99	24.9	89.9	5.59	1.91	5.69	2.84	54,789	39
8 or more	27.1	71.1	12.57	1.09	15.03	2.86	37,901	30

explain the otherwise anomalously high levels of insurance holdings among those for who benchmark insurance is zero (and who also tend to have higher incomes).

Panel B of Table 4 contains the same information as panel A, except that insurance holdings (again, household aggregates) are expressed as ratios to household earnings. This serves as a rough control for earnings, and permits us to evaluate the extent of the relationship between benchmark and actual levels of insurance that exists independently of the mechanical "income effect" hypothesized in the preceding paragraph. Once again, we divide the sample into four groups of roughly equal sizes (as indicated in the second column). Though the percent insured rises slightly with the ratio of benchmark insurance to earnings across the first three groups, it is lowest for the most vulnerable group (those with ratios of benchmark insurance to earnings in excess of 8). Moreover, regardless of whether one looks at medians (columns four and five) or means (columns six and seven), the correlation across groups between the ratio of benchmark insurance to earnings and the ratio of actual insurance to earnings is weak at best. These observations suggests that the small correlation between benchmark and actual insurance levels noted in our discussion of panel A results from an income effect and not from deliberate evaluation of other factors affecting vulnerabilities.

Table 5 summarizes the relationship between benchmark and actual levels of life insurance (household aggregates) through some simple univariate regression models. Panel A presents results based on levels of life insurance, and panel B examines ratios of life insurance to earnings. In each panel, we estimate the equation of interest using three different methods: Ordinary least squares (OLS), tobit (to check for potential biases arising from the sizable fraction of households who hold no insurance), and median regression (to check for sensitivity to outliers). In addition, we also estimate a probit model explaining the likelihood that a household's actual insurance holdings are strictly positive.

In panel A, we find that a $1 increase in the benchmark tends to coincide with a 22-cent increase in actual insurance holdings for the household. For each of the three methods employed, we estimate this coefficient with considerable precision. Households with higher levels of benchmark insurance are also more likely to hold strictly positive insurance, but the difference is not statistically significant. In panel B we find that the correlations noted in panel A result almost entirely from systematic differences in income. The ratio of benchmark insurance to earnings is almost completely unrelated to the ratio of actual insurance to earnings and is in fact negatively related to the likelihood of holding insurance.

TABLE 5
Regression Results, Actual Life Insurance versus Benchmark Life Insurance, Household Aggregates

	Constant	Benchmark
Panel A: Level of Actual Household Insurance		
OLS	109,294.5	0.2181
	(10,904.63)	(0.0029)
Tobit	61,187.04	0.2187
	(12,799.07)	(0.0033)
Median regression	20,275.87	0.2202
	(11,529.65)	(0.0453)
Probit (probability scaled, using benchmark insurance/10^5)		0.0010
		(0.0021)
Panel B: Ratio of Life Insurance to Earnings		
OLS	2.432	0.024
	(0.2598)	(0.0275)
Tobit	1.911	−0.042
	(0.3057)	(0.0341)
Median regression	1.487	−0.026
	(0.1316)	(0.0091)
Probit (probability scaled)		−0.0093
		(0.0015)

Note: Standard errors are given in parentheses.

The logic of examining simple univariate regressions is that benchmark insurance functions as a sufficient statistic for the full range of factors—ages of spouses, number of children, division of earnings between spouses, and so forth—that determine underlying vulnerabilities. Certainly, it would not be sensible to control for all of the factors affecting vulnerabilities, because this would leave no residual variation in benchmark insurance from which to identify the covariation with actual insurance.

There are, however, valid reasons to investigate the correlation between benchmark and actual levels of insurance conditional on a subset of the factors that determine vulnerabilities. As mentioned at the outset, Bernheim et al. (2003) previously documented the absence

of a correlation between insurance holdings and vulnerabilities among older workers. One of our primary objectives is to distinguish between two alternative explanations for this pattern: that insurance purchases are largely uncorrelated with vulnerabilities, or that households purchase long-term policies based on vulnerabilities relatively early in life but subsequently fail to adjust coverage appropriately because of inertia or other psychological considerations. To accomplish this objective, we must study the relationship between life insurance and underlying vulnerabilities conditional on age.

The final column of panels A and B in Table 4 provide a separate reason to control for age. Note that both the absolute value of benchmark insurance and the ratio of the benchmark to household earnings decline sharply with the couple's average age (defined as the average of the husband's and wife's ages). This stands to reason, because younger individuals generally have more future human capital to protect. Now consider the following two-part hypothesis: (1) Young individuals tend to procrastinate with respect to decisions concerning life insurance, and (2) when they finally obtain insurance, the amount purchased is closely related to vulnerabilities. Because younger households have greater vulnerabilities, the first portion of the hypothesis would tend to create a negative correlation between insurance holdings and underlying vulnerabilities. In principle, this could obscure the positive correlation implied by the second portion of the hypothesis.

Table 6 contains various regressions in which we control for the effects of age. In each case, the dependent variable is the ratio of actual life insurance (household aggregate) to household earnings.[13] There are two separate sets of results. For the first four regressions, we include linear and quadratic terms in the couple's average age (already defined), as well as an interaction term between age and the ratio of benchmark life insurance (household aggregate) to household earnings. The interaction term permits us to investigate the hypothesis that the correlation between insurance holdings and vulnerabilities changes with age. If households initially purchase appropriate levels of insurance but subsequently fail to adjust their holdings as circumstances change, we would expect to obtain a positive coefficient for the benchmark insurance term and a negative coefficient for the interaction term. For the last four regressions in Table 6, we adopt a more flexible functional specification. In particular, we divide the population into five age groups and control both for age-group dummies and for interactions between these dummies and the benchmark insurance–to-earnings ratio. This is equivalent to estimating a separate univariate regression for each age group. At the bottom of the table, we report test statistics for two hypotheses. The first is that there is no relation between actual insurance and benchmark insurance at any age. The second is that the relation between actual insurance and benchmark insurance does not change systematically with age.

In the first three specifications of Table 6, the coefficient of the benchmark insurance–to-earnings ratio is small, slightly negative, and statistically insignificant, whereas the insurance–age interaction term is indistinguishable from zero. This implies that there is essentially no relationship between actual insurance and benchmark insurance at any age. The test statistics at the bottom of the table confirm this impression. For the fourth specification (a probit model), the coefficient of the benchmark insurance–to-earnings ratio is negative and marginally significant, and the coefficient of the interaction term is positive but insignificant. One rejects the hypothesis that vulnerabilities (as measured by benchmark insurance) have no effect on the likelihood of holding insurance at any age, but the direction of the effect is counterintuitive: The probit coefficients imply that the likelihood of holding insurance is negatively correlated with benchmark insurance for households under 100 years of age.

Regression results based on a more flexible functional specification tell a similar story. The age-specific benchmark insurance coefficients are generally small, frequently negative, and almost always insignificant. For the OLS and tobit specifications, one cannot reject the joint hypothesis that there is no significant correlation between benchmark insurance and actual insurance at any age. In every specification, we

13. Scaling the insurance variables by earnings is, of course, a perfect control for earnings only if the earnings elasticity of insurance purchases is unity. When an earnings variable (either levels or logs) is added to the list of explanatory variables for any specification in Table 5, the results change relatively little.

TABLE 6

Additional Regression Results, Actual Life Insurance versus Benchmark Life Insurance, Household Aggregates

	OLS	Tobit	Median Reg.	Probit	OLS	Tobit	Median Reg.	Probit
Benchmark insurance to household income	−0.0485	−0.1923	−0.0524	−0.0102				
	(0.1057)	(0.1355)	(0.0337)	(0.0059)				
Average age of couple	0.0540	0.1374	0.0854	0.0195				
	(0.1390)	(0.1701)	(0.0507)	(0.0077)				
Average age of couple squared	−0.0012	−0.0023	−0.0012	−0.0002				
	(0.0013)	(0.0017)	(0.0005)	(0.0001)				
Benchmark insurance to income interacted with age	−0.0009	0.0034	0.0009	0.0001				
	(0.0034)	(0.0042)	(0.0012)	(0.0002)				
Age < 22 (group 1)					0.15054	−1.7201	0.00	0.1367
					(3.8696)	(5.1116)	(1.7100)	(0.2005)
Age 22–40 (group 2)					3.4352	3.1355	2.2503	0.2959
					(0.5222)	(0.6250)	(0.2103)	(0.0295)
Age 41–55 (group 3)					2.2872	1.8239	1.5278	0.3293
					(0.4705)	(0.5415)	(0.2209)	(0.0305)
Age 56–70 (group 4)					1.6729	0.7231	0.6167	0.2340
					(0.5756)	(0.6674)	(0.1167)	(0.0330)
Age 71+ (group 5)					0.7924	−1.2351	0.2759	0.1109
					(1.4840)	(1.7915)	(0.1381)	(0.0750)
Benchmark insurance to earnings ratio interacted with age group 1					0.0458	−0.0405	0.00	−0.0077
					(0.1213)	(0.1706)	(0.0814)	(0.0069)
Benchmark insurance to earnings ratio interacted with age group 2					−0.0306	−0.1098	−0.0556	−0.0091
					(0.0408)	(0.0510)	(0.0128)	(0.0222)
Benchmark insurance to earnings ratio interacted with age group 3					0.0738	0.0400	0.0362	−0.0083
					(0.1261)	(0.1464)	(0.0779)	(0.0076)
Benchmark insurance to earnings ratio interacted with age group 4					−0.0066	−0.0052	0.0348	0.0003
					(0.2239)	(0.2575)	(0.0632)	(0.0131)
Benchmark insurance to earnings ratio interacted with age group 5					−0.0368	0.2251	0.0750	0.0389
					(0.7363)	(0.8592)	(0.1057)	(0.0443)
Constant	2.7713	0.8826	0.3540					
	(3.43)	(4.19)	(1.3460)					
p-*value for tests*								
No relationship between actual and benchmark at any age	0.840	0.184	0.123	0.009	0.958	0.435	0.002	0.001
The relationship between actual and benchmark is the same at all ages	0.793	0.420	0.456	0.598	0.925	0.864	0.287	0.794

Note: Standard errors are given in parentheses.

fail to reject the hypothesis that this correlation does not vary significantly with age.

Analysis of Vulnerabilities for Individual Spouses

We turn next to an individual-specific analysis of financial vulnerabilities. This requires us to impute the fraction of a household's insurance held to protect each spouse, as described in section III.

It is useful to begin once again with summary statistics. Panel B of Table 3 indicates that the average level of insurance actually held on the lives of husbands and primary earners diverged sharply from the average benchmark, whereas the average level of insurance actually held on the lives of wives and secondary earners was quite close to the average benchmark. This implies that on average, wives and secondary earners faced substantial uninsured vulnerabilities, whereas husbands and primary earners did not.

In the second-to-last line of panel B (Δ liv. std., ignore ins.), we tabulate means and medians for *IMPACT* calculated as if each household held no life insurance. This variable measures underlying financial exposure. Without insurance, the average living standards for surviving husbands, wives, primary earners, and secondary earners would have been, respectively, 5.7%, 28.6%, 3.8%, and 30.5% below their benchmark levels. Because the corresponding medians for husbands and primary earners are zero, we can infer that more than half of these individuals were not at risk of a reduction in sustainable living standard. Likewise, because the corresponding median living standard reductions for wives and secondary earners were 21.5% and 24.6%, respectively, we conclude that more than half of these individuals confronted significant underlying vulnerabilities.

In the final line of panel B (Δ liv. std., imputed), we tabulate means and medians for *IMPACT* based on imputed insurance holdings. This variable measures uninsured financial exposure. Life insurance reduced the mean value of *IMPACT* by roughly 40% for wives and secondary earners. The decline in the median value of *IMPACT* for secondary earners is a bit larger, whereas the proportional decline in median *IMPACT* for wives is closer to two-thirds. For husbands and primary earners, life insurance had a relatively minor absolute effect on mean and median *IMPACT*, but the average underlying vulnerabilities were small for these groups to begin with.

As before, the averages in Table 3 obscure the mismatch between vulnerabilities and insurance holdings. Table 7 provides further information on the distributions of the *IMPACT* variables. As discussed previously, ignoring insurance, *IMPACT* is never strictly greater than zero. This reflects the fact that we have imposed a nonnegativity constraint on life insurance purchases. An individual's living standard may rise on a spouse's death; however, without life insurance, it cannot exceed the living standard that the he or she would enjoy as a survivor assuming implementation of the (constrained) benchmark financial plan. Note also that *actual IMPACT* is exactly equal to zero for a substantial fraction of the population. Generally, these are individuals for whom actual and (model-generated) benchmark levels of insurance protection are both zero.

From Table 7, it is again evident that insurance holdings match up rather poorly with financial vulnerabilities. Overall, 36.8% of wives, 64.1% of husbands, 32.0% of secondary earners, and 68.8% of primary earners had strictly positive life insurance protection despite the fact that they would have experienced increases in living standards on the deaths of their spouses (in Table 7, those with *actual IMPACT* greater than zero). Insurance reduced the fraction of individuals at risk for severe financial consequences (defined as a decline in living standard of 40% or greater) from 31.0% to 19.3% for wives, from 32.5% to 20.6% for secondary earners, from 3.5% to 2.9% for husbands, and from 1.9% to 1.6% for primary earners. Similarly, insurance reduced the fraction of individuals at risk for significant financial consequences (defined as a decline in living standard of 20% or greater) from 51.9% to 37.9% for wives, from 55.9% to 41.7% for secondary earners, from 10.0% to 8.5% for husbands, and from 6.0% to 4.7% for primary earners. Roughly speaking, only 15% to 25% of households with significant financial exposures held sufficient life insurance to avert significant consequences for survivors.

Table 8 provides further information on the extent to which insurance mitigates financial vulnerabilities. We subdivide the population based on underlying financial exposure, measured by the value of *IMPACT*, ignoring

TABLE 7
Distribution of Changes (%) in Living Standard for Surviving Spouses

	Surviving Spouses are			
	Wives		Husbands	
IMPACT	Ignoring Insurance	Imputed Insurance	Ignoring insurance	Imputed Insurance
Panel A: Husbands and Wives				
<−40%	30.98	19.26	3.48	2.90
−40% to −20%	20.91	18.68	6.49	5.61
−20% to 0%	19.36	19.94	17.42	15.10
0%	28.75	5.32	72.60	12.29
0% to 20%	—	31.85	—	63.12
20% to 40%	—	3.39	—	0.68
>40%	—	1.55	—	0.29
Observations	1,033	1,033	1,033	1,033
	Surviving Spouses Are			
	Secondary earners		Primary earners	
IMPACT	Ignoring Insurance	Imputed Insurance	Ignoring Insurance	Imputed Insurance
Panel B: Primary and Secondary Earners				
<−40%	32.53	20.62	1.94	1.55
−40 to −20%	23.33	21.10	4.07	3.19
−20% to 0%	20.52	22.17	16.26	12.88
0%	23.62	4.07	77.73	13.55
0% to 20%	—	27.49	—	67.47
20% to 40%	—	3.00	—	1.06
>40%	—	1.55	—	0.29
Observations	1,033	1,033	1,033	1,033

insurance. For each subgroup, we then calculate the means of both *IMPACT* variables, as well as the percentage of households holding no life insurance, and averages for benchmark and actual insurance.

We draw three conclusions from Table 8. First, life insurance had, at best, a moderate impact on financial exposures among the at-risk population. For example, among severely at-risk wives, insurance reduced the average consequences of a spouse's death (*mean IMPACT*) by only 27% (17.9 percentage points), from −65.5% to −47.6%. This is a far cry from the one-half overall reduction in mean *IMPACT* for wives noted in Table 3. For this same group, households would have needed to hold an average of $630,079 in total life insurance to ensure both the husband and wife of an undiminished living standard. In fact, on average, they held only one-quarter of this amount ($166,628). This implies an even larger discrepancy (more than $460,000) than the overall difference (just under $240,000) noted in Table 3.

Second, for a fixed level of financial exposure, households were more inclined to protect women than men. For example, among severely at-risk husbands, insurance reduced the average consequences of the wife's death (*mean IMPACT*) by only 6% (4.3 percentage points), from −68.4% to −64.1%. This contrasts sharply with the corresponding figures for wives, already noted. One might question the reliability of this finding on the grounds that we have imputed the division of insurance between spouses based on patterns in the HRS. However, the gender difference is plainly not an artifact of the imputation procedure. For example, couples with severely at-risk wives on average held in total $190,388 of life insurance (versus $166,628 imputed to the protection of the wife), whereas couples with severely

TABLE 8
Effect of Life Insurance on Changes in Living Standards for Surviving Spouses by Level of Vulnerability

		Mean *IMPACT*		Insurance Holdings		
Surviving Spouses Are	**Range for *IMPACT*, Ignoring Insurance**	**Ignoring Insurance (%)**	**Imputed Insurance (%)**	**Percent Uninsured**	**Mean Benchmark**	**Mean Actual**
Wives	<−40%	−65.5	−47.6	22.8	630,079	166,628
	−40% to −20%	−30.1	−17.2	14.8	908.146	265,210
	−20% to 0%	−10.3	1.68	10.5	107,633	133,122
	0%	0.0	6.14	18.2	0	116,059
Husbands	<−40%	−68.4	−64.1	22.2	291,568	24,827
	−40% to −20%	−27.1	−23.4	22.4	210,961	23,056
	−20% to 0%	−8.7	−4.4	18.3	78,528	37,625
	0%	0.0	2.9	16.7	0	42,545
Secondary earners	<−40%	−65.8	−49.1	22.3	614,989	159,234
	−40% to −20%	−29.8	−18.2	14.9	830,155	236,420
	−20% to 0%	−10.3	−0.3	13.2	104,850	116,975
	0%	0.0	6.9	16.8	0	125,237
Primary earners	<−40%	−66.4	−53.2	30.0	274,272	35,598
	−40% to −20%	−26.8	−21.5	26.2	243,488	44,115
	−20% to 0%	−8.6	−3.1	15.5	79,962	51,181
	0%	0.0	2.9	17.2	0	44,609

at-risk husbands on average held in total only one-third of this amount ($62,689 versus $24,827 imputed to the protection of the husband).

Third, the likelihood of holding insurance bears little if any relation to underlying vulnerabilities, as measured by *IMPACT* (*ignoring insurance*). Indeed, the fraction of couples without life insurance is generally largest for those with severe financial exposures (*IMPACT ignoring insurance* less than −40%).

Thus far we have measured the consequences of a spouse's death in terms of the proportional change in sustainable living standard. For many individuals, the potential financial consequences of a spouse's death are also severe in absolute terms. With maximal consumption smoothing (benchmark insurance), sustainable consumption for 3.58% of surviving wives and 2.61% of surviving husbands would fall below the 1995 poverty thresholds published by the U.S. Census Bureau. Taking into account actual levels of insurance coverage, poverty rates would have been 10.45% among surviving wives and 4.16% among surviving husbands. These findings imply that 66% (6.87 of 10.45 percentage points) of poverty among surviving women and 37% (1.55 out of 4.16 percentage points) of poverty among surviving men resulted from a failure to ensure survivors of an undiminished living standard through insurance. Ignoring insurance, poverty rates would have been 13.17% among surviving wives and 4.26% among surviving husbands. Consequently, insurance eliminated only 28% of the avoidable poverty among surviving widows (2.72 out of 9.59 percentage points) and only 6% of the avoidable poverty among surviving men (0.1 out of 1.63 percentage points).

Results for Population Subgroups

Table 9 provides disaggregated results. For various population subgroups, we report the ratio of actual-to–benchmark life insurance calculated two ways (mean to mean and median to median). These figures pertain to the entire household and therefore do not rely on the imputed division of insurance between spouses. For primary and secondary earners, we also report the percentage of each subgroup with severe and significant exposures based on *IMPACT* with actual insurance ("frequency"), as well as the fractional reduction in these exposures resulting from actual

TABLE 9

Life Insurance Coverage and Financial Consequences for Surviving Spouses: Selected Population Subgroups

	Ratio of Actual-to-Benchmark Insurance (Household Totals)		Consequences for Secondary Earners				Consequences for Primary Earners			
			Severe (>40%)		Significant (>20%)		Severe (>40%)		Significant (>20%)	
	Mean-to-Mean	Median-to-Median	Frequency (%)	Fraction Addressed	Frequency (%)	Fraction Addressed	Frequency (%)	Fraction Addressed	Frequency (%)	Fraction Addressed
Full sample	0.465	0.346	20.6	0.366	41.7	0.254	1.55	0.201	4.74	0.210
HH earnings < $15K	0.145	0.041	46.2	0.143	53.9	0.159	5.13	0.080	5.13	0.000
HH earnings $15–45K	0.333	0.116	30.1	0.246	51.6	0.161	2.73	0.080	7.65	0.125
HH earnings $45–100K	0.465	0.381	15.0	0.481	40.0	0.292	0.65	0.000	3.91	0.308
HH earnings > $100K	0.522	1.62	9.52	0.555	22.0	0.449	0.60	0.000	0.60	0.496
Dual earners	0.508	0.345	17.1	0.385	41.9	0.273	1.69	0.214	5.99	0.250
Single earners	0.428	0.491	26.7	0.290	41.4	0.217	1.31	0.510	2.62	0.090
Earnings diff. 1-1 to 2-1	0.503	0.384	9.07	0.426	35.5	0.300	2.59	0.167	9.07	0.238
Earnings diff. over 4-1	0.440	0.364	28.0	0.344	43.6	0.219	1.02	0.171	2.04	0.093
Age of survivor < 22	0.056	0.000	60.0	0.000	90.0	0.000	0.00	0.000	9.09	0.000
Age of survivor 22–39	0.276	0.196	31.1	0.344	62.6	0.203	1.67	0.301	6.70	0.200
Age of survivor 40–55	0.825	0.813	11.9	0.500	27.6	0.320	0.71	0.245	2.12	0.000
Age of survivor 56–70	2.57	NA*	7.64	0.293	13.4	0.298	0.67	0.000	2.01	0.250
Age of survivor > 71	16.8	NA*	8.82	0.000	11.8	0.197	16.7	0.000	26.7	0.000
No children	0.605	0.538	19.2	0.286	37.2	0.200	1.98	0.168	6.14	0.162
One or more child	0.411	0.349	22.0	0.280	46.0	0.290	1.14	0.070	3.41	0.281
Whites	0.463	0.403	20.1	0.389	40.9	0.266	1.17	0.287	4.21	0.217
Nonwhites	0.483	0.150	23.0	0.256	45.5	0.198	3.37	0.000	7.30	0.188

Note: *Median benchmark insurance is zero for these groups.

insurance holdings ("fraction addressed"),[14] in each instance using the imputed division of insurance (results for husbands and wives, not reported in the table, are similar).

Uninsured financial vulnerabilities were more common among low-income households, couples with disparate earnings, relatively young households, couples with dependent children, and nonwhites. Thus, factors that are highly correlated with underlying vulnerabilities (particularly earnings disparities, age, and children) are also highly correlated with uninsured vulnerabilities. This finding is, of course, implied by the poor correlation between vulnerabilities and insurance holdings.

According to Table 9, conditional on the existence of a significant or severe vulnerability, households with lower incomes, greater income disparities between spouses, no children, and nonwhites are less likely to moderate the financial consequences of a spouse's death through life insurance (see the column labeled Fraction Addressed). The relationship between age and the conditional likelihood of addressing a significant or severe vulnerability was hump-shaped, with a peak in the 40–55-year age range. Households were generally less likely to address vulnerabilities for primary earners. Note that a low proclivity to address exposures can coincide either with high (as in the case of lower-income households) or low (as in the case of older individuals and primary earners) levels of underlying vulnerability.

When one compares households of similar ages, our results are quite close to those derived from the HRS data in Bernheim et al. (2003).[15] For example, we find that 13.4% of secondary earners between the ages of 56 and 70 have significant uninsured vulnerabilities, whereas 7.6% have severe uninsured vulnerabilities; the comparable figures in Bernheim et al. (2003) for 60- to 69-year-old survivors are 14.1% and 8.9%. Moreover, the data exhibit the same qualitative patterns with respect to the magnitudes of vulnerabilities, as well as with respect to the propensity to address vulnerabilities, across household categories.

It is important to reiterate that data limitations precluded previous studies from shedding light on the insurance holdings of younger households. Consequently, our findings concerning age are particularly noteworthy. According to Table 9, for 22- to 39-year-olds, household life insurance holdings averaged only 27.6% of the benchmark, compared with 82.5% for 40- to 55-year-olds. Nearly two-thirds of secondary earners between the ages of 22 and 39 have significant financial vulnerabilities (projected reductions in living standards exceeding 20%), and nearly one-third have severe vulnerabilities (projected reductions exceeding 40%). Moreover, for this age group, only one in five households with at-risk secondary earners (those who would experience significant or severe declines in living standard on the death of their spouse, ignoring insurance) held sufficient life insurance to avert significant or severe financial consequences.

V. ROBUSTNESS

Given the limitations of the SCF data, it is important to examine the extent to which measured vulnerabilities are sensitive to changes in key assumptions and parameters. We summarize a variety of alternative calculations in Table 10. As in Table 9, we report household-level ratios of actual to benchmark life insurance, as well as the fractions of primary and secondary earners with severe and significant uninsured exposures. For purposes of comparison, we reproduce our base-case results in the first line of the table.

Changes in the real interest rate, baseline wage growth rate, and maximum lifespan have relatively little effect on the fractions of spouses with severe and significant vulnerabilities.[16] In each case, this reflects the opposing effects of offsetting forces. With higher interest rates, a given level of life insurance coverage generates higher real income. However, because survivors are typically more dependent on long-duration life annuities than intact couples, the present discounted value of their resources tends to decline by a larger proportion when the rate of return rises. For older

14. Formally, fraction addressed = [(Frequency given insurance = 0) – (Frequency given actual insurance)]/(Frequency given insurance = 0).

15. The age-adjusted frequencies of severe and significant uninsured financial exposures reported here and in Bernheim et al. (2003) are lower than Auerbach and Kotlikoff's (1987; 1991a; 1991b) estimates. Possible explanations for the disparity include increases in female labor force participation since the 1960s, changes in patterns of insurance coverage, and methodological differences.

16. For our base case, we assume that the inflation rate and real interest rate are both 3%.

TABLE 10

Frequency of Severe and Significant Financial Consequences for Surviving Spouses: Robustness (Full Sample)

	Ratio of Actual-to-Benchmark Insurance (Household Totals)		Consequences for Secondary Earners (Imputed Insurance) (%)		Consequences for Primary Earners (Imputed Insurance) (%)	
	Mean-to-Mean	**Median-to-Median**	**Severe (>40%)**	**Significant (>20%)**	**Severe (>40%)**	**Significant (>20%)**
Base case	0.465	0.346	20.6	41.7	1.55	4.74
Real interest rate = 1%	0.333	0.239	23.9	44.2	1.92	5.19
Real interest rate = 5%	0.610	0.462	19.3	38.5	1.35	4.15
Baseline wage growth rate = 0%	0.537	0.401	19.8	40.0	1.65	4.55
Baseline wage growth rate = 2%	0.403	0.306	22.6	42.9	1.45	5.14
Maximum lifespan = 85	0.473	0.352	20.5	41.2	1.45	4.66
Consumption growth rate = 1%	0.484	0.366	20.2	39.9	1.45	4.17
Consumption growth rate = –1%	0.450	0.342	21.1	42.5	1.55	5.02
No ecs. of shared living ($\alpha = 1$)	0.877	1.14	14.7	26.2	0.77	1.84
Survivor receives 50% pension benefits	0.435	0.300	22.3	45.3	1.65	4.84
Housing completely fungible	0.552	0.513	10.8	28.3	0.38	2.00
Survivor downsizes house by 30%	0.503	0.422	18.5	36.9	1.16	4.07
Survivors fully employed	0.533	0.569	11.5	29.6	1.45	4.65
All life insurance assigned to the primary earner	0.465	0.346	19.9	38.8	1.94	6.00
Wealth reduced by 20%	0.446	0.337	21.3	43.1	1.54	4.73
Wealth increased by 20%	0.484	0.357	20.2	40.9	1.45	4.45

workers, the rate of wage growth is relatively unimportant because it affects comparatively few years of earnings. Moreover, although a given rate of growth produces a larger absolute increase in earnings for primary earners, secondary earners tend to be younger and therefore benefit from higher growth over a longer time frame. A reduction in maximum lifespan reduces the resources that a survivor needs to achieve a given living standard but increases the living standard that the intact couple can achieve from available resources.

The consumption growth rate refers to steepness of the sustainable living standard trajectory. For our base case, we compute the highest living standard that is sustainable throughout life in all contingencies; this corresponds to a consumption growth rate of zero. For sufficiently patient (impatient) households, it may be more natural to construct benchmarks based on a rising (falling) living standard trajectory. The proportional effects of a change in the consumption growth rate on the resource needs of survivors and intact couples are approximately equal. Our results are therefore robust with respect to changes in this parameter.

The frequencies of exposure to significant and severe financial consequences are noticeably lower in the absence of household scale economies (an extreme and somewhat implausible assumption). It is somewhat higher when we reduce the rate of pension survivor benefits from 100% to 50%, but the change is much smaller than that reported in Bernheim et al. (2003). The discrepancy is presumably attributable to differences between the age distributions of the SCF and HRS samples. For younger households, pension survivor benefits make less of a difference both because they are further in the future and thus more heavily discounted and because younger workers have not yet accumulated substantial pension entitlements.

As mentioned previously, our base-case assumptions concerning housing are consistent with empirical evidence indicating that individuals avoid changing residences (increasingly so as they age) and that they resist using housing equity to finance ordinary living expenses. However, because some widows do move, we examine sensitivity to two alternative assumptions. For the first, we adopt the extreme position that housing consumption is completely and continuously flexible and that housing equity is a perfect substitute for other forms of wealth. For the second alternative, we adopt an intermediate position: A survivor downsizes the couple's primary residence by 30%, but thereafter avoids using housing equity to finance ordinary living expenses.[17] Though the first alternative dramatically reduces the estimated frequencies of individuals at risk of severe or significant financial consequences, the effect of the second (and, we think, more plausible) alternative is relatively small.

As in Bernheim et al. (2003), our base-case calculations assume that survivors do not alter their labor force participation. Because nonworking wives approaching retirement age have limited employment options subsequent to their husbands' deaths, this assumption was appropriate for the HRS sample. It is somewhat more problematic for the younger households included in the SCF. Conceivably, some couples may hold apparently low levels of insurance on a primary earner because, in the event of his or her death, they expect the surviving spouse to find full-time employment. We note, however, that this explanation is inconsistent with a finding reported in section IV: Fixing the level of underlying vulnerability, couples are less likely to obtain insurance for the protection of primary earners (and husbands) than for secondary earners (and wives). Because primary earners are already (in most cases) fully employed, secondary earners have greater latitude to increase labor force participation. Consequently, for a given level of vulnerability, one would expect to observe greater protection of primary earners, which is not the case.

To evaluate the sensitivity of our results with respect to possible changes in labor force participation, we consider an extreme alternative assumption: All survivors, whether out of the labor force or employed part-time, shift to full-time employment. We impute full-time earnings based on regressions of earnings on demographic characteristics, estimated separately for fully employed men and women. A survivor's contingent earnings are set equal

17. In this exercise, we assume that the financing for the new house is the same as the continuation financing for the old house. Consequently, on a spouse's death, the decline in home equity equals the reduction in the value of the home, and there is an offsetting increase in nonhousing assets; mortgage payments are unchanged, but other housing expenses fall by 30%.

to the maximum of imputed earnings and actual earnings. Due to familiar sample selection problems, this procedure tends to overstate potential earnings for nonworkers; it therefore understates survivors' financial vulnerabilities. As indicated in Table 10, the estimated frequencies of financial vulnerability are reasonably sensitive to this alternative assumption. This contrasts with the findings of Bernheim et al. (2003). The difference is not particularly surprising, because the ability to alter labor supply is more important for younger workers. Note, however, that substantial uninsured vulnerabilities are reasonably widespread even with this extreme alternative assumption. This, too, is not surprising, given the high levels of vulnerabilities among dual-earner households noted in Table 9.

Not surprisingly, this alternative assumption concerning conditional changes in labor force participation has a larger impact on measured vulnerabilities among younger households. It does not, however, reverse the qualitative conclusions that uninsured vulnerabilities decline with age, or that substantial uninsured vulnerabilities are widespread among young households. In particular, the ratio of means for actual to benchmark household insurance is 0.314 for 22- to 39 year olds, 0.975 for 40- to 55-year-olds, 2.76 for 56- to 70-year-olds, and 16.8 for those over 70; the corresponding ratio of medians is 0.235 for 22- to 39-year-olds, 1.45 for 40- to 55-year-olds, and not defined for older households (because the median benchmark is zero). Moreover, the fraction of secondary earners with severe (significant) vulnerabilities is 15.5% (42.9%) for 22- to 39-year-olds, 7.7% (19.6%) for 40- to 55-year-olds, 5.1% (12.1%) for 56- to 70-year-olds, and 8.8% (11.8%) for those over 70.

In principle, shifts in nonlabor income might also cushion the financial impact of a spouse's death. Our analysis makes no allowance for this possibility. Presumably, the most important source of potential support is assistance from relatives. Bernheim et al. (2003) report that between the first two waves of the HRS, only 6.2% of new widows received any assistance of this type. Between the second and third waves, the figure was 7.5%; and between the third and fourth waves, it was only 2.5%. In addition, support may have been modest or temporary in many of these cases. Consequently, there is little evidence that external support payments are significant in practice. Although the HRS only provides information on older widows, we conjecture that these figures are lower for young widows (except perhaps where young children are present).

As emphasized in section III, the SCF collects information only on total life insurance for each household and not on the proportions attributable to each spouse. Our base-case results apportion this insurance using the imputation method described previously. If in some cases these imputations falsely imply that insurance is held to protect spouses who are not financially vulnerable, then our calculations will overstate uninsured vulnerabilities for at-risk individuals. To examine the potential significance of imputation error, we consider an extreme alternative assumption: All life insurance is held to protect the secondary earner (who is, in general, considerably more vulnerable). As indicated in Table 10, the effects on our results are minimal. This is not surprising for two reasons: Total household life insurance is typically small relative to benchmark insurance, and, in any case, our imputations attribute most life insurance to the primary earner.

If important economic variables are measured with error, our calculations may overstate the thickness of the upper and lower tails of the distribution of $IMPACT$, thereby exaggerating the frequencies of significant and severe financial vulnerabilities. Measures of household assets tend to be particularly noisy. However, as illustrated in the final two rows of Table 10, our findings are not sensitive to moderate changes in the values of wealth (a 20% increase or decrease).

As mentioned in section II, we have ignored the possibility that remarriage might cushion the impact of a spouse's death. Bernheim et al. (2003) report that remarriage occurs with low frequency among the HRS sample, but it is probably more common for the younger households included in the SCF. Because it is difficult to model the consequences of remarriage, we did not conduct pertinent sensitivity analysis. Consequently, we cannot rule out the possibility that low levels of insurance among young households are attributable to the expectation that a survivor will benefit financially from remarriage. However, this hypothesis strikes us as odd. Even young individuals must consider the risk that

they may be unable or unwilling to remarry, the value of retaining the option not to remarry, and the effect of financial status on remarriage prospects.

Although our measures of uninsured financial vulnerabilities are sensitive to certain critical assumptions, it is important to emphasize that the poor correlation between coverage and vulnerability is robust. From Table 10, we see that the change in the estimated incidence of vulnerability is largest when we assume that there are no economies of shared living, that housing expenditures are completely flexible, or that all survivors work full time. However, for all of these alternative scenarios, regressions analogous to those reported in section IV continue to indicate that households with greater vulnerabilities (measured by benchmark insurance) typically have no more coverage on average, at any age, than those with greater vulnerabilities.

VI. CONCLUSIONS

Using the 1995 SCF and an elaborate lifecycle model, we have quantified the potential financial impact of each individual's death on his or her survivors, and we have measured the degree to which life insurance moderates these consequences. We have found that life insurance is essentially uncorrelated with financial vulnerability at every stage of the life cycle. This finding is consistent with the hypothesis that life insurance bears little relation to needs at the time of purchase; however, it tends to refute the hypothesis that households purchase long-term contracts with initially appropriate insurance coverage but fail to adjust this coverage over time as their circumstances change.

The impact of insurance among at-risk households is modest, and substantial uninsured vulnerabilities are widespread, particularly among younger couples. Nearly two-thirds of secondary earners between the ages of 22 and 39 have significant financial vulnerabilities (projected reductions in living standards exceeding 20%), and nearly one-third have severe vulnerabilities (projected reductions exceeding 40%). Moreover, only one in five of these at-risk households held sufficient life insurance to avert significant or severe financial consequences. Combining all age groups, 56% of secondary earners and 6% of primary earners would have experienced significant or severe declines in living standard on the death of a spouse in the absence of life insurance. Actual life insurance holdings reduce these figures to, respectively, 42% and 5%. Thus, the overall impact of life insurance holdings on financial vulnerabilities among at-risk households is modest. Roughly two-thirds of poverty among surviving women and more than one-third of poverty among surviving men results from a failure to ensure survivors of an undiminished living standard through insurance.

We have also provided evidence of a significant gender bias in life insurance holdings. Specifically, for any given level of financial vulnerabilities, couples provide significantly more protection for wives than for husbands. For example, couples with severely at-risk wives on average hold $166,628 of life insurance, whereas couples with severely at-risk husbands on average hold only 15% of this amount ($24,827).

APPENDIX: DATA IMPUTATION

Nonasset Income

Our calculations require data on each spouse's past and future covered earnings as well as future total (covered and uncovered) earnings. We assume that all earnings are covered. For respondents who were working at the survey date, we have 1995 self-reported labor earnings. To impute past and future earnings we use a model that assumes that the cross-section age-earnings profile for fully employed workers remains constant through time. We allow real wages for all ages to grow over time, using the historic Social Security real wage growth for past years and a 1% overall real wage growth factor for future years. (In our robustness analysis, was also look at a 0% and 2% overall real wage growth factor for future years.) In estimating past earnings we assume that the first year of employment is the maximum of 1951 and the year the person was age 22.[18] Households in which one of the spouses was temporarily not working, as opposed to out of the labor force, were dropped from our sample.

The SCF provides information on other kinds of nonasset income. We treat some of these income sources, such as veteran's benefits, Social Security, disability income, welfare, child support, and regular help from friends or relatives, as nontaxable. Except for Social Security disability income and child support, we assume these income streams continue, with full adjustments for inflation, until the respondent's death. Social Security disability income is assumed to end at age 62, when the recipient becomes eligible for Social Security retirement benefits. We divide child support received by the number of children to obtain

18. For workers who were under age 22 in 1995, we assume that 1995 was their first year of employment.

child support per child and assume it is received until the child in question reaches age 18.[19] We treat other kinds of special receipts, such as income from trust funds and royalties, as taxable. We assume they will be received for 10 years beyond the survey date and that the payments will be constant in nominal terms. Relatively few respondents receive these kinds of income flows, and the amounts are generally small relative to average earnings. We assume that SCF respondents retire at their stated intended ages of retirement or age 70, whichever is smaller. For those who fail to say when they will retire, we use age 65.

Pension Plans, Retirement Accounts, and Social Security

The SCF provides information on nominal benefits currently received from defined benefit pension plans as well as expected nominal benefits for future pension recipients. We assume that all pensions are indexed to inflation and that a surviving spouse would receive 100% of the monthly benefit or lump-sum distribution. We further assume that employer-sponsored defined contribution plans and all private retirement accounts (IRAs and Keoghs) provide for tax-deductible contributions and tax-deferred accumulation. Contributions in all future years up to age 59 are set equal, in real terms, to contributions in the survey year. If total contributions are greater than the legal limits ($30,000 or 25% of income), contributions are truncated. The proportion given by the employer remains constant. Any contributions (by the employee or the employer) over the legal limit are included in employee nondeductible and tax-favored contributions.

The SCF contains information on IRA account balances but not annual contributions. We impute contributions based on tobit regressions from the Consumer Expenditure Survey. Contributions are calculated as a function of marital status, work status, age, earnings, and family size.

If an individual is already receiving Social Security benefits, we assume that benefits have already started. Otherwise, we impute the initial age of benefit receipt as follows. If the individual is still working, we assume that benefits will start at his or her projected retirement age (but not earlier than age 62). If the individual is retired, we use the reported start date for those currently receiving benefits; for those not yet receiving benefits, we assume benefits will start at age 62 for those currently under 62, and at the current age for those over 62. In all cases, the initial age of benefit receipt is between 62 and 70. For respondents currently receiving Social Security disability benefits, we assume that they switch to retirement benefits at age 62.

Our calculations also require information on the age at which individuals begin to receive private pension benefits. For those not yet receiving benefits, we use the age at which the individual expects benefits to begin as reported in the SCF.

Individuals with previous marriages lasting more than 10 years and ending in divorce or separation and individuals with previous marriages lasting more than 9 months and ending in the spouse's death are eligible to receive Social Security benefits based on the earnings history of their prior spouse. This presents a problem because we do not have any information about prior spouses. We assume that all such individuals receive benefits based on either their own earnings history or that of their current spouse.

Housing

Our calculations require information on a variety of specific housing expenditures, including mortgages, home insurance premiums, property taxes, and other recurring expenses. Association fees, homeowner or condo/co-op/townhouse association fees, and rent on the site for households owning and living in mobile homes are added to home insurance premiums to form recurring house expenditures. When one owns part of a farm, one is classified as a homeowner. The rent paid is then also added to the insurance premium. The SCF does not contain information on home insurance premiums. We imputed annual home insurance premiums by multiplying the home value by 0.0025.

If the mortgage payment (minus property taxes and insurance premium if respondent states these are included in their payments) is negative then the observation is dropped. If the annual property tax is greater than 5% of the home value, the observation is dropped.

With regard to mortgages, the SCF reports the balance remaining, the number of years remaining, the interest rate, and the payment. To ensure consistency, we imputed the balance remaining on the mortgage based on the years remaining and the interest rate and payment.

In some instances, rental payments reported in the sample include heat and electricity expenses; in such cases, respondents were not asked separately about these utility payments. We apportion the reported number into separate components by assuming that the ratio of rent to utilities is the same for these respondents as the average ratio computed from the HRS. If rent includes all utilities, rent is set to $0.77 * \text{rent}$. If rent includes some utilities, rent is set to $0.89 * \text{rent}$. We have no information on utility expenditure if it is not included in rent. The SCF does not include any information concerning property taxes paid on second homes. We assume that this property is taxed at the same rate as the primary home. Finally, we set monthly rental payments equal to zero for the few respondents who report that they live in a house or apartment that they neither rent nor own.

In addition, for our base case we assume that all households plan to remain in the same house before and after retirement. One test of robustness allows widows and widowers to move to smaller homes. For this case we assume that the move to homes that are 70% of the size of their previous homes.

Other Variables

For confidentiality reasons, the SCF does not report the respondent's date or month of birth or state of residency. We assume that each respondent was born on June 15. For the purposes of computing state taxes, we use Massachusetts law. We set the maximum age of life to 95 for all individuals. Many households have adult children living with them. For the purposes of this project, only children age 18 or under are included.

19. The SCF reports only the sum of child support and spousal support. However, we confine our attention to couples, 98% of which are married. Because spousal support generally ends on remarriage (and also declines somewhat on average when individuals become unmarried partners), we can safely assume that the entire reported amount is child support.

We assume for all respondents a fixed amount for funeral expenses, which is set equal to the median of the reported expenses ($5,000) for HRS spouses who actually died in 1991 (90 observations). The HRS reports information on actual funeral expenses and legal fees of deceased spouses. We set intended bequests equal to zero.

The SCF allows mortgages to end with a balloon payment. When there is a balloon payment, we assume that they refinance for the amount of the balloon payment with a 15-year mortgage (8% interest rate). There is no space in ESPlanner for future mortgages, so these are included in special expenditures. Interest payments on the first home are included in deductible special expenditures. Payments on the balance are included in nondeductible special expenditures. Nondeductible special expenditures also include child support or alimony payment and support to other family members. These are assumed to be paid in the current year and the next four years (a total of five years).

As a measure of a household's net worth, we use total nonhousing assets minus total nonhousing liabilities. Total nonhousing assets include checking and saving accounts, money market funds, certificates of deposit, government saving bonds, T bills, stocks, mutual funds, investment trusts, business equity, bonds, bond funds, real estate other than primary and vacation homes, the cash value of life insurance policies, and some miscellaneous items. Total nonhousing liabilities include personal loans, student loans, credit card balances, car loans, installment loans, and other nonhousing debt. Housing debt (mortgages and equity lines of credit) are considered separately. We assume that apart from mortgages and other outstanding housing debt, households cannot borrow against future income. For our base case, we use a 3% rate of inflation and a 3% real pretax rate of return.

We assume respondents' borrowing limit is set equal to zero.

The expected change in living standard after retirement or in case of death of one partner is set equal to zero.

REFERENCES

American Council of Life Insurance. *1994 Life Insurance Fact Book*. Washington, DC: American Council of Life Insurance, 1994.

———. *1999 Life Insurance Fact Book*. Washington, DC: American Council of Life Insurance, 1999.

Auerbach, A. J., and L. J. Kotlikoff. "Life Insurance of the Elderly: Its Adequacy and Determinants," in *Work, Health, and Income among the Elderly*, edited by G. Burtless. Washington, DC: Brookings Institution, 1987.

———. "Life Insurance Inadequacy—Evidence from a Sample of Older Widows." National Bureau of Economic Research Working Paper No. 3765, 1991a.

———. "The Adequacy of Life Insurance Purchases." *Journal of Financial Intermediation*, 1(3), 1991b, 215–41.

Bernheim, B. D. "The Economic Effects of Social Security: Towards a Reconciliation of Theory and Measurement." *Journal of Public Economics*, 33(3), 1987, 273–304.

Bernheim, B. D. L. Forni, J. Gokhale, and L. J. Kotlikoff. "The Mismatch between Life Insurance Holdings and Financial Vulnerabilities: Evidence from the Health and Retirement Survey." *American Economic Review*, 93(1), 2003, 354–365.

Caplin, A. "The Reverse Mortgage Market: Problems and Prospects," in *Innovations in Managing the Financial Risks of Retirement*, edited by Z. Bodie, B. Hammond, O. Mitchell, and S. Zeldes. Pension Research Council, Wharton School, University of Pennsylvania, 2001.

Gokhale, J., L. J. Kotlikoff, and M. Warshawsky. "Comparing the Economic and Conventional Approaches to Financial Planning," in *Essays on Saving, Bequests, Altruism, and Life-Cycle Planning*, edited by L. J. Kotlikoff. Chicago, IL: MIT Press, 2001, 489–560.

Holden, K. C., R. V. Burkhauser, and D. A. Myers. "Pensioners' Annuity Choice: Is the Well-Being of Their Widows Considered?" University of Wisconsin Institute for Research on Poverty Discussion Paper 802–86, 1986.

Hurd, M. D., and D. A. Wise. "The Wealth and Poverty of Widows: Assets before and after the Husband's Death," in *The Economics of Aging*, edited by D. Wise. Chicago: University of Chicago Press, 1989, 177–99.

Kennickell, A. "Imputation of the 1989 Survey of Consumer Finances: Stochastic Relaxation and Multiple Imputation." Mimeo, Board of Governors of the Federal Reserve System, 1991.

Kotlikoff, L. J., and A. Spivak. "The Family as an Incomplete Annuities Market." *Journal of Political Economy*, 89(2), 1981, 372–91.

Lundberg, S. "Family Bargaining and Retirement Behavior," in *Behavioral Economics and Retirement Policy*, edited by H. Aaron. Washington, DC: Brookings Institution, 1999.

Mitchell, O. S., J. M. Poterba, M. J. Warshawsky, and J. R. Brown. "New Evidence on the Money's Worth of Individual Annuities." *American Economic Review*, 89(5), 1999, 1299–318.

Nelson, J. A. "Methods of Estimating Household Equivalence Scales: An Empirical Investigation." *Review of Income and Wealth*, 38(3), 1992.

Ringen, S. "Households, Standard of Living, and Inequality." *Review of Income and Wealth*, 37, 1991, 1–13.

Venti, S., and D. Wise. "Patterns of Housing Equity Use among the Elderly," in *Innovations in Managing the Financial Risks of Retirement*, edited by Z. Bodie, B. Hammond, O. Mitchell, and S. Zeldes. Pension Research Council, Wharton School, University of Pennsylvania, 2001.

Williams, R., D. Weiner, and F. Sammartino. "Equivalence Scales, the Income Distribution, and Federal Taxes." Technical Paper Series 1999–2, Congressional Budget Office, October 1998.

Yaari, M. "Uncertain Lifetime, Life Insurance, and the Theory of the Consumer." *Review of Economic Studies*, 32, 1965, 137–50.