

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

Volume Title: Themes in the Economics of Aging

Volume Author/Editor: David A. Wise, editor

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-90284-6

Volume URL: <http://www.nber.org/books/wise01-1>

Publication Date: January 2001

Chapter Title: Are the Elderly Really Over-Annuitized? New Evidence on Life Insurance and Bequests

Chapter Author: Jeffrey Brown

Chapter URL: <http://www.nber.org/chapters/c10322>

Chapter pages in book: (p. 91 - 126)

Are the Elderly Really Over-Annuitized?

New Evidence on Life Insurance and Bequests

Jeffrey R. Brown

3.1 Introduction

It is well established in the economics literature that annuities ought to be of substantial value to life-cycle consumers who face an uncertain date of death. Yaari (1965) proved that a life-cycle consumer with an uncertain lifetime and no bequest motives would find 100 percent annuitization the optimal investment. More recent work has quantified the potential utility gains to such a life-cycle consumer. For example, a sixty-five-year-old male life-cycle consumer would be willing to give up nearly one-third of his wealth to gain access to an actuarially fair market for annuities (Mitchell et al. 1999).

Buying a life insurance contract is analogous to selling an annuity. Life insurance is generally viewed as an appropriate product for working-age individuals who seek to protect their families against the loss of future labor earnings (Lewis 1989). However, it appears to serve little purpose in the portfolio of a retired life-cycle consumer whose sole concern is self-financing retirement out of his or her accumulated wealth. With no labor earnings to insure, an elderly individual should be purchasing annuities in order to provide a certain consumption stream in retirement, not selling annuities through the purchase of life insurance. Even if the individual

Jeffrey R. Brown is assistant professor of public policy at the John F. Kennedy School of Government at Harvard University, and a faculty research fellow of the National Bureau of Economic Research.

For helpful comments and discussions, I wish to thank Anne Case, Courtney Coile, Peter Diamond, Jon Gruber, Jerry Hausman, Jim Poterba, Harvey Rosen, Scott Weisbenner, participants in the MIT Public Finance lunch, and participants in the NBER Aging Conference. The author gratefully acknowledges the financial support of the NBER and NIA aging Fellowship and the National Science Foundation.

wishes to leave a portion of wealth to his or her heirs in the form of gifts or bequests, this can be achieved by investing this portion of wealth in ordinary bonds or other non-annuitized assets. In fact, if life insurance premiums were higher than actuarially fair, holding riskless bonds would dominate life insurance as a form of wealth transfer.

Yet elderly households in the United States overwhelmingly hold life insurance, while only a small fraction hold privately purchased annuity contracts. In the Asset and Health Dynamics among the Oldest Old (AHEAD) survey, which consists of households aged seventy and older, privately purchased annuity contracts (excluding private pensions) are held by fewer than 8 percent of couples, while 78 percent of couples age seventy and older own a life insurance policy on at least one member. According to the Life Insurance Ownership Study (Life Insurance Market Research Association [LIMRA] 1993), ownership of individual (non-group) life insurance policies is actually higher among the group aged sixty-five and older than any other age cohort. While this difference is offset by much lower coverage by group (usually employer-based) policies, the overall incidence of coverage among the elderly is quite high by any measure.

Two major alternative hypotheses have been explored in the literature to explain the patterns of life insurance coverage among the elderly. Auerbach and Kotlikoff (1987, 1989) explored the idea that married couples use life insurance to reallocate annuity streams across survival states of the couple. However, they found virtually no support for the notion that older households were using life insurance to protect potential widows and widowers against severe drops in living standards upon the death of the other spouse.

The second hypothesis, suggested by Bernheim (1991), is that life insurance is being held by elderly households to offset an excessive level of mandated annuitization in the form of Social Security. He estimates that 25 percent of elderly households have too much of their wealth annuitized and that they are using term life insurance to sell these annuities in order to leave a bequest. To the extent that this “annuity offset model” is true, it has at least two important implications. First, this would be indicative of very strong bequest motives, which is an issue of perennial controversy in the economics literature (e.g., Kotlikoff and Summers 1981; Modigliani 1988; Hurd 1987; Laitner and Juster 1996). Second, if individuals are over-annuitized due to these strong bequest motives, this would indicate a potential welfare gain from lessening the extent of the mandated annuitization. This is potentially important in the debate about whether individuals would be required to annuitize individual account accumulations as part of a reformed Social Security system. If a significant fraction of households are over-annuitized, allowing individuals some discretion over the disposition of the assets in their individual accounts could be welfare enhancing.

This paper reexamines the annuity offset model using more recent and better data than were available for the original empirical tests. The four empirical implications of the model that this paper tests are (1) that no individual would hold both term life insurance and private annuities, (2) that the level of Social Security benefits and term life insurance ownership should be negatively correlated, (3) that term life insurance should behave as an inferior good because it is a negative annuity and annuities are normal goods, and (4) that individuals who hold term life insurance must have a Social Security benefit in excess of desired retirement consumption. These implications will be explained in more detail in the next section. This paper presents results that are inconsistent with all four of these empirical implications, and thus concludes that life insurance coverage is *not* a good indicator of the extent of over-annuitization.

This paper proceeds as follows: Section 3.2 summarizes the annuity offset model as posited by Bernheim (1991), and outlines the empirical implications of the model to be tested. Section 3.3 presents and critiques the empirical results from the previous literature, with particular attention on the distinction between types of life insurance. Section 3.4 discusses the data used in this paper, from the AHEAD study. Section 3.5 presents empirical results. Section 3.6 discusses some alternative explanations for life insurance holdings among the elderly, and section 3.7 concludes.

3.2 The Annuity Offset Model of Life Insurance Demand

The basic insight behind of the annuity offset model of life insurance demand is that individuals can purchase term life insurance in order to sell a government mandated annuity. Bernheim suggests a simple two-period model that demonstrates this point. Assume that an individual possesses total wealth W_0 , which the individual is able to divide between two types of investments. It is assumed that the investment decision is taking place after consumption in period 0 has already occurred. The first type of asset, A , is a life-contingent annuity contract that yields a return of α in period 1 if the individual is alive, and 0 otherwise. The second type of asset, B , is a traditional (bequeathable) financial asset that yields a return of β in period 1 regardless of whether the individual is alive. If the individual lives, his or her period 1 resources are $W_1 = \alpha A + \beta B$. If the individual dies, his or her heirs receive βB . Because actuarially fair annuities pay a “mortality premium” equal to the probability of dying, $\alpha > \beta$. Utility of the individual is assumed to be a function of total resources and bequeathable resources in period 1, $U = U(B, W_1)$. Call A^* and B^* the quantities of the two assets that correspond to the optimal division of total wealth, subject to the constraint that $W_0 = A + B$.

Now suppose the government confiscates A^s in period 0 and returns αA^s in period 1, conditional on the individual’s survival. In other words, the government mandates a minimum level of annuitization. If $A^s < A^*$, then

the individual simply decreases his or her private purchase of annuities by an amount equal to A^g , or alternatively, the individual buys private annuities in the amount of $A^* - A^g$. If $A^g > A^*$, then the individual wishes to sell annuities. This can be done through the purchase of a term life insurance contract. When markets for annuities and life insurance are actuarially fair, then the government mandate has no effect on the individual's division of wealth between A and B . Private insurance contracts offset the government annuity dollar for dollar. If insurance is not actuarially fair, then the offset is less than dollar for dollar, but the basic story is unchanged. Individuals who wish to hold more annuities than the government mandates will own private annuity contracts. Individuals who wish to hold less in annuities will own private life insurance contracts. No individual will hold both private annuities and life insurance, since they are offsetting transactions, each with a positive load factor. Some individuals will hold neither, if A^g is sufficiently close to A^* .

There are four major empirical implications that must hold if the annuity offset model is the reason that the elderly hold life insurance. These are as follows.

1. *No individual will hold both private annuities and private term life insurance contracts.* Given the existence of significant load factors in annuity markets (Mitchell et al. 1999), no one would rationally purchase annuities above the actuarial cost only to sell them back below the actuarial cost.

2. *An increase in the level of the mandated annuity will increase the demand for term life insurance.* Recall that an individual will hold term life insurance in the amount of $\max\{0, A^g - A^*\}$. Holding W_0 fixed, an increase in A^g will increase the total amount of life insurance coverage among those who already hold it. It will also cause some individuals who did not hold life insurance before to purchase it.

3. *Term life insurance will behave as an inferior good.* If B and W_1 are normal goods, then an increase in the individual's total resources will increase the demand for annuities. This is because a person with more resources will wish to buy more annuities in order to increase retirement consumption. Since term life insurance is a negative annuity, an increase in the demand for annuities corresponds to a decrease in the demand for term life insurance. Therefore, term life insurance will decline with total resources, and thus behave as an inferior good.

4. *The Social Security annuity income flow must exceed consumption for life insurance owners.* If an individual is over-annuitized due to bequest motives, it must be because his or her desired consumption is less than the annuity income from Social Security. So long as optimal consumption exceeds the level of the Social Security benefit, there is no need to offset Social Security. Rather, one would want to supplement Social Security through the purchase of private annuities. An equivalent way to state this

implication is that an individual who purchases life insurance to offset an annuity will not consume out of his or her non-Social Security resources. The individual will save these resources for bequests, and will in fact supplement this bequest with the term life insurance.

It should be noted some of the reasons an individual might be over-annuitized have nothing to do with bequest motives. Hurd (1987) points out that when an individual's optimal consumption path is constrained by an exogenously given annuity stream, he or she may be willing to give up annuitization at an actuarially fair rate in order to loosen this liquidity constraint. This is especially likely if the individual has little nonannuitized wealth. However, over-annuitization in this case is driven by a desire to reallocate consumption across one's lifetime, not to reallocate between consumption and bequests. Another example is the case in which an individual wishes to hold a buffer stock of assets to cover unforeseen expenditure shocks (e.g., health expenditures). In such a situation, the individual may wish to hold some of his or her wealth in a nonannuitized form. Once again, the role of the nonannuitized wealth in this case is still to provide for own consumption, not to leave a bequest to one's heirs. In this case, the way to undo the excessive annuitization, however, is not to purchase life insurance, since these proceeds will be unavailable for future consumption. Rather, the individual would wish to purchase insurance against the risky future event (e.g., health insurance) or alter his or her saving behavior in order to provide for a buffer stock. The tests that I propose in this paper are meant to test for over-annuitization that derives from bequest motives, not these other factors.

3.3 Discussion of Previous Empirical Results and Data Contamination

Bernheim tested the first three implications of the annuity offset model using the 1975 cross-section of the Retirement History Survey (RHS), and found support for two of them. The most robust finding was that higher Social Security benefits were correlated both with a higher probability of owning life insurance, and with the amount of coverage conditional on owning a policy. His interpretation of this finding is that individuals are using the life insurance to offset excessive levels of Social Security. He also found some evidence to suggest that life insurance coverage was a decreasing function of lifetime resources, which is consistent with the "inferior good" implication, though this finding was not robust across specifications.

The first implication, that no person would hold both life insurance and annuities, was clearly at odds with the data, because 36 percent of the RHS sample reported both in-force life insurance holdings and the receipt of pension annuities. He attributes this latter result to data contamination,

namely, the fact that there is no way in the RHS to distinguish whole from term life insurance. Because much of Bernheim's analysis was focused on trying to overcome this data handicap, it is useful to discuss the types of life insurance in more detail.

3.3.1 Term versus Whole Life Insurance

The distinction between term and whole life insurance is quite important to the annuity offset model. The difference between the two policy types is quite simple, but has important economic implications. Term life insurance contracts provide insurance protection for a specified limited period. The face amount of the policy is payable to the beneficiaries only if the insured dies within the term specified. Common term periods include one-year, five-year, ten-year, and twenty-year. Most term policies have options allowing an individual to guarantee renewability at the end of the term specified. This means that an individual is not at risk for losing coverage if he or she is diagnosed with a serious health problem, so long as he or she pays the contract premium. Because the price of a term insurance contract is a function of the probability of the individual's dying during that term, premiums are an increasing function of the insured's age.

Whole life policies, on the other hand, are not limited in duration, but rather protect "the whole of life" (Graves 1994). Unlike term insurance contracts, which represent pure insurance, the typical whole life contract is a combination of insurance and tax-deferred savings. The typical "ordinary life" product has fixed- (nominal) level premiums and a fixed (nominal) death benefit or face value. As demonstrated in figure 3.1, the cash values of these policies rise over time, while the pure insurance component declines. The standard practice among life insurers is for the cash value to equal the face value by age ninety-five or 100 (Graves 1994). According to the 1995 Survey of Consumer Finances (SCF), the median whole life insurance policy held by individuals aged seventy and up had a cash value that was 67 percent of the face value. This means that only one-third of the reported face value of whole life policies represents insurance. Most whole life policies have provisions that enable the individual to borrow against the cash value of his or her policy, and thus provide some degree of liquidity. Importantly for the annuity offset model, the cash value of a policy is *not* a negative annuity, but rather represents a nonannuitized financial asset, much like a saving account. While it is true that the cash value of a life insurance policy may be left to one's heirs as a bequest, a large cash value policy would not be indicative of over-annuitization any more than would the holding of a large savings account.

As important as this distinction may be between term and whole life insurance, previous empirical work on the elderly was unable to distinguish between them. The RHS provided data only on the total face value




Fig. 3.1 Proportion of saving and insurance elements in an ordinary whole life insurance contract

Source: Graves (1994)

Note: Calculated based on issue age of twenty-five using 1980 Commissioners Standard Ordinary female lifetable and a 4.5 percent rate of interest.

of all life insurance policies. Thus, to the extent that ownership of whole life more closely resembles tax-deferred savings than it does insurance, previous researchers were unable to disentangle these two potentially different effects. For example, suppose high-income individuals are more likely to purchase whole life insurance as a form of tax-deferred savings. Because these individuals are high income, they also receive a higher Social Security benefit at retirement. This could lead to a spurious correlation between total life insurance holdings and the level of the Social Security benefits. As we shall see, this commingling of insurance and tax deferred savings has an important impact on the results.

3.3.2 Group versus Individual Coverage

Another relevant distinction between types of life insurance that may be important to this model is between group and individual coverage. Group life insurance policies are commonly provided through employers or unions. An example of a typical group life policy is one that insures an employee for a fixed multiple of his or her salary. Individual contracts, on the other hand, are purchased directly from the insurance company, most often through an insurance agent or broker.

The primary distinction between these policies is that individual life

coverage is clearly a “choice” variable, whereas group coverage is often automatic with employment. While in many instances group coverage simply substitutes for individual coverage that would have been purchased anyway, it will also cover some individuals who may have chosen to hold no life insurance if not covered through their employer. Another reason this distinction is relevant is that, since group coverage is usually tied to employment, its purpose is often to protect an employee’s family from the loss of future earnings. This purpose for holding life insurance is distinctly different from using life insurance to offset a retirement annuity.

Group coverage is less common among retired elderly households, since most group coverage is tied to employment. Neither the RHS, nor the AHEAD data used in this study, allow for this distinction between group and individual coverage. However, by making use of information about the current employment status of an individual, it is possible to extract some information about the effects of this difference.

3.3.3 Previous Empirical Results

Previous empirical support of the annuity offset model rests on two key results. The first is that there exists a strong positive correlation in a cross-section of households between the level of Social Security benefits and the holdings of life insurance. Bernheim estimates that approximately 25 percent of households own term life insurance, and based on his model, are therefore over-annuitized. His central results indicate that they are using term insurance to offset Social Security by roughly twenty cents on the dollar.

Second, Bernheim finds mixed evidence to suggest that a portion of the total life insurance holdings are negatively correlated with total lifetime resources, and thus represents an inferior good. Importantly, in his most direct specifications, he finds that life insurance coverage is actually increasing with resources for married couples with children. Only when he imposes more structure on the problem to overcome problems of data contamination does he find a consistently negative and significant resource effect. However, this approach is unable to identify directly the effect of resources on term life insurance, and instead relies on modeling total holdings as the sum of two separate processes (one representing term and one representing whole, but each unidentified in advance) and testing the sign of various coefficient combinations.

The difficulty with these results is that the potential for bias is quite high due to the inability to identify directly the term insurance component of total life insurance holdings. Suppose that individuals purchase insurance during their working lives in order to protect their spouses and children from the loss of their human capital in the event of an early death. Individuals can choose between term and whole life insurance to meet this insurance need. The annual premium on a whole life contract is higher than the

premium on a term life contract because some of the additional premium essentially goes into a savings account that benefits from tax-deferred inside buildup. Because of this, the whole life contract is more attractive, all else equal, to an individual in a higher marginal tax bracket. Therefore, high earners (who therefore face higher marginal rates) have the most to gain from purchasing whole life contracts. High earners will also be paying more in Social Security payroll taxes, and will thus have a higher benefit upon retirement. Thus, to the extent that whole life contracts held by the elderly represent “residue” from decisions made early in life to protect human capital, this would induce a spurious positive relationship between Social Security benefits (SSB) and whole life insurance coverage.

Two pieces of evidence suggest that this scenario is a strong possibility. First, according to a life insurance ownership study conducted by LIMRA, the median age of the oldest life insurance policy held by individuals aged seventy and up was forty-two years, suggesting that most policies were in fact purchased during the individual’s working life. Second, there is a clear positive relationship between ownership of whole life insurance and income during one’s working life. For example, if we focus on working-age individuals (aged twenty-two to sixty-five) in the 1995 SCF, we find that only 20 percent of individuals with incomes under \$30,000 own a whole life insurance policy. Of those with annual incomes between \$30,000 and \$60,000, 33 percent own a whole life policy. Nearly half (48 percent) of those earning over \$60,000 per year own a whole life policy. This relationship is not biased by the age-earnings profile, as a nearly identical trend emerges when one examines ownership patterns conditional on age. Thus, whole life insurance ownership during one’s working life is clearly correlated with income, and thus with OASDI contributions. If individuals continue to hold these policies after retirement, this will lead to a positive correlation between the level of Social Security benefits and whole life insurance ownership. Newly available data allow for a separation of total life insurance into whole versus term life policies, and as such provides a more direct test of the model.

A second potential source of spurious correlation is that some individuals in the Bernheim study were still in the work force. His 1975 RHS sample was comprised of individuals aged sixty-five to sixty-nine. According to Department of Labor statistics, in that year the labor force participation rate of individuals aged sixty-five to sixty-nine was 31.7 percent. High labor force participation can lead to bias in the annuity offset test for two reasons. First, individuals still in the workforce still have positive human capital to protect, and may hold life insurance for this reason. If these individuals also have higher Social Security benefits due to their strong attachment to the labor force, this could induce a positive correlation between benefits and insurance coverage. Second, employed workers are more likely to be automatically covered by group insurance plans.

Therefore, even if the person has no demand for insurance, he or she may be insured. If employed workers are more likely to have higher Social Security benefits, a spurious positive correlation would result.

3.4 Data and Methodology

This paper uses data on elderly households from the first wave of the AHEAD survey. Fielded in 1993–94, this survey collected detailed financial and demographic data on community-based individuals born in 1923 or earlier, so they were aged seventy and up at the date of the survey. The questionnaire collects detailed information on economic and demographic variables, health, work status, and importantly for this study, life insurance coverage.

There are several advantages to the use of this data over the earlier work completed using the RHS. First, the data allow for the important distinction between term life and whole life insurance coverage. While they still do not permit the decomposition of whole life into its cash versus insurance value, the fact that we can distinguish between pure term policies and whole policies represents an important improvement over the total face value of all insurance. Second, because the data consist of individuals aged seventy and up, nearly all of them are retired. This is important both because this means that the individuals no longer carry life insurance to protect against the loss of human capital, and because it is significantly less likely that the individual will be covered by a group life insurance plan through the employer. Therefore, a test of the annuity offset hypothesis will not be contaminated by work-related reasons for insurance coverage. Third, the data are much more recent than the RHS, which is potentially important due to the clear long-term decline in the life insurance coverage of households over the past three decades (LIMRA 1993). Fourth, because of the advanced age of the cohorts, there are large enough samples to investigate the behavior of widows and couples separately. This may be an important distinction because at least one alternative to the annuity offset hypothesis is relevant to couples but not to widows. This is the notion that elderly couples may use life insurance to reallocate wealth across states of spousal survival.

The primary disadvantage of the AHEAD data is the fact that they do not currently contain information on the earnings histories of respondents. As a result, it will not be possible to replicate precisely the specification of lifetime resources as used in Bernheim's work on this subject. However, the information on current income from Social Security and pension plans is quite detailed, and along with information on financial assets it is possible to construct a very good measure of resources available to the household from the date of the survey forward.

This analysis will focus on two subsets of households in the AHEAD

data set. The first is married couples in which both spouses were interviewed, and for which we therefore have complete information about important characteristics of both spouses. The second set consists of male widowers and female widows (hereafter often referred to collectively as widows), i.e., formerly married individuals who lost their spouses to death. Excluded from this analysis are never married individuals, both because of small sample sizes and because they are less likely to have children or grandchildren to which they may wish to bequeath. Also excluded are single divorcees, due to small sample sizes and the fact that the survey lacks important information about their former spouses. The resulting sample size for married couples ranges from 1,750 to 1,950 households, and from 2,600 to 2,800 widows and widowers. The range of households arises from missing data for some versions of the dependent variable. For example, an individual may state that he or she does not own a whole life policy, but that he or she does own a term life policy with an unknown face value. My decision rule was to include the maximum number of households possible, so this person would be included in the whole life regressions, but excluded from the term life and total life regressions due to missing data. I have conducted extensive checks to ensure that the results were not sensitive to this selection process, and found that the basic results are unchanged.

In order to test for the effect of Social Security benefits and total resources on the holdings of life insurance, I use the following econometric specification:

$$(1) \quad LI_i = \max\{0, \beta_0 + \beta_1 SSB_i + \beta_2 LR_i + \beta_3 X_i + \varepsilon_i\}$$

LI is the face value of life insurance. In some specifications, this will represent total face value, while in others I will limit it to term life or whole life only, in order to account for the cash value bias discussed earlier. SSB represents the annual flow of benefits from Social Security. LR is a vector of characteristics that attempts to capture components of lifetime resources. It includes the variable PVR (present value of resources), which equals the expected discounted present value of resources, including net worth, social security wealth, and pension wealth. Because lifetime earnings records are not yet available in the data set that I use, the LR vector also includes a number of variables which proxy for the effect of lifetime earnings. These include nine occupation indicators and four educational attainment indicators. For specifications involving couples, these indicator variables are all included separately for each spouse. X is a vector of other relevant demographic characteristics, including age, gender, race, and whether the respondents have any living children.

I will show results using two different estimation procedures. For comparability with Bernheim's study, I will first assume the normality of ε and

report results from a tobit specification. One might be concerned about the possibility of heteroscedasticity in the unobservables in the demand for life insurance, which would render the tobit results inconsistent. Therefore, I will also report results using simple ordinary least squares (OLS) regressions with White-corrected standard errors. As the results will indicate, to the extent that heteroscedasticity biases the tobit results, it appears to do so in a direction that favors the annuity offset model. Further specification checks using a censored least absolute deviations (LAD) estimator, or modeling the heteroscedasticity in a multiplicative form, similarly indicate that any such biases tend to work in favor of the annuity offset model. This phenomenon is captured by the OLS estimates, so I limit reported results to tobit and OLS.

Equation (1) closely approximates the main specification used by Bernheim in his test of the annuity offset model in the RHS, with three primary differences. The first is that Bernheim was restricted to using total face value of all life insurance as his dependent variable, whereas the current study can examine whole and term separately. The second difference is in the construction of the measure of total resources. The definition used here, PVR, is net worth plus the present value of future income from Social Security and pensions, and thus represents resources available from today forward. Bernheim's measure was the present value of lifetime earnings plus the present value of Social Security and pensions, and thus represented total lifetime resources. The third difference is that the current study examines behavior of widows and couples in separate regressions. Bernheim ran his model on all households, with appropriate indicators for marital status, but excluded individuals who had been widowed more than six years.

3.5 Results

Table 3.1 presents summary statistics on life insurance ownership among households aged seventy and up in the AHEAD data. Several features of the data are worth noting. First, men are more likely to hold all types of life insurance than are women. Nearly 62 percent of widowed men own a life insurance policy, versus only 49 percent of widowed women. Among currently married couples, 72 percent of men are covered by at least one policy, versus only 55 percent of married women. Looking at term and whole life ownership separately, the same basic pattern emerges, in that men are always more likely to hold insurance than women. In addition, men always hold more insurance conditional on owning, than do women.

A second feature of the data is that most policies tend to be quite small, though the distribution is fairly skewed. The median married household owns a total of \$10,756 of life insurance, a figure that includes all types of

Table 3.1 Life Insurance Coverage in the AHEAD Data

	Widows/widowers		Married Couples		
	Men	Women	Men	Women	Household
Pr(Owns Any LI)	0.6184	0.4868	0.7176	0.5540	0.7791
Amount Owns Any					
Median (\$)	5,000	2,500	9,000	3,000	10,756
Mean (\$)	14,280	5,250	25,481	10,718	31,541
Pr(Owns Term LI)	0.3730	0.3310	0.4174	0.3014	0.4958
Amount Owns Term					
Median (\$)	5,000	2,000	5,000	3,000	7,000
Mean (\$)	9,028	3,841	12,238	7,564	15,313
Pr(Owns Whole LI)	0.2749	0.1505	0.4940	0.2577	0.5659
Amount Owns Whole					
Median (\$)	6,000	3,000	10,000	4,500	14,000
Mean (\$)	18,297	7,189	33,503	13,221	36,119

Source: Authors' tabulations from AHEAD survey, using household weights.

Notes: PR(Owns LI) is the fraction reporting ownership of that life insurance contract type. Amount | Owns LI is the mean or median policy size conditional on ownership.

life insurance on both spouses. Among widowed households, it is even smaller, with a median value of \$5,000 for men, and \$2,500 for women. However, the means are roughly two to three times larger than the medians, which is driven by the fact that a small fraction of households own very large policies. For example, the 95th percentile of total household coverage among married couples (conditional on owning) is \$113,000. The 95th percentile of coverage for male widowers is \$50,000.

The third broad pattern to recognize is that marital status is an important margin along which insurance coverage differs. Married individuals are much more likely to own life insurance than are widows or widowers of the same gender, and hold more of it conditional on owning. There are many reasons that this could be true, including reasons that might bear upon the relative importance of using life insurance to protect a spouse versus providing a bequest. However, a large part of these differences is undoubtedly attributable to differences in the financial status of married versus widowed households, which is not captured in these simple tabulations.

3.5.1 Test of Implication no. 1: No Simultaneous Holdings

The first implication of the annuity offset model, and the one easiest to test in the data, is the notion that no individual would choose to hold life insurance and annuities simultaneously. This is because they are offsetting transactions, each of which may cause the individual to incur transactions costs due to the fact that private insurance markets are not actuarially fair.

This assumption is clearly violated by the data in table 3.2. This is par-

Table 3.2 Cross-Ownership Patterns for Life Insurance and Annuities

	Owns Private Pension		Owns Private Annuity (excludes pensions)	
	Yes	No	Yes	No
<i>Married Couples (total household coverage)</i>				
Owns any life insurance				
Yes	.501	.278	.066	.713
No	.093	.128	.011	.210
Owns term life insurance				
Yes	.332	.164	.038	.458
No	.260	.244	.038	.466
<i>Widows and Widowers</i>				
Owns any life insurance				
Yes	.211	.298	.030	.479
No	.137	.354	.026	.465
Owns term life insurance				
Yes	.133	.192	.039	.309
No	.216	.460	.017	.636

Source: Authors' tabulations from AHEAD survey, using household weights.

Note: Proportion of population holding both products.

ticularly notable if one follows the Bernheim approach of treating annuities from pension plans as voluntarily purchased. Of all married households, 50 percent own both a private pension and some form of life insurance. Among widows and widowers, 21 percent own both private pension annuities and life insurance. There are reasons to suspect that private pensions are not strictly voluntary, especially among those aged seventy and up who were likely covered for most of their careers in traditional defined benefit plans. However, even if we restrict ourselves to privately purchased, nonpension annuities, 6.6 percent of married couples own both. Since only 7.7 percent of the sample own such an annuity, however, this means that 86 percent of those married households who have purchased a private, nonpension annuity also own life insurance.

These numbers are not surprising, since in Bernheim's own sample 36 percent of households, which included both married and widowed individuals, owned both pensions and life insurance. He attributed this finding to data contamination, namely the fact that he was unable to distinguish between term and whole life insurance. If the 36 percent of people holding both were really holding whole life policies with cash values approaching their face values (i.e., they contained very little insurance), then this finding would not be inconsistent with the annuity offset model. However, using the AHEAD data, we can see that this is not the explanation. Roughly one-third of married households own both straight term life in-

insurance policies and a private pension, as do 13 percent of widows. Perhaps the purest test of the model is to use term life insurance and nonpension annuities. In this case, 3.8 percent of couples hold both. Importantly, one-half of all married households that own a nonpension annuity also have life insurance coverage. This is clearly inconsistent with the annuity offset model.

3.5.2 Test of Implication no. 2: Positive Correlation between Insurance and Social Security

The second, and arguably the most important, implication of the annuity offset hypothesis is that there should exist a positive correlation between term life insurance coverage and the level of SSB. The heart of this hypothesis, as outlined in section 3.2, is that when individuals have higher SSB, they want to buy fewer private annuities and more life insurance.

Table 3.3 reports tobit results for equation (1) in the combined sample of widows and widowers. Column (1) reports tobit coefficients for the case in which total face value of all life insurance holdings (term plus whole) is the dependent variable. Column (2) reports coefficients for the OLS specification. Columns (3) and (4) repeat this analysis with the dependent variable limited to term life insurance, and columns (5) and (6) limit the dependent variable to the face value of whole life insurance.

The coefficient on annual SSB is the coefficient of interest for testing this implication of the model. If the annuity offset model is correct, the coefficient should be significantly positive. Looking first at column (1), we can see that this relationship does hold for total life insurance coverage in the tobit specification, with a coefficient of 0.48 that is highly significant. Using the well-known approximation that the marginal effect can be obtained by scaling the parameters by the probability in the uncensored region yields a marginal effect of another dollar of SSB of approximately \$0.22 of life insurance coverage. Column (2) repeats the analysis using OLS, and finds a nearly identical marginal effect of 0.22, though the large (White-corrected) standard errors render this coefficient insignificant.

We can translate the life insurance face value into an annuity flow by dividing by the appropriate annuity factor, i.e., the actuarial present value of a \$1.00 annuity flow. Using a real interest rate of 3 percent, this factor is approximately 10 for the average individual in the AHEAD sample. Therefore, we find that life insurance is offsetting the flow of SSB on the margin by approximately 2.2 cents on the dollar. This offset is much lower than the 10–20 cent offset that Bernheim found because the current sample is of widows and widowers only, while Bernheim's results were for a mixed sample. Results for couples, discussed below, show a somewhat larger offset that falls in the lower end of the Bernheim offset range.

Columns (3) through (6) of table 3.3 make the important distinction between term and whole life insurance. Columns (3) and (4) report the

Table 3.3 Tobit and OLS Results for Widows and Widowers

Dependent Variable Model	Total Tobit (1)	Total OLS (2)	Term Tobit (3)	Term OLS (4)	Whole Tobit (5)	Whole OLS (6)
SSB	0.4751*** (0.1530)	0.2163 (0.1621)	0.0440 (0.1028)	0.0285 (0.0324)	1.2214*** (0.3020)	0.1680 (0.1433)
PVR	0.0093*** (0.0019)	0.0078** (0.0037)	-0.0012 (0.0013)	0.0000 (0.0003)	0.0122*** (0.0032)	0.0068 (0.0035)
Working	2.178.6 (1,855.6)	1,363.0 (1,427.4)	-147.37 (1,250.3)	-114.88 (446.70)	3,936.3 (3,527.2)	1,155.5 (1,315.2)
Age	-441.40*** (78.97)	-99.683** (49.534)	-310.39*** (52.850)	-80.387*** (19.030)	-448.34*** (1672.02)	-16.375 (43.593)
Female	-8,181.7*** (1,366.8)	-4,836.5*** (1,056.9)	-3,257.7*** (917.73)	-1,898.3*** (532.95)	-12,538.7*** (2,635.1)	-2,738.1*** (837.44)
Nonwhite	1,640.2 (1,302.9)	644.77 (619.89)	1,184.1 (848.02)	167.22 (245.36)	-3,500.7 (2,778.8)	140.78 (521.08)
Kids	3,078.4** (1,401.5)	653.55 (685.74)	1,616.1 (932.16)	311.77 (309.70)	3,076.3 (2,858.7)	158.92 (565.06)
N	2,605	2,605	2,738	2,738	2,811	2,811

Source: Authors' calculations from the AHEAD survey.

Notes: Standard errors in parentheses (OLS standard errors are White-corrected). Specifications also include indicator variables for occupation and education. Dependent variable is the dollar value of total, term, or whole life insurance coverage. "SSB" is the annual Social Security benefit, and "PVR" is the present value of resources, including the actuarial value of pensions, Social Security, housing wealth, and financial net worth. "Working" is an indicator variable that equals 1 if respondent is employed.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

results for the case in which only term life insurance is treated as the dependent variable. Relative to the results for columns (1) and (2), the difference is striking. The tobit coefficient falls to 0.04, and is statistically no different from zero. The tobit coefficient for SSB in the whole life insurance specification in column (5), on the other hand, is large and significant—the tobit coefficient on SSB is 1.22 and is highly significant. The OLS results are again similar to the marginal effects that arise out of the tobit coefficients, but are not significant. It therefore appears, at least in the sample of widows and widowers, that the central implication that SSB will be positively correlated with term life insurance coverage does not hold. While there is a significant positive relationship found between total life insurance coverage and Social Security in the tobit specifications, this relationship appears to be driven more by whole life insurance than by term coverage, and even this relationship is not always significant.

As discussed by Bernheim and earlier in this paper, the annuity offset model is really a model about term life insurance. Yet the primary implication of this model, that term life insurance ownership will be an increasing function of the level of SSB, is clearly rejected by the data. The relationship between whole life insurance and Social Security, on the other hand, is much stronger but subject to numerous sources of bias. As discussed, by age seventy, whole life insurance consists primarily of tax-advantaged savings, and therefore does not serve to undo annuitization any more than holding other financial assets. Since the individuals who have the most to gain from the inside build-up associated with whole life policies are people who were in higher tax brackets while younger (and who therefore are also likely to have higher Social Security earnings), the observed relationship may be driven more by tax consequences than by a desire to offset a government annuity.

Table 3.4 repeats the analysis for the male widowers only, and finds a similar pattern. In the specification using total life insurance as the dependent variable, there is a significant positive effect of SSB in the tobit specification, and it is much larger in magnitude. Once again, however, when we decompose total life insurance holdings into the two types, we find that the positive relationship is being driven entirely by whole life policies. The OLS results again show similar, though insignificant, marginal effects. Table 3.5 repeats the same analysis for female widows. In this case, the coefficient on SSB in the total life insurance regression falls short of significance, and is smaller in magnitude than for males. More importantly, the coefficient on SSB in the term life insurance specifications continues to be small and insignificant. The only specification in which the SSB coefficient is significant is for the tobit specification in the whole life insurance regression.

Table 3.6 presents results for married couples. Focusing on column (1), we again see that the coefficient on SSB in the “total life insurance”

Table 3.4 Tobit and OLS Results for Male Widowers Only

Dependent Variable Model	Total Tobit (1)	Total OLS (2)	Term Tobit (3)	Term OLS (4)	Whole Tobit (5)	Whole OLS (6)
	SSB	1.1216*** (0.4331)	0.7016 (0.5906)	-0.0905 (0.3289)	0.0536 (0.0823)	2.4462*** (0.6940)
PVR	0.0159*** (0.0044)	0.0136 (0.0092)	0.0006 (0.0027)	-0.0001 (0.0007)	0.0155*** (0.0060)	0.0109 (0.0085)
Working	-7,381.7 (6,359.4)	-5,561.8 (3,414.7)	-13,058** (5,232.0)	-4,333.8** (2,107.5)	-824.04 (9,310.5)	-2,672.8 (2,982.9)
Age	-612.18** (265.73)	-186.26 (228.62)	-727.18*** (204.82)	-242.08*** (82.403)	-290.26 (417.16)	58.199 (195.12)
Nonwhite	-2,853.6 (4,496.5)	-495.99 (1,382.7)	-3,212.2 (3,319.4)	-1,170.1 (780.10)	-12,367 (7,414.0)	-983.99 (1,012.2)
Kids	8,625.4 (4,473.3)	3,201.0 (2,460.2)	9,285.2** (3,647.9)	1,754.8*** (594.73)	-2,253.6 (6,728.3)	-13,368 (2,357.1)
N	453	453	472	472	487	487

Note: See table 3.3 for source, notes, and significance levels.

Table 3.5 Tobit and OLS Results for Female Widow Only

Dependent Variable Model	Total		Term		Term		Whole	
	Tobit (1)	OLS (2)	Tobit (3)	OLS (4)	Tobit (5)	OLS (6)	Tobit (5)	OLS (6)
SSB	0.2402 (0.1434)	0.0543 (0.0963)	0.0778 (0.0961)	0.0139 (0.0342)	0.6842** (0.3100)	0.0339 (0.0748)	0.6842** (0.3100)	0.0339 (0.0748)
PVR	0.0036 (0.0021)	0.0037 (0.0027)	-0.0034 (0.0019)	-0.0001 (0.0002)	0.0081** (0.0039)	0.0035 (0.0024)	0.0081** (0.0039)	0.0035 (0.0024)
Working	3,848.9** (1,631.9)	2,589.0 (1,573.8)	1,642.4 (1,074.3)	706.83 (427.54)	5,025.3 (3,413.0)	1,869.2 (1,467.9)	5,025.3 (3,413.0)	1,869.2 (1,467.9)
Age	-381.36*** (69.716)	-96.511*** (26.810)	-213.34*** (45.934)	-46.397*** (15.778)	-480.66*** (158.66)	-42.135** (19.736)	-480.66*** (158.66)	-42.135** (19.736)
Nonwhite	2,289.5** (1,138.2)	964.84 (697.72)	1,675.1** (729.73)	474.24 (257.79)	-932.74 (2,668.6)	454.24 (608.16)	-932.74 (2,668.6)	454.24 (608.16)
Kids	2,033.6 (1,247.0)	287.91 (553.19)	471.99 (802.27)	-31.699 (352.60)	4,630.9 (2,888.5)	294.18 (387.59)	4,630.9 (2,888.5)	294.18 (387.59)
N	2,152	2,152	2,266	2,266	2,324	2,324	2,324	2,324

Note: See table 3.3 for source, notes, and significance levels.

Table 3.6 Tobit and OLS Results for Married Couples with Total Household Coverage

Dependent Variable Model	Total Tobit (1)	Total OLS (2)	Term Tobit (3)	Term OLS (4)	Whole Tobit (5)	Whole OLS (6)
	SSB	0.9087*** (0.3225)	0.5744 (0.3768)	0.2668 (0.1644)	0.1036 (0.1459)	1.4557*** (0.4526)
PVR	0.0289*** (0.0038)	0.0255*** (0.0082)	0.0001 (0.0017)	0.0014 (0.0013)	0.0327*** (0.0051)	0.0209*** (0.0071)
Working husband	6,502.8 (4,646.6)	6,889.3 (4,739.7)	-1,602.4 (2,468.6)	956.93 (1,583.4)	9,439.2 (6,478.5)	4,744.8 (4,016.6)
Working wife	9,676.4 (5,537.6)	6,512.9 (7,429.9)	5,635.4** (2,868.0)	2,984.2 (1,922.0)	2,217.8 (7,905.2)	3,270.9 (6,234.8)
Age of husband	-1,035.9*** (349.25)	-294.27 (249.98)	-750.98*** (180.27)	-284.04*** (76.812)	-559.88 (528.26)	5,791.0 (217.85)
Age of wife	-257.55 (312.22)	-309.40 (295.41)	66.711 (162.37)	-10.409 (79.747)	-501.19 (467.29)	-277.42 (268.06)
Nonwhite	-7,728.8 (5,774.9)	-3,864.2** (1,758.1)	1,244.8 (2,868.3)	-269.17 (715.50)	-50,519*** (10,364)	-4,134.2*** (1,378.5)
Kids	7,910.6 (5,976.5)	6,200.2*** (2,268.0)	1,091.7 (3,128.3)	819.06 (1,033.0)	7,681.5 (8,835.8)	5,065.3*** (1,848.8)
N	1,751	1,751	1,893	1,893	1,937	1,937

Note: See table 3.3 for source, notes, and significance levels.

specification is a positive and significant 0.91. This offset is similar to what Bernheim found, though at the lower end of his range. Once again, however, the split of total insurance into its two types yields dramatically different results. The coefficient for term insurance in column (2) is only 0.27, and is not significant at the 95 percent level. The coefficient on SSB in the whole insurance specification (column [3]), on the other hand, is a significant 1.46. Repeating the analysis with OLS regressions, we again find no significant correlation between SSB and life insurance ownership. As in the case with widows and widowers, the significant tobit results appear to be driven primarily by a whole life insurance, not the term insurance that the model is meant to represent.

Table 3.7 repeats the analysis on the subsample of married couples in which neither spouse is currently in the workforce. This distinction is quite important, as even the tobit coefficient on total life insurance is no longer significant. There are two important reasons to think that working couples may differ from nonworking. First, an employed individual still has some (albeit small) human capital to protect, just as a younger working-age individual does. Second, a worker is more likely to be covered by a group insurance plan through the employer. In either case, if workers also have higher Social Security benefits because of a stronger attachment to the labor force, this will induce a positive correlation between SSB and life insurance, even in the absence of a desire to undo annuitization.

Table 3.8 reports results similar to those in table 3.6, except that the dependent variable is life insurance coverage on the husband only. The pattern of coefficients on SSB is similar to those found in table 3. 6. Again, any positive correlation is limited to the total or whole life specifications, and is significant only in the tobit specifications. Table 3.9 shows the results for the sample of married women, and again we see the familiar pattern of coefficients.

In short, there is no evidence to suggest that term life insurance ownership among retired elderly households exhibits the correlation with SSB levels that the annuity offset model demands.

3.5.3 Test of Implication no. 3: Term Insurance as an Inferior Good

The third implication of the annuity offset model is that term life insurance will behave as an inferior good with respect to lifetime resources. This is simply because retirement consumption is viewed as a normal good, and therefore the demand for annuities should be an increasing function of resources. Since term insurance is to behave as a negative annuity, this means that the demand for term insurance should be declining with total lifetime resources.

I am unable to replicate exactly Bernheim's measure of lifetime resources because access to Social Security earnings records is unavailable. However, we can observe other components of resources, including the

Table 3.7 Tobit and OLS Results for Married Couples with Total Household Coverage (nonworking sample only)

Dependent Variable Model	Total		Term		Whole	
	Tobit (1)	OLS (2)	Tobit (3)	OLS (4)	Tobit (5)	OLS (6)
SSB	0.0514 (0.3146)	-0.1214 (0.3152)	0.0850 (0.1722)	-0.0426 (0.1103)	0.4301 (0.4412)	-0.0359 (0.2594)
PVR	0.0245*** (0.0038)	0.0216*** (0.0083)	0.0001 (0.0023)	0.0017 (0.0013)	0.0308*** (0.0052)	0.0190*** (0.0078)
Age of husband	-1,029.7*** (331.78)	-348.94 (277.59)	-704.88*** (184.08)	-277.62*** (76.702)	-500.96 (502.45)	-1,4535 (249.83)
Age of wife	-481.67 (300.29)	-533.43 (284.52)	3.0825 (166.94)	-36.418 (71.699)	-724.23 (451.59)	-484.06 (259.29)
Nonwhite	-10,281 (5,301.6)	-5,453.6*** (1,773.4)	-170.96 (2,799.5)	-835.96 (732.26)	-48,199*** (9,721.2)	-5,221.2*** (1,389.7)
Kids	3,042.9 (5,466.9)	3,961.5 (2,307.3)	-1,374.8 (3,035.8)	-317.32 (1,148.5)	2,415.4 (8,135.7)	3,593.9** (1,679.1)
N	1,393	1,393	1,505	1,505	1,536	1,536

Note: See table 3.3 for source, notes, and significance levels.

Table 3.8 Tobit and OLS Results for Married Men

Dependent Variable Model	Total Tobit	Total OLS	Term Tobit	Term OLS	Whole Tobit	Whole OLS
	(1)	(2)	(3)	(4)	(5)	(6)
SSB	0.7851*** (0.2640)	0.4080 (0.2674)	0.1934 (0.1448)	0.0370 (0.0960)	1.4209*** (0.4154)	0.3889 (0.2221)
	0.0208*** (0.0031)	0.0181*** (0.0061)	-0.0006 (0.0016)	0.0009 (0.0009)	0.0259*** (0.0046)	0.0154*** (0.0057)
Working husband	7,700.2** (3,827.4)	7,240.7 (4,047.6)	168.17 (2,171.6)	1,259.7 (1,235.7)	11,689** (5,883.0)	5,092.5 (3,582.7)
Working wife	-3,929.9 (4,593.4)	-1,769.2 (4,318.2)	2,128.6 (2,543.2)	975.72 (1,197.2)	-11,394 (7,368.1)	-2,225.4 (3,810.9)
Age of husband	-1,041.5*** (294.07)	-270.94 (210.71)	-744.40*** (161.99)	-252.96*** (63.281)	-791.92 (492.02)	0.8571 (189.65)
Age of wife	-184.67 (260.38)	-236.69 (246.53)	252.68 (145.05)	89.866 (50.504)	-547.23 (432.07)	-324.11 (229.46)
Nonwhite	-10,877** (4,798.9)	-3,678.9*** (1,236.4)	454.91 (2,504.2)	-568.51 (471.63)	-50,921*** (10,142)	-3,155.0*** (1,042.6)
Kids	6,575.6 (4,997.1)	5,123.3*** (1,676.8)	877.24 (2,762.8)	831.61 (848.74)	6,074.2 (8,231.1)	4,188.6*** (1,402.9)
N	1,841	1,841	1,946	1,946	1,979	1,979

Note: See table 3.3 for source, notes, and significance levels.

Table 3.9 Tobit and OLS Results for Married Women

Dependent Variable Model	Total Tobit	Total OLS	Term Tobit	Term OLS	Whole Tobit	Whole OLS
	(1)	(2)	(3)	(4)	(5)	(6)
SSB	0.3202 (0.1687)	0.2037 (0.1552)	0.0562 (0.1173)	0.0619 (0.0654)	0.6058** (0.2771)	0.1420 (0.1163)
PVR	0.0045*** (0.0016)	0.0035 (0.0022)	0.0000 (0.0013)	0.0004 (0.0005)	0.0060*** (0.0023)	0.0030 (0.0018)
Working husband	1,373.4 (2,503.8)	828.63 (1,497.5)	-2,215.4 (1,794.2)	-251.48 (683.79)	4,510.9 (3,959.9)	973.33 (1,229.1)
Working wife	9,016.2*** (2,896.4)	5,942.7 (3,335.7)	4,203.2** (1,988.4)	1,730.3 (1,038.0)	9,043.8 (4,719.6)	3,937.9 (2,924.3)
Age of husband	-365.95 (188.88)	-12.422 (84.895)	-237.89 (128.56)	-23.745 (31.576)	-296.10 (321.91)	18.801 (74.479)
Age of wife	-232.24 (170.72)	-143.55 (113.04)	-160.91 (115.72)	-98.718** (46.716)	-112.64 (291.01)	-33.212 (98.039)
Nonwhite	-1,548.3 (2,983.3)	-1,170.1 (688.73)	3,330.4 (1,898.2)	284.84 (378.71)	-25,548*** (6,410.8)	-1,426.5*** (516.74)
Kids	2,183.6 (3,199.7)	1,302.6 (865.54)	364.61 (2,207.3)	133.78 (406.76)	-1,333.0 (5,295.1)	1,012.2 (670.07)
N	1,896	1,896	1,982	1,982	2,008	2,008

Note: See table 3.3 for source, notes, and significance levels.

actuarial present value of pensions, Social Security, housing wealth, and financial net worth. I construct the variable PVR to be the sum of these resource variables. In addition, I am able to use indicator variables for education and occupation to proxy for lifetime earning effects.

Using a measure of lifetime resources that included lifetime earnings and the present value of pensions and Social Security, Bernheim found mixed results in his test of this implication. Specifically, in his basic tobit results, he found that the lifetime resource effect was negative for the average childless couple, but positive and insignificant for couples with children. He finds better support for the notion that at least some portion of total life insurance demand behaves as an inferior good by conducting refined estimates that model total life insurance holdings as the sum of two distinct, but separately unidentified, processes. Based on these refined estimates, he concludes that the term part of total life insurance ownership is the component that is behaving like an inferior good.

Looking at the coefficient on PVR among widows and widowers (tables 3.3, 3.4, and 3.5) and among married couples (tables 3.6–3.9), we find no significant relationship between PVR and term life insurance. While the sign of the coefficient is negative in some of the term life insurance specifications, it is not even approaching significance at any standard level of confidence. The coefficient on PVR in the whole life insurance specifications, and as a result in some of the total life insurance specifications, is strongly positive. This latter finding is consistent with Bernheim's conclusion that term and whole life insurance respond rather differently to variation in total resources. In the AHEAD data, however, there is no evidence that term insurance is behaving like an inferior good.

3.5.4 Test of Implication no. 4: Term Insurance Owners Consume Less than Social Security Income

The fourth and final empirical implication of the annuity offset model derives from the definition of being “over-annuitized” by Social Security. The basic notion behind this model is that household bequest motives are sufficiently strong that their desired consumption level is less than the annuity provided by Social Security, and that they would therefore prefer to keep some of their wealth unannuitized in order to leave it to their heirs.

Conceptually, this is a straightforward implication to test, since it requires the simple comparison of consumption to the income provided by Social Security. If households consume more than the Social Security benefit, then they are not over-annuitized. However, this implication is difficult to test directly in the AHEAD data due to the fact that a good measure of consumption is difficult to construct with currently available data. Therefore, I will rely on less direct methods to infer the extent to which households wish to consume less than their Social Security income.

It is useful first to consider a household's dynamic budget constraint:

$$(2) \quad W_{t+1} = (W_t - C_t + \text{SSB}_t + Y_t)(1 + r),$$

where W_t represents financial wealth at period t , C_t is consumption in period t , SSB_t is the income flow from Social Security, and Y_t is the income flow from other (non-Social Security) sources. If it is true that individuals are over-annuitized by Social Security, it must be the case that $\text{SSB}_t \geq C_t$. If not, then the constraint of the mandated Social Security annuity is nonbinding, and it need not be offset by life insurance. Since we do not directly observe consumption in the AHEAD survey, this test must necessarily be indirect. To be over-annuitized by Social Security requires $\text{SSB}_t - C_t \geq 0$. This implies

$$(3) \quad W_{t+1} - W_t(1 + r) \geq Y_t.$$

In other words, we need the amount of net saving done by households to exceed the levels of non-Social Security income that they receive during the period. That is, they must be saving some fraction of their Social Security payments in addition to all non-Social Security income. According to the annuity offset model, households that own term life insurance should be saving all non-Social Security income, and then supplementing this bequeathable saving with life insurance.

One simple way to test for this is to make use of a question asked in the first wave of the AHEAD survey:

Not counting any money or assets that you may have given children or others, did you [and your (husband/wife/partner)] use up any of your investments or savings during (1992/1993) to pay for expenses?

If households are spending down their existing nonannuitized assets in order to pay for current consumption expenses, then they must be consuming at least as much as their current *total* income, and therefore at least as much as their Social Security income. Therefore, these individuals would have no reason to hold life insurance.

Table 3.10 shows that approximately one-fourth of all households spent down assets in 1992–93. Importantly, the overwhelming majority of these households own life insurance, and in particular, term life insurance. In fact, the proportion of those owning term life insurance who spent down assets does not appear to be very different from the proportion of those not owning term insurance who spent down assets, for both widows and couples. Specifically, 24 percent of widows and 25 percent of couples who own term life insurance engaged in a spend-down of financial assets. This test clearly understates the proportion of term life insurance owners who are consuming more than their SSB levels, as it does not account for consumption out of non-Social Security income. If a person also has pension or investment income, for example, the individual may consume in excess of Social Security, and yet still be a net saver.

Table 3.10 Asset Spend-Down Versus Life Insurance Ownership (fraction of population)

	Own Any Life Insurance?		Own Term Life Insurance?	
	No	Yes	No	Yes
	<i>Widows/Widowers</i>			
Spend-down?				
No	.383	.383	.516	.250
Yes	.105	.128	.157	.078
<i>Married Couples</i>				
Spend-down?				
No	.173	.590	.389	.371
Yes	.047	.190	.115	.125

Source: Author's tabulations from the AHEAD survey.

3.5.5 Summary of Annuity Offset Tests

All four of the major implications of the annuity offset model fail empirical testing in the AHEAD data. As a result, it seems clear that this model does not explain life insurance behavior of elderly households. This leads to the obvious next question, "What is the alternative hypothesis?" This is the subject of the next section.

3.6 Alternative Explanations

There are a number of plausible alternative hypotheses that could explain why elderly individuals and couples hold life insurance. These alternatives share the common feature that none of them rely on the four empirical implications of the annuity offset model. That is, these hypotheses are still quite plausible even knowing the results of section 3. 5. It is not the goal of this paper to conduct a definitive test to select from among these alternative hypotheses. I will, however, present some suggestive evidence to provide direction for further research. The four alternative hypotheses I discuss below include

1. *Couple protection.* Elderly couples use life insurance to insure against loss of pension or Social Security benefits upon the death of the first spouse.

2. *Inertia.* Life insurance holdings are simply "residue" from attempts earlier in life to insure human capital.

3. *Estate tax planning.* Life insurance is used to assist with estate tax planning (e.g., to provide liquidity) in wealthier households.

4. *Funeral expenses.* Many elderly view life insurance as their burial money.

3.6.1 Couple Protection

The first of these alternatives, the couple protection model, assumes that married couples are purchasing life insurance in order to reallocate life-contingent incomes. For example, suppose a husband has a pension plan which is being paid out as a “joint and 50 percent contingent” annuity. This type of annuity treats the spouses asymmetrically. If the wife dies first, the husband continues to receive the full benefit. If the husband dies first, on the other hand, the pension income paid to the wife drops by 50 percent. If the couple decide that they would like to reallocate income from the husband-only state to the wife-only state, one way to do this is to purchase a term life insurance policy on the husband.

The evidence on this alternative is mixed. First, it cannot explain the fact that 62 percent of widowers and 49 percent of widows hold life insurance policies. Second, Auerbach and Kotlikoff (1987, 1989) tested this model of couples using several data sets, including the RHS, and found little support for the model’s implications. Specifically, they calculated the decline in resources that a married individual would face upon the death of his or her spouse, and used this variable as a predictor of life insurance ownership on the spouse. They found that most households do not adequately insure spouses against the potential resource loss associated with widowhood.

On the other hand, 95 percent of husbands in the AHEAD sample who own term life insurance name their spouses as the policy beneficiaries. If the insurance is truly being held to leave as a bequest to children, there is no obvious reason to leave the policy to the surviving spouse instead. Further exploration of this alternative hypothesis is left for future research.

3.6.2 Inertia

The second alternative hypothesis is that the elderly hold life insurance while old only because they held it when they were young. This could reflect irrational or rational behavior, as when an individual rationally keeps a policy because it represents a good value from today forward. This could be because the policy is already paid up, or because someone else is paying for the policy (e.g., a child or a former employer). For example, roughly 40 percent of the individuals in the AHEAD data who are covered by a term life insurance policy are currently paying no premiums. Alternatively, many have had a multiyear term policy with flat or level premiums, and therefore the policy is a better than actuarially fair deal from this time forward because the individual has essentially prepaid.

There are also nonrational reasons that one might hold on to a policy that was bought earlier in life. Samuelson and Zeckhauser (1988) provide evidence of status quo bias in decision making. They point out that almost every decision, such as an elderly individual’s decision about how much

life insurance to hold, has a status quo alternative of “doing nothing” or “maintaining one’s current or previous decision.” Using evidence from a series of experiments, as well as data on retirement plan choice, they show that individuals have a strong propensity to stick with the status quo. They attribute this to the presence of “(1) . . . transition costs and/or uncertainty; (2) cognitive misperceptions; and (3) psychological commitment stemming from misperceived sunk costs, regret avoidance, or a drive for consistency.”

There are several pieces of evidence that suggest inertia may affect a significant fraction of the sample. First, data from LIMRA’s Life Insurance Ownership Study indicates that most policies held by the elderly are quite old. When asked the age of the newest life insurance policy held, the median response among those aged seventy and up was thirty-two years. The median age of the oldest policy was forty-two years. Fully 30 percent of these elderly individuals bought their *newest* insurance policies before the age of thirty, and have not purchased any additional insurance since that time. Half of those owning insurance have not bought a policy since the age of forty-three. It thus appears that the majority of policy owners have not purchased any insurance for many decades, at least raising the possibility that their continued ownership is due to a failure to cancel.

On the other hand, 17 percent of those who own life insurance bought their most recent policies since the age of sixty-five. According to a LIMRA Buyer Study (1996), only 8 percent of all life insurance policies sold by agents to individuals aged sixty-five and up were term policies. Most of the rest were whole life policies, which are commonly used for estate planning purposes. The average size of the policies sold to those aged sixty-five and up was \$92,800, with an annual premium of \$4,698. These are quite large policies compared to the average policy size found in the AHEAD data, indicating that these individuals are likely to be wealthier than average and more concerned with estate planning. While these households may well be concerned about bequests, it is highly unlikely that they would be purchasing large cash-value policies in order to offset Social Security.

3.6.3 Estate Tax Planning

The third alternative hypothesis is that individuals hold life insurance to aid in estate planning. There are several reasons that a wealthy household that would be subject to estate taxation upon death would use life insurance as part of an estate planning strategy. First, the owners of a family business may wish to provide heirs with sufficient liquidity to pay for the estate taxes associated with the value of the business operation, in order to avoid the need to liquidate business assets. Holtz-Eakin, Phillips, and Rosen (1999) explore this point in detail. They find that, other things equal, business owners purchase more life insurance than other individuals.

While it is undoubtedly true that some high-wealth households use life insurance as an effective estate planning tool, this simply cannot explain more than a small fraction of households in the AHEAD data. Fewer than 5 percent of households in the data have a combined net worth and life insurance face value in excess of \$600,000, which was the point at which the estate tax became an issue for households at the survey date.

3.6.4 Funeral Expenses

The fourth alternative hypothesis is that elderly individuals view life insurance policies as their “burial money.” This could be a mental-accounting approach to portfolio choice (Thaler 1985) or a rational way to circumvent the probate process. Either way, this burial money notion may explain the preponderance of small face value policies in the sample, since according to the National Funeral Directors Association, the average cost of a funeral in 1997 was \$4,782. Hurd and Smith (this volume) show that total out-of-pocket death expenses (which include out-of-pocket medical and funeral expenses) for decedents in the AHEAD data average \$8,934. For comparison, the median amount of total life insurance coverage is \$5,000 among male widowers, \$2,500 among female widows, \$3,000 among married women, and \$9,000 among married men. It seems reasonable to suspect that many of these small policies are held for the purpose of paying for final death-related expenses. This notion is present in popular financial planning books as well. The author of one such book tells the story of a conversation with a widow who asked him to review her finances. She was financially well off, with more than \$600,000 in net worth and annual living expenses of only \$30,000. When he asked her why she was carrying a term life insurance policy that was costing her several hundred dollars a month in premiums, she replied, “That is to bury me” (Gardiner 1997).

A LIMRA study confirms that life insurance is frequently purchased with the intention of using the proceeds to pay for one’s burial. Eighty-three percent of widows report using the life insurance proceeds of their deceased spouses primarily to pay for death-related expenses. The LIMRA also reports that paying for death-related expenses is the most commonly cited reason that consumers give for purchasing life insurance.

One reason that life insurance is a popular device for paying for death expenses is that it avoids probate if paid to a named beneficiary (Graves 1994). Probate proceedings can tie up ordinary assets for many months, so that family members are unable to use these assets to pay for funeral or other death-related expenses. The proceeds from a small life insurance policy, because it avoids the probate proceedings, can provide the decedent’s family with timely access to funds with which to pay for these expenses.

Table 3.11 Determining Fraction of Sample Subject to Over-Annuitization Due to Bequests

	Widows/Widowers (%)	Married Men (%)
Full sample	100.0	100.0
Fraction holding any insurance	49.9	70.7
Fraction holding term insurance	31.5	41.6
Minus those purchasing private annuity	29.9	39.3
Minus those with private pension annuity	17.7	14.5
Minus those spending down financial assets	13.9	10.7
Minus those with zero term premium	9.2	8.6
Minus those with term premium < 0.5 actuarially fair	8.9	7.7
Minus those naming spouse as beneficiary	—	0.8
Minus policies under \$5,000 (funeral expenses)	2.0	0.5

Source: Author's tabulations from the AHEAD survey.

3.6.5 Putting It All Together

Once we account for all the behavior that is directly inconsistent with the annuity offset model or is potentially explained by alternative hypotheses, what fraction of households exhibit behavior that can be explained only by the desire to offset annuities? A simple running tabulation presented in table 3.11 shows that it is likely to be a trivial fraction of the population—far less than the 25 percent figure resulting from earlier analyses.

Table 3.11 starts with the full population of widows and widowers in the left-most column, and married men in the right. As the chart shows, approximately half of all widows and 71 percent of married men own life insurance policies. However, the annuity offset model is really a model of term life insurance, which means we are really concerned about the 31.5 percent of widows and the 41.6 percent of married men who own term policies. Next we can subtract those households which purchase a life annuity, since these households would clearly not purchase life insurance to offset Social Security only to turn around and annuitize additional resources. If we follow Bernheim's lead in treating life-contingent pension annuities as a choice variable, we can further reduce the sample to only 17.7 percent of widows and 14.5 percent of married men.

Next, we can eliminate those households who are spending down their financial assets for consumption, since these individuals are also clearly not constrained by the Social Security income floor. At this point, we have 13.9 percent of widows and 10.7 percent of married men still in the pool. Now let us account for individuals whose term life insurance coverage costs them nothing. The reasoning here is that if an individual can receive a policy at zero marginal cost, then it is perfectly rational for him or her

to keep it regardless of whether he or she has bequest motives. In any case, since they do not have to use the Social Security benefits to pay for the premium, they are not offsetting the annuity in any direct way. This leaves us with 9.2 percent of widows and 8.6 percent of married men. Using similar logic, we can eliminate those for whom the term premium is an actuarially advantageous deal. Specifically, I exclude those whose term premiums are less than half of the actuarially fair term premium of a seventy-year-old in 1993.

For married couples, I eliminate those who name their spouses as the beneficiaries, since such policies may be held more for spousal protection than for bequest purposes. Finally, let us take the funeral expense notion seriously, and assume that any policy with a face value of under \$5,000 is essentially the individual's burial money. This reduces the sample to 2.0 percent of widows and 0.5 percent of married men.

The calculations in table 3.11 are meant to be illustrative only, and one can certainly quibble with any one of the above exclusion restrictions. These figures demonstrate, however, that one can explain away the finding that individuals use life insurance to offset Social Security. In short, with a few simple assumptions one can show that only a small fraction of households may be over-annuitized by Social Security because they have strong bequest motives.

3.7 Summary and Future Directions

This paper has presented substantial evidence that the reason the elderly hold life insurance is *not* to offset mandated annuitization in the form of Social Security in order to leave a bequest. Four empirical implications of the annuity offset model were developed and tested, and all four were found to be inconsistent with the behavior of elderly households in the AHEAD data set.

This finding is relevant to the current debate over the future of the Social Security system because it bears upon the question of whether mandatory annuitization is desirable. Were it the case that a substantial fraction of elderly households were over-annuitized by the existing Social Security system due to the existence of strong bequest motives, this would be evidence in favor of allowing choice over the annuitization decision. The results of this paper suggest that households are not over-annuitized by Social Security for bequest reasons. Therefore, the simple fact that many elderly households own term life insurance is not a sufficient reason to argue against mandatory annuitization of retirement resources. This finding is consistent with the idea that annuities are of substantial value in the retirement portfolios of elderly individuals (Mitchell et al. 1999; Brown 1999; Friedman and Warshawsky 1988). As a result, mandatory annuitization may be desirable to overcome adverse selection in the annuity market.

However, this conclusion should be tempered by the acknowledgment that individuals can be over-annuitized for reasons other than bequest motives, as suggested in work by Hurd (1987, 1989).

This paper then suggests several alternative hypotheses for explaining the large fraction of elderly households who own life insurance. While these alternatives were not subjected to formal empirical testing in this paper, informal evidence suggests that some of these alternatives may be relevant. For example, the fact that the vast majority of policies have been held for several decades suggests that many holdings may be due to inertia from insurance decisions earlier in life. This would be consistent with the “status quo bias” in decision making that has been documented by Samuelson and Zeckhauser (1988), among others. It may also be the case that many small policies are held as a method of prepaying death expenses, such as funeral expenses.

It has also been found that the majority of policies held by married individuals name the spouses as the beneficiaries rather than the children. This is at least suggestive that the purpose of these policies may be to provide an adequate consumption stream for a widowed spouse. While this hypothesis found little support in earlier empirical work by Auerbach and Kotlikoff, those authors suggested that this might be due in part to the poor quality of their data. Further investigation of these and other alternative hypotheses are left to future research.

References

- Auerbach, A. J., and L. J. Kotlikoff. 1987. Life insurance of the elderly: Its adequacy and determinants. In *Work, health, and income among the elderly*, ed. Gary Burtless. Washington, D.C.: Brookings Institution.
- . 1991. The adequacy of life insurance purchases. *Journal of Financial Intermediation* 1 (3): 215–41.
- Bernheim, B. D. 1991. How strong are bequest motives? Evidence based on estimates of the demand for life insurance and annuities. *Journal of Political Economy* 99:899–927.
- Brown, J. R. 2001. Private pensions, mortality risk, and the decision to annuitize. *Journal of Public Economics*, forthcoming.
- Friedman, B. M., and M. J. Warshawsky. 1988. Annuity prices and saving behavior in the United States. In *Pensions and the U.S. economy*, ed. Z. Bodie, J. Shoven, and D. Wise, 53–77. Chicago: University of Chicago Press.
- Gardiner, R. M. 1997. *Dean Witter guide to personal investing*. New York: Penguin.
- Graves, E., ed. 1994. *McGill's life insurance*. Bryn Mawr, Penn.: The American College.
- Holtz-Eakin, D., J. Phillips, and H. Rosen. 1999. Estate taxes, life insurance, and small business. NBER Working Paper no. 7360. Cambridge, Mass.: National Bureau of Economic Research, September.

- Hurd, M. D. 1987. Savings of the elderly and desired bequests. *The American Economic Review* 77:298–312.
- . 1989. Mortality risk and bequests. *Econometrica* 57 (4): 779–813.
- Kotlikoff, L. J., and L. H. Summers. 1981. The role of intergenerational transfers in aggregate capital accumulation. *Journal of Political Economy* 89 (4): 706–32.
- Laitner, J., and T. Juster. 1996. New evidence on altruism: A study of TIAA-CREF retirees. *The American Economic Review* 86 (4): 893–908.
- Lewis, F. D. 1989. Dependents and the demand for life insurance. *The American Economic Review* 79:452–67.
- Life Insurance Market Research Association [LIMRA]. 1993. Profiles in coverage: The widening gap in U.S. life insurance ownership. Windsor, Conn.: LIMRA.
- . 1996. The 1995 Buyer Study: A market study of new insureds and the ordinary life insurance purchased. Windsor, Conn.: LIMRA.
- Mitchell, O., J. M. Poterba, M. J. Warshawsky, and J. R. Brown. 1999. New evidence on the money's worth of individual annuities. *The American Economic Review*, 89 (5): 1299–1318.
- Modigliani, F. 1998. The role of intergenerational transfers and life cycle savings in the accumulation of wealth. *Journal of Economic Perspectives* 2 (2): 15–40.
- National Funeral Directors Association. 1998. Funeral price information from the NFDA 1997 general price list survey. Available at [<http://www.nfda.org>].
- Samuelson, W., and R. Zeckhauser. 1988. Status quo bias in decision making. *Journal of Risk and Uncertainty* 1:7–59.
- Thaler, R. H. 1985. Mental accounting and consumer choice. *Marketing Science* 4:199–214.
- Yaari, M. 1965. Uncertain lifetime, life insurance, and the theory of the consumer. *Review of Economic Studies* 32:137–50.

Comment Anne Case

This is a sensible, straightforward paper that answers some old questions, and raises interesting new ones. Brown brings new data to an old question—whether we can find evidence of Social Security's causing the elderly to hold more annuities than they would otherwise choose, forcing them to offset their annuitization by purchasing life insurance. Using newly available data from the Asset and Health Dynamics of the Oldest Old (AHEAD) survey, Brown finds no evidence in favor of the hypothesis that households are over-annuitized by Social Security.

Brown is able to revisit evidence presented in Bernheim (1991), evidence that, according to Bernheim, “strongly suggests that an increase in Social Security old-age insurance (OAI) benefits tends to shift household resources from a regime in which they obtain annuities from private sources . . . and ultimately into a regime in which they purchase life insurance” (900). The statistic often cited from the Bernheim paper is that for “more

than one-fourth of all households, transfer motives are so strong that the compulsory provision of annuities through Social Security actually reduces bequests below the first-best levels" (900). As Brown lays out quite coherently, there are many reasons that the Bernheim results may not be robust. Bernheim, using the Retirement History Survey (RHS), did not have the ability to separate term life insurance from whole life insurance, and the latter often represents a tax-deferred savings instrument to wealthier people. Even without details on data deficiency, the Bernheim results seem rather implausible. The part of the income distribution most concerned with bequests would be the part *least* likely to be over-annuitized by Social Security. Social Security income represents only a small fraction of consumption for wealthy households.

While Brown's results seem reasonable in their totality, they are more problematic in some of their details. In the conference version of the paper, the results presented in the key tables had been estimated using tobits. It is highly likely that there will be heteroscedasticity in the unobservables in the demand for life insurance equation, which would bias the coefficient estimates. Wealthy households have greater scope for idiosyncratic behavior, and it would be surprising if we did not see a fanning out of the variance of unobservables with household income or wealth. In tables 3.3–3.9, Brown presents results on the demand for whole life insurance, term life insurance, and total life insurance for different groups of interest (widowers and widows, married couples, nonworking married couples). If the variance of the unobservable component of the demand for insurance differed between these groups, or between whole and term insurance, then we would expect to see differences in the Social Security benefit (SSB) coefficient due to differences in the variances, and not necessarily due to behavioral differences. (I think the next version of the paper will speak to this, so there seems little point in belaboring it here.)

There are some puzzles in the results presented in tables 3.3–3.9. For example, Brown suggests that because the coefficient on whole life insurance is positive and significant, while that on term insurance is smaller and insignificant, that whole insurance is driving the relationship between the level of SSB and life insurance that one finds when one looks at total insurance and SSB. However, something here does not quite add up. Fifty-seven percent of married couples own whole life insurance, and 50 percent own term insurance. Thus it would seem that one should find the tobit estimate of the effect of SSB on life insurance holdings for married couples to be *larger* than either of the two estimates taken individually, all else equal. That the effect is smaller than that for whole insurance estimated separately (0.91 in place of 1.46) suggests there are other forces at work here, making it difficult to state that whole insurance is driving the relationship between SSB and total life insurance.

There are many interesting findings in Brown's results that he does not

discuss in the paper. For example, nonwhite widows appear to hold a good deal of term life insurance. Why would this be the case? Are these women living with children and grandchildren for whom they are the main providers?

The variable “Kids” is an indicator that the respondent has any living children (although it is unclear why one would not use the number of living children in these tables rather than an indicator variable). In tables 3.6 and 3.7, the “Kids” indicator is generally an insignificant predictor of life insurance holdings. This seems surprising. I would have thought we would see whole life insurance holdings respond to children.

In the term life insurance equations for the full sample of married couples, an indicator that the husband is working is not a significant predictor of holding term life insurance, while an indicator that the wife is working is. What explains this? I would have thought term insurance was there to protect the income flow into the household upon the unexpected demise of a working person. If, as is likely to be the case, a husband is earning more than a wife, it would seem that that source of income would be more likely to merit protection.

The paper promises that there will be a lot of interesting work ahead. There is a need to examine life insurance in a coherent model. (Brown mentioned at the conference that the saying in the business is that “life insurance isn’t bought, it’s sold.” That seems an important part of the story still to be incorporated.) Even for a conference of economists researching aging, Brown spent two pages of his paper clarifying the difference between whole and term insurance. It seems both confusion and market imperfections may be worth modeling.

Finally, it seems the time has come to revisit the literature on spousal (and family) protection. Some of the Auerbach and Kotlikoff work cited analyzes data from the late 1960s. Much has changed—both in terms of those savings mechanisms that are available and in use, and in terms of the data available with which to analyze decisions made. It would seem an ideal time to pursue this work.

Reference

- Bernheim, B. D. 1991. How strong are bequest motives? Evidence based on estimates of the demand for life insurance and annuities. *Journal of Political Economy* 99:899–927.