NBER WORKING PAPER SERIES

HOW STRONG ARE BEQUEST MOTIVES? EVIDENCE BASED ON ESTIMATES OF THE DEMAND FOR LIFE INSURANCE AND ANNUITIES

B. Douglas Bernheim

Working Paper No. 2942

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 1989

I would like to thank the Center for Economic Policy Research at Stanford University, and the National Science Foundation (Grant No. SES-8607630) for financial support. Laurence Levin provided invaluable research assistance. Previous drafts of this paper bore the titles "Life Insurance, Annuities, and Bequests," and "Annuities, Bequests, and Private Wealth." This paper is part of NBER's research program in Taxation. Any opinions expressed are those of the author not those of the National Bureau of Economic Research.

NBER Working Paper #2942
April 1989

HOW STRONG ARE BEQUEST MOTIVES? EVIDENCE BASED ON ESTIMATES OF THE DEMAND FOR LIFE INSURANCE AND ANNUITIES

ABSTRACT

This paper presents new empirical evidence in support of the view that a significant fraction of total saving is motivated solely by the desire to leave bequests. Specifically, I find that Social Security annuity benefits significantly raise life insurance holdings and depress private annuity holdings among elderly individuals. These patterns indicate that the typical household would choose to maintain a positive fraction of its resources in bequeathable forms, even if insurance markets were perfect. Evidence on the relationship between insurance purchases and total resources reinforces this conclusion.

B. Douglas Bernheim
Kellogg Graduate School of Management
Northwestern University
2001 Sheridan Road
Evanston, IL 60208

1. Introduction

The past decade has witnessed the emergence of an enormous literature that examines the various economic roles of intergenerational transfers. Unfortunately, empirical analysis of gifts and bequests has proven problematic, and several critical questions remain unanswered. Economists disagree sharply about the quantitative importance of intergenerational transfers in the U.S. By some estimates, these transfers are responsible for approximately four-fifths of aggregate capital accumulation (see Kotlikoff and Summers [1981] and Kotlikoff [1988]); others place this figure below one-fourth (see Modigliani [1988]). There are also several competing schools of thought concerning transfer motives. Some argue that bequests are accidental, and result from precautionary responses to the absence of perfect insurance markets (Hurd [1987]). Others argue that bequests are intentional, and either reflect altruism (Tomes [1981]), or self-interested exchange with one's heirs (Bernheim, Shleifer and Summers [1985]).

This paper presents new empirical evidence in support of the view that a significant fraction of total saving is motivated solely by the desire to leave bequests. This evidence is based upon the premise that bequest motives fundamentally alter attitudes towards insurance.

Consider first a life cycle saver who has no desire to enrich his heirs. This individual will always take advantage of opportunities to purchase annuities,[1] as long as marginal annuities pay a rate of return in excess of that received on conventional assets.[2] In contrast, an individual who is motivated in part by the desire to leave a bequest must weigh the benefit of higher survival-contingent returns against the costs of a reduced estate. In general, such an individual will wish to convert less that 100% of his assets into annuities, even if annuities are available at actuarially fair rates. If the government converts some fraction of his resources into annuities, he will attempt to readjust his portfolio, purchasing annuities if the fraction is too small, and selling them if it is too large (an individual sells annuities by purchasing term life insurance -- see Yaari [1965]).

The evidence presented in this paper strongly suggests that an increase in Social Security Old Age Insurance (OAI) benefits tends to shift households from a regime in which they obtain annuities from private sources (e.g. employers), to a regime in which they acquire neither additional annuities nor life insurance, and ultimately into a regime in which they purchase life insurance, shunning private annuities. Furthermore, an increase in OAI benefits tends significantly to depress the level of private annuities for households within the first regime, and to raise the level of life insurance purchases within the third regime.

These patterns indicate that the typical household would choose to convert less that 100% of its resources into annuities, even if insurance markets were perfect. This conclusion is at odds with the accidental bequest hypothesis (as well as several other versions of the life cycle hypothesis that appear to explain consumers' reluctance to fully annuitize without appealing to bequest motives), but is in harmony with the view that most individuals are motivated in part by the desire to leave bequests. Indeed, the evidence indicates 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 first-best levels. Evidence on the relationship between insurance purchases and total resources reinforces these findings.

The paper is organized as follows. In section 2, I elaborate on the relevant implications of the life cycle and bequest motive hypotheses. Section 3 contains a discussion of data sources, and describes the construction of several key variables. Section 4 discusses various institutional factors that affect my interpretation of the data. In sections 5 and 6, I estimate empirical models describing private purchases of life insurance and annuities, and I interpret the results.

2. Theoretical Preliminaries

In this section, I elaborate upon the theoretical relationship between Social Security OAI benefits and private insurance holdings under several alternative models of personal savings behavior. I argue that the presence of a bequest motive leads to distinctive behavior; these distinctions form the basis for subsequent empirical tests.

A. The Bequest Motive Hypothesis

First, consider the following parsimonious model of an individual who is, in part, motivated by the desire to leave bequests. He possesses total initial resources W_0, which he divides between purchases of annuities, A, and purchases of conventional assets, B. For simplicity, I assume that the individual consumes nothing in the current period (one may also think of this individual as having already chosen current consumption, and selecting an optimal division of assets conditional on this choice).

One dollar invested in annuities yields a return of $\$\alpha$ in the subsequent period if the individual lives, and nothing if he dies in the interim. One dollar invested in conventional assets yields $\$\beta$ one year hence, whether or not he survives. Thus, having chosen a portfolio consisting of (A,B), the individual either survives, in which case his resources in period 1 are given by

(1) $$W_1 = \alpha A + \beta B,$$

or he dies, in which case his heirs receive βB.

The individual chooses A and B subject to a budget constraint,

(2) $$W_0 = A + B,$$

and non-negativity constraints

(3) $$B,\ W_1 \geq 0$$

(note that we do not constrain $A \geq 0$ -- he may sell annuities). His utility is given as a function of conditional returns and conditional bequests:

(4) $$U = U(B,\ W_1).$$

One can think of (4) as a "reduced form" expression for the utility of an altruist, given optimal disposition of period 1 resources, as well as the contingent plans of potential heirs. Equation (4) could also embody a bequest motive that arises from the desire to facilitate intrafamily exchange (e.g., the strategic bequest motive formulated by Bernheim, Shleifer, and Summers [1985]). Other interpretations are also possible; I assume only that both B and W_1 are "goods", and that the indifference curves generated by (3) have the "usual" shape (i.e. convex to the origin).

The individual described above faces a standard consumer choice problem, which is illustrated in Figure 1. Despite the premium available on annuities, he generally chooses to hold positive amounts of bequeathable assets (i.e. the non-negativity constraints, (3), do not bind).

Now suppose that the government provides this individual with a compulsory annuity. That is, it confiscates A^g in period 0, returning αA^g in period 1, conditional upon survival. Clearly, this action does not change the individual's opportunity locus. If he previously chose to invest B^* in conventional assets and $A^* = W_0 - B^*$ in annuities, he will now maintain his previous conventional investments, and reduce his private annuity purchases to $A = A^* - A^g$. Figure 2, line L illustrates annuity purchases as a function of compulsory annuities. The point is simple: when government provides annuities at the market rate, it simply displaces private purchases of annuities dollar for dollar.

When $A < 0$, the individual sells annuity claims. Although a contract to sell an individual's rights to his stream of Social Security benefits is legally unenforceable, he can also accomplish this objective by purchasing life insurance (see Yaari [1965] for a general discussion). One can see this as follows. Consider a term life insurance policy that pays an amount $\$\gamma$ per dollar invested if the individual dies within the coverage period, and nothing if he survives. With perfect insurance markets, pricing is actuarially fair, so $\gamma = (1 - \beta/\alpha)^{-1}$. Now compare the sale of annuities to the purchase of life insurance. If the individual sells one dollar of annuities and invests the proceeds into conventional assets, his survival-contingent income falls by $\$(\alpha - \beta)$, and his death-contingent bequest rises by \$1. Suppose instead that he sells $\$(\alpha - \beta)/\beta$ of conventional assets, and purchases a life insurance policy. Then his survival-contingent income again falls by $\$(\alpha - \beta)$, and his death-contingent bequest rises by

$$\frac{(\alpha - \beta)}{\beta(1 - \beta/\alpha)} - \frac{(\alpha - \beta)}{\beta} = 1.$$

Thus, these two investment strategies have identical affects on state-contingent consumption.

The extension of this analysis to multiperiod settings is straightforward. In the Appendix to this paper, I show that a consumer can perfectly offset the forced acquisition of an annuity that pays \$1 per year by purchasing appropriate term life insurance contracts. Specifically, at each age t, he acquires $\$q_t/(1-\pi_t)$ of term life insurance, where q_t is the price of a \$1 annuity stream at age t, and π_t is the probability of surviving another year, conditional upon having reached age t. If survival probabilities decline with age, then so will term insurance holdings. Purchases of insurance are financed by sales of conventional assets. Specifically, the offsetting portfolio adjustment requires the consumer to reduce his conventional wealth by $\$q_t/(1-\pi_t)$ at each age t.

Of course, this simple analysis has a strongly counterfactual implication: it predicts that all individuals should (except by accident), hold either annuities (other than OAI benefits), or term life insurance. Casual empiricism suggests that many consumers hold neither; data presented in subsequent sections confirm this impression. This phenomenon is, however, not terribly mysterious, and undoubtedly arises from imperfections in insurance markets.

Suppose in particular that individuals have some private information concerning longevity. Then equilibrium in insurance markets may be characterized by adverse selection. The nature of equilibrium will depend in part upon the structure of contracts offered by private insurance firms. I will assume that firms compete in price alone (as in Akerlof [1970]), so that contracts do not involve quantity constraints (as in Rothschild and Stiglitz [1976]). Similar arguments apply to cases in which underwriters supply insurance at prices that vary with quantity.

Adverse selection induces a "wedge" between the buying price ($1/\underline{\alpha}$), the actuarially fair price ($1/\alpha$), and the selling price ($1/\bar{\alpha}$), of an actuarial claim on one dollar one period in the future. In particular,

(5) $$\underline{\alpha} < \alpha < \bar{\alpha}.$$

Intuitively, individuals who have low survival probabilities tend to buy life insurance, driving the implicit zero-profit selling price of annuities down, and individuals who have high survival probabilities tend to buy annuities, driving the zero-profit buying price up. Since we are not primarily concerned with the effect of Social Security provision on insurance market equilibrium, we will simply take $\underline{\alpha}$ and $\bar{\alpha}$ as given, and abstract from further details (the interested reader is referred to Abel [1985] and Eckstein, Eichenbaum and Peled [1985]).

Suppose then that the government, through coercion, can supply annuities at actuarially fair rates, but that consumers receive a return of only $\underline{\alpha}$ on private purchases, and must pay a return of $\bar{\alpha}$ on private sales. Then for a fixed level of total initial resources W_0 ($= B + A + A^g$), the consumer's budget constraint is given by

(6) $$W_1 = \min\,(\underline{W}_1(B),\ \bar{W}_1(B))$$

where

(7) $$\underline{W}_1(B) = \underline{\alpha}(W_0 + (\alpha/\underline{\alpha} - 1)A^g) - (\underline{\alpha} - \beta)B$$

and

$$(8) \qquad \overline{W}_1(B) = \bar{\alpha}(W_0 + (\alpha/\bar{\alpha} - 1)A^g) - (\bar{\alpha} - \beta)B$$

That is, market imperfections induce a kink in the individual's budget constraint, located at the point where $A = 0$ ($B = W_0 - A^g$). The government's choice of A^g places the kink at the corresponding point on line J in Figure 3. Note that there may be a wide range of provision levels for which the consumer purchases neither annuities, nor term life insurance ($W_0 - \overline{B}$ is the minimum such level; $W_0 - \underline{B}$ is the maximum).

Suppose the government provides annuities at the level $A^g < W_0 - \overline{B}$. Then the individual purchases positive amounts of annuities -- the relevant portion of his budget constraint is $\underline{W}_1(B)$. Equation (7) is written in a way designed to illustrate that a rise in A^g shifts the relevant portion of the budget constraint outward -- since $\underline{\alpha} < \alpha$, the consumer's "effective" wealth ($W_0 + (\alpha/\underline{\alpha} - 1)A^g$) increases with A^g (graphically, consider the result of shifting from M to M′ in figure 4.) Thus, the individual will substitute private for public annuities as before, and in addition will experience a positive "income effect", the magnitude of which will depend on the wedge between α and $\underline{\alpha}$. If both W_0 and B are normal goods, then A will fall, but by less than the rise in A^g.

Next, suppose that the government provides annuities at some level $A^g > W_0 - \underline{B}$. Then the individual purchases positive amounts of term life insurance -- the relevant portion of his budget constraint is $\overline{W}_1(B)$. In this case, a rise in A^g induces a negative income effect, the magnitude of which depends upon the wedge between α and $\bar{\alpha}$ (graphically, consider the shift from N to N′ in figure 4). Thus, while the individual substitutes private for public annuities as before, the income effect induces him to settle for a lower overall level of annuities, thereby reinforcing his response.

The effect of government annuity provision at actuarilly fair rates is depicted in figure 2, where we have superimposed lines L' and L". When $A^g = 0$, the consumer holds only annuities, and no term life insurance. An increase in A^g causes these purchases to decline, but by less than dollar for dollar (i.e. in Regime I, he moves along L'). When $A^g = W_0 - \bar{B}$, the consumer withdraws from private insurance markets entirely (Regime II). He remains in this regime until $A^g = W_0 - \underline{B}$. Further increases in A^g cause him to sell annuities on private markets (buy term life insurance), although here the offset is more than dollar for dollar (i.e. in Regime III he moves along line L").

For the most part, the empirical portion of this paper focuses on the relationship between private insurance purchases and public insurance provision. However, the relationship between private insurance purchases and total private resources also merits consideration. Turn once again to the simple model described above. Suppose we increase W_0, holding A^g fixed. How will this affect private insurance holdings?

Suppose as before that both B and W_1 are normal goods. Then, at a fixed price, the total demand for annuities rises with W_0. With A^g fixed, this will tend to (1) increase private annuity purchases by those previously holding private annuities, (2) decrease term life insurance purchases by those previously holding term life insurance, and (3) move individuals from Regime III to Regime II, and from Regime II to Regime I. Note in particular that term life insurance gives the appearance of being an <u>inferior</u> good -- this strong implication is clearly testable.

In the next subsection, I contrast the implications derived above with those obtained from standard formulations of the life cycle hypothesis, and variants thereof.

B. The Life Cycle Hypothesis

While the addition of uncertainty concerning length of life to the standard life cycle model helps to account for the apparent persistence of high levels of wealth subsequent to retirement (see Davies [1981]), it also generates very strong implications concerning attitudes towards insurance. In particular, individuals will always take advantage of opportunities to convert resources into annuities which, on the margin, yield rates of return in excess of those available on conventional assets (see Yaari [1965]). Thus, as long as the government does not change an individual's market opportunities through the provision of Social Security, it will not affect his private insurance holdings. While the economy-wide level of OAI benefits may indeed influence private opportunities (by altering equilibrium in insurance markets), it is less likely that cross-sectional variation in benefits would have the same effect. However, regardless of whether or not the provision of OAI benefits alters market opportunities, it should never induce a life cycle saver to purchase life insurance, for the simple reason that he does not care what happens after his death. Thus, under the standard life cycle hypothesis, we might or might not expect to find a relationship between public annuity provision and private annuity purchases, but we certainly would not expect to find a systematic relationship between public annuity provision and life insurance. Indeed, the simplest versions of the life cycle hypothesis are hard pressed to explain purchases of life insurance to begin with.

Proponents of the life cycle hypothesis might argue that it is possible to reconcile more sophisticated versions of the basic model with positive holdings of life insurance, as well as with the behavioral patters described in subsection A. In particular, a variety of additional factors might cause a life cycle saver to prefer to hold less than 100% of his resources as annuities. Presumably, an individual would then act to offset government policies which displace him from his interior optimum. While this is the essence of the behavior described in subsection A, the actual mechanics of behavior will differ according to the theory under consideration. Thus, it is indeed possible to rule out many, if not all of these competing explanations. I will discuss three specific additional factors, including the desire to insure spouses, the provision of annuities through families, and precautionary motives to provide for emergencies.

The first alternative concerns the behavior of couples. A life cycle couple cares about three distinct sets of contingent claims: the joint, husband-only, and wife-only survival contingent consumption streams (see Auerbach and Kotlikoff [1985]). As with a lone individual, the couple wishes to fully annuitize as long as annuities yield any premium over conventional assets. However, it is very concerned about the distribution of resources between the three sets of survival contingent streams. In particular, if the joint survival contingent stream is too large relative to either of the other streams, the couple may wish to redress this imbalance through the purchase of life insurance. Thus, for example, if the couple's only asset is the husband's human capital, it will purchase a life insurance policy on the husband. Similarly, the couple may respond to an exogenous increase in a pension which only pays benefits conditional upon the husband's survival by reducing other annuity purchases, or by taking out life insurance on the husband. Since our focus is on intergenerational transfers, it would be inappropriate to categorize the desire to provide for one's spouse as a "bequest" motive.[3]

These predictions appear quite similar to those discussed in subsection A. This is as it should be, since the couple effectively has a bequest motive with regard to its figurative "child" (the surviving spouse). However, there are important differences. Suppose we focus on retired individuals, so that human capital is not an issue. As long as the government provides annuities in a way that causes no imbalance between the three survival-contingent streams, the life cycle couple will have no cause to purchase life insurance.

There is reason to believe that OAI benefits are well balanced, in this sense. The joint-survival contingent stream consists of either the sum of personal benefits, or 150% of the largest personal benefit (whichever is greater). Each lone-survival contingent stream consists of the largest personal benefit. Thus, survivors receive between 1/2 and 2/3 of the joint benefit. For roughly 75% of retirement age couples in the mid 1970's (the period of our data sample), this ratio was in fact 2/3 (many wives had very little covered employment, and those that did were typically paid substantially less than their husbands).

One should also remember that most retired couples hold conventional assets in addition to annuities, and that the stream of income derived from these assets is not contingent upon the survival of either spouse. For the sample of retired couples studied in subsequent sections, conventional wealth roughly equals the actuarial value of annuities. When one adds up Social

Security and other assets, it therefore appears that the joint-survival contingent income stream for the average retired couple is only about 20% higher than either lone-survival contingent income stream. Unless household economies of scale are extraordinarily large and insurance market imperfections are very small, it is difficult to believe that a couple in this position would purchase life insurance in order to raise lone-survival contingent incomes still further relative to joint-survival contingent incomes. If anything, redistributions between survival-contingent streams should flow in the opposite direction (towards the joint-survival contingent stream).

At least three other considerations support the position that life insurance holdings among retired couples do not reflect attempts to achieve appropriate survival-contingent income profiles for spouses. First, until recently, it was very common for individuals <u>not</u> to elect survivorship options for private pensions. Indeed, this state of affairs precipitated legislation requiring pension plans to specify survivorship as the default option, to be waived only with the consent of the employee's spouse. Failure to elect survivorship suggests that typical couples are not overly concerned about the relative size of lone-survival contingent income streams.

Second, recall that Social Security treats spouses symmetrically (each spouse would receive the same benefit as a survivor). One would therefore expect a retired life cycle couple to hold roughly the same amount of life insurance on both spouses. Yet casual inspection of the data reveals that elderly couples primarily insure husbands.

Third, Auerbach and Kotlikoff [1985] have explicitly estimated and tested a model of the demand for life insurance by the life cycle couple, and found little or no empirical support for the implications of this model. My results should be interpreted in light of their paper: life insurance holdings among the elderly are much better explained as vehicles for transferring resources from couples, than as internal arrangements between spouses.[4]

A second competing explanation for the reluctance of life cycle consumers to convert all of their resources into annuities has been proposed by Kotlikoff and Spivak [1981]. They argue that, when confronted with annuity market imperfections, life cycle savers may form small private groups to pool risks stemming from uncertain longevity, and that the natural locus of such activity is the family. Parents may, for example, effectively purchase annuities from their children by demanding financial support in exchange for the privilege of being named heir. Even if annuities command a premium over conventional assets, individuals will purchase annuities from their relatives until the shadow price of family annuities rises to the market price. Family annuities show up in data as holdings of bequeathable wealth (coupled with intrafamily transfers).

While this theory does account for consumers' reluctance to convert all assets into annuities, as well as for the displacement of private annuities by OAI benefits (as in Regime (I) above), it is not consistent with the notion that individuals should offset excess annuities through the purchase of life insurance. An individual who is compelled to hold a certain level of annuities would never sell them off at a price below that which is acctuarially fair ($1/\bar{\alpha} < 1/\alpha$), only to buy them back from the family at a higher price (since the family pools risks imperfectly, the price of a family annuity must exceed $1/\alpha$). Thus, we can empirically distinguish the family annuity model from both the pure life cycle and bequest motive formulations, through estimation of the life insurance relationship described above.

A third variant of the life cycle hypothesis suggests that individuals hold conventional assets as a precaution against emergencies, such as illnesses and accidents. The reluctance to annuitize is explained as a consequence of adverse selection in annuity markets. Specifically, when confronted with an emergency which requires a large expenditure of resources, the heavily annuitized consumer is forced to liquidate a portion of his annuity claims at unfavorable terms. Furthermore, if the emergencies are by and large medical, the desire to sell an annuity may convey additional information concerning survival probabilities, which would exacerbate the adverse selection problem.

While concern for potential emergencies does account for consumers' reluctance to annuitize marginal assets, as well as for the displacement of private annuities by OAI benefits, it also implies that the individual should offset excess annuities by purchasing medical or disability insurance (see Levin [1988]), rather than life insurance. The term component of life insurance holdings avails the consumer nothing during an emergency, as long as he survives. Term life insurance benefits cannot, for example, cover hospital or nursing home expenses. Although the holder of a whole life policy may borrow against cash value, or even surrender the policy in the event of an emergency, were this his only objective, he would be better off accumulating resources in some other form that does not include a costly term insurance component. Moreover, if life insurance is simply a precautionary reserve, then one would expect the size of this reserve to rise with the overall level of resources. Thus, the "potential emergency" theory, like the "family annuitization" theory, can be distinguished from the bequest motive hypothesis on the basis of evidence concerning the relationships between life insurance, and both publicly provided annuities and private resources.

3. Data

The data sample is drawn from the Longitudinal Retirement History Survey (RHS), which followed more than 11,000 retirement-aged individuals (58 to 63 in 1969) for a period of ten years, starting in 1969. I also obtained some information from matching administrative records.

I drop observations for the following reasons. First, I exclude all individuals listed as widows or widowers in 1969. Since the survey does not collect information on the earnings history of a deceased spouse, data on such individuals is necessarily incomplete. Second, I drop all observations reporting non-positive bequeathable wealth (which includes the net value of owner occupied housing, equity in a business or farm, the net value of other property, and financial assets, minus debt). Such households may be "fully annuitized", in which case we would expect them to substitute private for public annuities, even under the pure life cycle hypothesis. Alternatively, non-positive assets may simply reflect poor financial planning. Either way, one would expect such households to behave atypically. Third, I exclude very poor households (those with lifetime resources, defined below, of less than $50,000). Again, I suspect that the behavior of such households is unrepresentative of the larger population. Fourth, I exclude households that have accumulated more than $500,000 in assets by 1969.[5] It is quite unlikely that Social Security plays a significant role in the retirement plans of such households. Finally, I drop all individuals with government pensions (including Civil Service, Military, and Railroad Retirement). Since individuals are not covered by Social Security during their period of employment with the government, there is a built in negative correlation between government pensions and Social Security benefits. Thus, one cannot treat government pensions as privately acquired annuities, even though the household's choice of employer is voluntary. Nor does it seem proper to construe government pensions as compulsory annuities, since there is a built in negative correlation between government employment and holdings of private pensions. By excluding such households, I avoid these considerations entirely. The final sample consists of 3,808 households.

The empirical analysis focuses on the 1975 wave of the LRHS (although I employ data from other waves to construct certain variables). This choice reflects a desire to minimize the role of future earnings in determining life insurance purchases. In 1975, the youngest respondent is 64 years old, while the oldest is 69. Consequently, the vast majority of respondents were retired, and few anticipate significant future labor force participation.

I measure life insurance holdings as the total face value of policies covering husband and wife in 1975. Unfortunately, survey responses do not allow one to separate out the cash values of these policies. Undoubtedly, many elderly individuals employ life insurance at least in part as a vehicle for saving. It is important to bear this problem in mind when evaluating the evidence presented below -- I will return to it at several points.

2,767 households report positive life insurance holdings in 1975. The average face value of household policies for this group is $9,716. It is interesting to note that this figure does not differ greatly from the 1969 average, as reported by Auerbach and Kotlikoff [1985], although relatively fewer households held life insurance in 1975 (73% vs. 87%).

I measure private annuity holdings as the total yearly benefit from private pensions received as an annuity in 1975 or thereafter (the survey questions allow one to distinguish yearly from lump sum payments). I omit direct private purchases of actuarial contracts from underwriters or their agents, as these are negligeable. Casual inspection of the data reveals suspiciously large variance in reported pension benefits for specific individuals, as well as non-reporting in particular years. For example, it is not uncommon to find a reported benefit stream which fluctuates unsystematically between 1971 and 1979, with zeros in one or two years. To minimize the effects of this apparent noise, I assume that private pensions pay a constant nominal stream, and calculate pension income by averaging positive values of reported benefits for each spouse. 1,731 households report positive private pension receipts. The average yearly benefit for this group is $2,652.

Having discussed sample construction and dependent variables, I turn finally to the independent variables. I calculate the household's Social Security benefit (SSB) using earnings histories from administrative records, and applying prevailing law in 1975. Independent variation in SSB presumably reflects considerations such as tenure in covered employment, the frequency with which an individual achieved the taxable maximum, and the division of lifetime earnings between spouses.

Lifetime Resources (LR) consist of the actuarial value of private pension and computed Social Security benefits, plus net lifetime earnings, accumulated to 1975. Given our assumption that individuals participate in insurance markets, it is appropriate to accumulate earnings at the implicit marginal rate of return on purchased annuities -- the actuarially fair rate (assuming 3% real interest) is used as an approximation. I employ earnings histories to calculate yearly earnings, as well as payroll taxes, which I net out. Unfortunately, the administrative records provide incomplete information, since yearly earnings are reported only up to the taxable maximum. Since these records also indicate the quarter in which the taxable maximum was reached, I was able to extrapolate yearly earnings using the method described by Fox [1976].

The empirical analysis in section 5 also employs a number of other independent variables. These are defined as follows:

KID: Equals one if the respondent has children; zero otherwise,

WID: Equals one if the respondent is a widow or widower; zero otherwise,

SINGLE: Equals one if the respondent is single or divorced; zero otherwise,

AGE: Age of respondent, minus the mean age for the entire sample.

4. Institutional Considerations

The theory presented in section 2A suggests that individuals will under certain conditions purchase either life insurance or private annuities, but never both. The data appear to be inconsistent with this implication: roughly 36% of the sample reports both in-force life insurance holdings and receipt of private pensions. One must, however, be very careful about interpreting this finding. There are many good reasons to believe that the annuity and life insurance variables described in section 3 measure discretionary insurance holdings imperfectly, in which case substantial overlap is to be expected. In this section, I discuss likely sources of data contamination, and I argue that these factors tend to bias my results against the bequest motive hypothesis. Thus, evidence of the behavioral patterns described in section 2A should be seen as all the more compelling.

A. Life Insurance

The model outlined in section 2A describes purchases of term life insurance, which pays a fixed benefit if the insured party dies, and nothing if he survives. According to the American Council of Life Insurance (ACLI) [1979], in 1974 term policies accounted for slightly more than one-half of all outstanding life insurance. The remainder primarily consisted of whole life and endowment policies, which combine term insurance with tax-favored saving.

It is quite likely that an even smaller fraction of policies held by elderly individuals are of the pure term variety. Specifically, the ACLI [1979] reports that in 1976, 17% of males over 65 held group life insurance (in the population as a whole, more than 99% of all group insurance policies were term). In contrast, 56% of males 65 and over held agent-marketed, individual life insurance policies (for the entire population, only about 30% of these policies were term). Thus, the majority of policies held by the individuals in the data sample used here were probably whole life or endowment.

The owner of a whole life or endowment policy can surrender his policy at any time for a cash value. The difference between the face value and cash value of the policy reflects the term insurance component. Since cash value rises as the policyholder ages, the term component shrinks. However, for whole life policies, cash value generally remains below face value until the individual reaches 100 years. Moreover, the ACLI reports that in 1974, more than 90% of all non-term insurance was whole life. Thus, the vast majority of individuals reporting life insurance holdings in the RHS sample probably had some level of term insurance.

Unfortunately, the RHS only provides data on the face value of life insurance policies -- cash value (and hence the term component) is unavailable. As a result, our life insurance variable reflects a mixture of term insurance and savings. How does this affect the empirical analysis in subsequent sections?

If the face value of life insurance includes savings, then, as long as consumption during retirement is normal, one would expect this component to rise with lifetime resources. Moreover, a large body of empirical work suggests that social security displaces private saving, at least to some extent (see Feldstein [1974], Feldstein and Pellechio [1979], Kotlikoff [1979], Munnell [1974], King and Dicks-Mireaux [1983], Diamond and Hausman [1984], and Bernheim [1987]). Although the impact on any given component of saving (such as life insurance cash value) is probably small, one would expect that, if anything, social security would depress this component. Indeed, this prediction follows directly from the theoretical analysis presented in the appendix to this paper. Thus, if the life insurance variable used here primarily measures saving, one should strongly reject the patterns predicted in section 2A. Contamination of the dependent variable biases the results against the bequest motive hypothesis.

When interpreting the data on life insurance, it is also important to bear in mind that some significant fraction of life insurance holdings may not reflect "rational" behavior. Many individuals may fail to cash out their policies either because of inconvenience or superstition. In those cases, policies held after retirement might simply be the residue of efforts to insure human capital earlier in life. Cynics might even adopt the stronger view that life insurance purchases are never guided by rationality -- indeed, an old adage asserts that life insurance is only sold, and not bought.

The cases for inertia and irrationality have been overstated. Consumers certainly have the discretion to structure policies in a variety of ways at the time of initial purchase. An individual who does not wish to maintain some element of term insurance after retirement could simply elect some sort of endowment policy. Yet these policies are relatively unpopular. Moreover, many individuals do cash out their policies prior to death. Indeed, the ACLI [1979] reported that in 1977, the value of surrendered policies was $4.3 billion, or roughly 44% of contemporaneous death benefits. In many cases, individuals chose to reconstitute their policies with lower coverage levels at the time of surrender. All of these factors would seem to reflect conscious choices.

Nevertheless, inertia and irrationality undoubtedly explain some fraction of life insurance holdings. The effect of this is to bias the empirical analysis against the bequest motive hypothesis. If life insurance reflects earlier efforts to protect human capital, then one would expect it to rise with lifetime resources. Moreover, Social Security survivor benefits would tend to substitute for life insurance.[6] If, on the other hand, most purchases result from high pressure sales, then one would expect salesmen to have more success with those who have resources to spend. Once again, insurance holdings would rise with resources. Since Social Security entitlements are presumably unrelated to gullibility, it would also be surprising to find any systematic relationship between Social Security and life insurance.

B. Pensions

In the U.S. economy, it is very rare for individuals to purchase annuities directly from insurance agents or carriers. Virtually all private annuities are provided through employee pension programs. In many cases, employees have little or no choice about the nature or extent of their participation in these programs. Consequently, one might well argue that private pensions should be thought of as a form of compulsory annuitization, akin to Social Security, rather than as voluntary private purchases of annuities.

This view undoubtedly overstates the inflexibility of the private pension system. First, pension programs may offer recipients various choices concerning the nature of benefits (for example, a TIAA-CREF participant may reduce his degree of annuitization by electing an "n year certain" policy -- see Bernheim, Shleifer, and Summers [1985]). Second, many defined contribution plans do allow participants to choose their levels of participation. Third, even if each employer offers an inflexible plan, individuals may still choose among various employers. When competitive forces work properly, different employers offer a range of compensation packages (consisting of wages, pensions, and other benefits), so that the tastes of each individual are accommodated.

I do not mean to push the case for flexibility too enthusiastically. The degree of flexibility in the pension system is an empirical issue; the extent to which this system substitutes for a private annuity market can only be determined through analysis of data, such as that undertaken in the following section. However, as with life insurance, one must be aware that individuals may hold private pensions for reasons unrelated to those discussed in section 2A, and that this may affect econometric estimates.

One would expect non-discretionary pension benefits to rise with lifetime resources (indeed, defined benefit plans base payments on terminal wages). Since all models discussed in section 2 have the same implication, this does not bias the analysis in one direction or another. One might also argue that Social Security benefits do not affect the level of compulsory pension participation, so that estimates of the relationship between Social Security and private annuitization should not be contaminated by spurious annuity holdings. Unfortunately, this argument ignores the existence of "integrated" pension plans, in which pension benefits are reduced automatically in response to increases in Social Security benefits. Integration provisions build in the response predicted by the bequest model, as well as several variants of the life cycle hypothesis.

Fortunately, integrated pension plans were rather uncommon for workers retiring in the early 1970's. Kotlikoff and Smith [1983] have found that in 1977, only about 12% of all plans were integrated. Since integration is a relatively new development, the relevant fraction for our sample of relatively elderly individuals is probably much lower. Moreover, the average offset in an integrated plan is about 50 cents on the dollar. Thus, for our sample, integrated plans account for substantially less than a 6 cent on the dollar rate of substitution between Social Security and private pension annuities.

5. Basic Estimates

In this section, I test between the bequest motive hypothesis and various versions of the life cycle hypothesis by estimating behavioral relationships describing holdings of life insurance and private pensions. For the time being, I make no attempt to correct for the various sources of data contamination described in section 4. I have already mentioned that this should bias my results against the bequest motive hypothesis. Section 6 contains estimates based upon a procedure that is designed to correct for data contamination.

A. Estimation Strategy

I begin by formulating an empirical model based upon the theoretical framework developed in section 2. This analysis suggests that holdings of life insurance and pensions both reflect a latent demand for annuities. I therefore posit a latent demand relationship:

$$A = f[LR(q), X, q, \theta] + \epsilon \ ,$$

where A is level of annuities demanded (in dollars of income per year), LR is lifetime resources, X is a vector of exogenous variables, q is the relevant annuity price, and ϵ is a stochastic error term. Note that the actuarial value of lifetime resources depends upon q. Specifically, the survival probabilities used to make the actuarial calculation should be those implicit in the annuity price. This is appropriate, since q measures the rate at which an individual can reallocate consumption intertemporally. I suppose throughout that the demand for annuities is decreasing in q.

When individuals hold private annuities (pensions), the relevant value of q is the price at which additional annuities can be purchased on the margin (henceforth labelled q^P). Due to market imperfections, q^P will generally exceed q^F, the actuarially fair price. When individuals hold life insurance, the relevant value of q is the implicit price for which annuities can be sold through the purchase of life insurance (henceforth labelled q^L). In this case, market imperfections imply that $q^L < q^F$.

Given this framework, the demands for life insurance and pensions are determined as follows:

$$\phi L = \max\{\ 0\ ,\ -\ f[LR(q^L),X,q^L,\theta] + SSB - \varepsilon\ \}$$

$$P = \max\{\ 0\ ,\ f[LR(q^P),X,q^P,\theta] - SSB + \varepsilon\ \},$$

where L is the face value of term life insurance, ϕ is a scalar, and P and SSB are the yearly income flows yielded by private pensions and social security, respectively. The analysis in the appendix implies that $\phi = (1 - \pi^L)/q^L$, where π^L is the single year survival probability implicit in the price of term life insurance. Note that either P or L may be positive, but never both.

It is possible to obtain consistent estimates of either relationship in isolation through the use of standard single equation techniques (see Heckman [1976]). Yet theory clearly implies that these estimates should satisfy a variety of cross-equation restrictions. In addition to testing for the general behavioral patterns discussed in section 2, one could also in principle construct more demanding tests of the underlying theory based upon these restrictions. In essence, life insurance and annuities should be two sides of the same coin -- theory suggests that there is only one underlying relationship. With appropriate restrictions, the qualitative choices described above could be modelled as an ordered probit, and some analogous generalization of the standard Tobit model should describe quantitative choices.

Unfortunately, it is readily apparent that cross-equation restrictions will be rejected. The most basic implication of these restrictions is that households will never hold both life insurance and pensions. We already know that this is severely at odds with the facts. As discussed in section 4, the pension and life insurance variables are both very imperfect measures of discretionary insurance holdings, and the likely sources of contamination are different in each case. In particular, pensions and the cash value portions of life insurance are both vehicles for saving. Consequently, one should not be terribly surprised if many variables push life insurance and pensions in the same direction, rather than in opposite directions (as predicted by theory).

These considerations suggest that it is best to focus on each relationship in isolation, and to test broad-brush behavioral patterns, rather than specific parametric restrictions. I adopt this strategy throughout the remainder of this paper. In the current section, I proceed by estimating relationships of the following form:

$$Y = \max\{0, \alpha + (\beta + X\gamma + \eta)LR + \delta SSB\}. \quad (9)$$

Y denotes the dependent variable, which in some cases is L, and in others is P. For equations involving L, the interpretation of this relationship is that $\phi^{-1}f(LR,X,q^L,\theta) = -\alpha - (\beta + X\gamma)LR$, and $\phi^{-1}\varepsilon = -LR\eta$. Moreover, one would expect to find $\delta = \phi^{-1}$. For equations involving P, the interpretation is that $f(LR,X,q^P,\theta) = \alpha + (\beta + X\gamma)LR$, and $\varepsilon = LR\eta$. In these cases, one would expect to find $\delta = 1$.

Several remarks are in order. First, note that the price of annuities, q, is not included in equation (9). This variable is simply unavailable. In writing (9), I have implicitly assumed that differences in q^L among those who purchase life insurance (as well as differences in q^P among those who purchase pensions) are either non-existent, or unrelated to other explanatory variables (and thus subsumed in η). I also calculate the actuarial value of lifetime resources using population survival probabilities, rather than those implicit in effective annuity prices. This should affect little more than the scale of some coefficients.

Second, note that in equation (9), the effect of the error term η, as well as that of each exogenous variable in the vector X, is proportional to lifetime resources. I chose this specification on the basis of the view that preferences may well be nearly homothetic. I do not, however, impose homotheticity. The intercept parameter α (which is omitted in some specifications) allows the income elasticity of asset demand to differ from unity. In addition, I always include LR as a component of X, so that asset demand is expressed as a quadratic function of resources.

Estimation of (9) requires a distributional assumption on η. If one assumes that η is normally distributed, then three alternative estimation strategies are available. Using only qualitative information, one can estimate a Probit relationship, and identify the parameters of (9) up to a multiplicative constant. Alternatively, one can incorporate quantitative information, and use either the Tobit or Heckit procedures (see Heckman [1976]). The Heckit procedure has the advantage of being robust with respect to a larger class of sample selection processes. With contaminated dependent variables, this is potentially important.

B. Estimates of the Life Insurance Relationship

Table 1 presents probit estimates explaining holdings of life insurance among the RHS sample. The most striking feature of this table is that the coefficient of social security benefits (SSB) is positive in all four specifications, and statistically significant in three of these. The level of statistical significance is extremely high in specifications 3 and 4, which allow α to be estimated freely. These findings are highly favorable to the bequest motive hypothesis.

The marginal effect of resources on life insurance holdings, henceforth labelled ω, is computed as follows:

$$\omega = \beta + 2\mu_{LR}\gamma_{LR},$$

where μ_{LR} is the sample mean of LR, and γ_{LR} is the element of γ corresponding to LR. For specifications 1 and 3, ω measures the derivative of life insurance holdings with respect to resources, evaluated at the sample means for age and resources (recall that AGE is measured as deviations from the sample mean, so its coefficient does not enter into this calculation). The same interpretation applies for specifications 2 and 4, with the qualification that the calculated value of ω describes the behavior of the average childless couple. For these specifications, marital status and the presence of children alter the marginal resource effect.

Note that the marginal resource effect is positive and statistically significant in all four specifications. Strictly speaking, this finding is inconsistent with the theory outlined in section 2A. One should recall, however, that the biases from contamination discussed in section 4 probably affect the coefficient of LR more than the coefficient of SSB (since the displacement of saving through life insurance by social security is probably very small, one would still expect to find a positive coefficient for SSB even in the presence of severe contamination). On balance, the evidence therefore seems favorable to the bequest motive hypothesis.

The effects of other variables are also of interest. The probability of holding life insurance falls with age. This could reflect dissaving, or the mechanical cashing in of policies shortly after retirement. On the other hand, this coefficient is also consistent with the effect predicted by the bequest motive hypothesis (I refer the reader back to the appendix). The presence of children raises the probability of holding life insurance. This seems to confirm the plausible hypothesis that individuals with children have stronger bequest motives than do those without children. On the other hand, those with children had higher living expenses, and hence greater incentives to insure human capital earlier in life. Thus, this effect can also be explained by the hypothesis that insurance holdings among the elderly reflect the unintended residue of decisions made much earlier in life. The negative coefficient of WID no doubt reflects the fact that life insurance pays off upon the death of a spouse, so that the only remaining in-force insurance is on the surviving spouse. The coefficient of SINGLE is positive and statistically significant. One possible interpretation is based on the hypothesis of intrafamily exchange: single individuals may be more dependent on others for various services, and therefore more likely to manifest a strategic bequest motive.

Table 2 contains Tobit estimates of the life insurance relationship. Once again, the coefficients of SSB are uniformly positive -- in this case, all of them are statistically significant at extremely high levels of confidence. One interprets magnitudes as follows. If life insurance purchases substitute perfectly for the sale of social security benefits, then the coefficient of SSB should be ϕ^{-1}. As noted earlier, $\phi^{-1} = q_L/(1-\pi_L)$. Given average life expectancies for individuals in their mid-60's, as well as standard loading factors for life insurance, and a real discount rate of 3%, it appears that the value of ϕ^{-1} should be in the neighborhood of 10. Accordingly, the estimates in Table 2 imply that life insurance purchases offset social security at a rate somewhere between ten and twenty cents on the dollar. Although this effect is fairly small, it is estimated very precisely.

I turn next to the marginal resource effect. Note that ω is negative in three of the four specifications. Although it is positive in the first specification, the estimate lacks statistical significance. In contrast, the negative values of ω in the second and fourth specifications are highly significant. While this finding is very favorable to the bequest motive hypothesis, recall that it refers to the behavior of the average childless couple. As before, the coefficients of KID are uniformly positive. Moreover, the magnitudes of these coefficients are such that the marginal resource effect turns out to be positive (although not statistically significant) for couples with children. Evidence based upon estimates of the marginal revenue effect is therefore mixed, and somewhat inconclusive.

The sign patterns for other coefficients are essentially the same as with the Probit estimates. Similar interpretations apply. The only substantial difference is that the coefficients of SINGLE are now statistically insignificant.

Table 3 contains Heckit estimates of the life insurance relationship. Note again that the coefficients of SSB are positive in all four cases, and highly statistically significant in three cases. Magnitudes are generally larger than for the Tobit estimates. On the basis of specification 3, one would conclude that life insurance purchases substitute for social security at a rate in excess of 40 cents on the dollar.

The estimated marginal resource effects in Table 3 are again negative in three of the four specifications. Note in particular that the negative values of ω in specifications 1 and 3 are statistically significant. This is of special interest, since the omission of KID, WID, and SINGLE implies that ω measures an average effect for the entire population. One can therefore conclude with a high level of statistical confidence that life insurance is on average an inferior good for the elderly population. This bolsters the bequest motive hypothesis. On the other hand, estimates based on specifications 2 and 4 imply that the marginal resource effect is positive for couples with children. Overall, the evidence on marginal resource effects remains somewhat ambiguous.

The sign patterns for AGE and WID are essentially the same as for the other estimation techniques, although the coefficient of WID becomes insignificant in specification 4. The estimated effect of children is now statistically insignificant, and the sign of this coefficient is perverse in specification 4. The effect of SINGLE is now opposite that found in tables 1 and 2, and the coefficient is significantly negative in specification 4. Generally, Heckit estimates for the fourth specification -- which provide the weakest support for the bequest motive hypothesis -- seem anomalous.

Overall, there is a strong positive relationship between life insurance and social security benefits. This finding is robust across specifications and estimation techniques. As argued in section 2, it is very difficult to account for this relationship without positing an effective bequest motive. In addition, there is some evidence that life insurance may be an inferior good for significant segments of the elderly population. Since alternative theories uniformly imply that life insurance holdings should rise with resources, this finding is difficult to explain unless one assumes that behavior is governed at least in part by bequest motives. The fact that life insurance is not universally inferior can easily be attributed to contamination of the dependent variable. This possibility is pursued in section 6.

C. Estimates of the Pension Relationship

Table 4 presents Probit estimates explaining pension holdings in the RHS sample. In specifications 1 and 2, the coefficients of SSB are negative and statistically significant, as predicted. Allowing α to be estimated freely substantially weakens this result. Indeed, the corresponding coefficient is essentially zero and has a very large standard error in specification 4. Taken together, the Probit estimates provide suggestive but inconclusive evidence for the view that Social Security substitutes for private pensions.

All theories discussed in section 2, and all sources of contamination mentioned in section 4, imply that pensions should rise with resources. It is therefore reassuring to find that the estimated marginal resource effects are all positive and extremely statistically significant. The coefficients for AGE are uniformly negative and significant. Since the private pension system grew rapidly during the working lives of the individuals in this data sample, this could conceivably reflect a cohort effect. The positive coefficients for KID are difficult to explain, and in particular appear to contradict the bequest motive hypothesis. Note however that these coefficients lack statistical significance. The failure of many individuals to elect survivorship options for private pensions may well account for the negative coefficients on WID. Finally, the positive and statistically significant effect of SINGLE is somewhat puzzling.

Table 5 contains Tobit estimates of the pension relationship. The most striking feature of this table is that the coefficients of SSB are all negative and statistically significant. The magnitudes of these coefficients imply that social security displaces private pensions at a rate somewhere between ten and twenty cents on the dollar. Coincidentally, this range is roughly the same as that obtained for life insurance using the Tobit procedure. While the effects are relatively small, they are estimated rather precisely. Note that all marginal resource effects are once again positive and highly statistically significant. The sign patterns for the other coefficients are essentially the same as in Table 4.

I present Heckit estimates of the pension relationship in Table 6. Once again, all four coefficients of SSB are negative, and three of the four are statistically significant. When α is estimated freely, I find rates of displacement in excess of thirty cents on the dollar. Estimated marginal resource effects remain positive and significant. The sign patterns for other variables are essentially unchanged.

Overall, there is a strong negative relationship between pension benefits and social security entitlements. This finding contradicts the implications of simple life cycle models, since it suggests that consumers do not wish to hold all of their assets as annuities (a life cycle consumer would exhaust all conventional assets before reducing annuities). However, the evidence does not distinguish between the bequest motive hypothesis and various more sophisticated versions of the life cycle hypothesis that blunt the strict preference for annuities.

6. Refined Estimates

While most of the evidence presented in section 5 favors the bequest motive hypothesis, there are still a few points of ambiguity. This is hardly surprising. As discussed in section 4, data contamination is probably severe, and this tends to bias the results against the bequest motive hypothesis. In this section, I design and implement a procedure for estimating the relationship of interest while simultaneously decontaminating the dependent variable. Since estimates of the pension relationship cannot distinguish between the bequest motive hypothesis and various other theories of interest, I focus exclusively on life insurance.

A. Estimation Strategy

I begin by defining two latent variables, L^1 and L^2. L^1 measures the latent demand for discretionary term insurance, while L^2 measures the latent demand for other components of measured life insurance (presumably, this is mostly cash value or saving, but it could also represent life insurance held due to inertia or irrationality). These variables are described by the following relationships:

$$L^k = \alpha^k + (\beta^k + X\gamma^k + \eta^k)LR + \delta^k SSB \,, \tag{10}$$

where k = 1,2. In practice, neither component of life insurance can ever be negative. I assume that it is impossible to observe the two components in isolation. One can only observe the truncated sum:

$$L = \max\{\, 0 \,,\, L^1 \,\} + \max\{\, 0 \,,\, L^2 \,\}. \tag{11}$$

If one imposes a distributional assumption on the η's, then it is possible to estimate this system by maximizing the appropriate likelihood function. Unfortunately, there is a severe identification problem. While estimation yields two sets of parameters, $\Theta^s = (\alpha^s, \beta^s, \gamma^s, \delta^s, \sigma_\eta^s)$, s=A,B, it is impossible to determine which parameters (Θ^A or Θ^B) correspond to the term insurance relationship (Θ^1), and which correspond to the relationship describing other components of observed insurance holdings (Θ^2).[7]

Fortunately, one does not require full identification to test the theories of interest. The central implications of the bequest motive hypothesis are that $\delta^1 > 0$, and $\omega^1 < 0$. Moreover, under the view that the cash value component of life insurance is discretionary saving, theory also implies that $\delta^2 < 0$ and $\omega^2 > 0$. Without knowing which of the estimated relationships is which, we can therefore test the following implications:

$$\lambda_1 = \omega^A \omega^B < 0 \tag{12}$$

(13) $$\lambda_2 = \omega^A \delta^A < 0$$

(14) $$\lambda_3 = \omega^B \delta^B < 0$$

(15) $$\lambda_4 = \delta^A \delta^B < 0.$$

To construct the likelihood function for the system described by equations (10) and (11), I assume that the η's are normally distributed. Although I attempted to estimate a parameter measuring correlation between the η's, this proved problematic. Fortunately, it was possible to obtain estimates by imposing values for the correlation coefficient, and the results were not particularly sensitive to changes in the assumed value of this parameter. For convenience, I have chosen to present results based on the assumption of no correlation.

B. Estimates

One can obtain a better understanding of the formal hypothesis tests by first studying the parameter estimates informally. These appear in Table 7. Each of the four estimated systems contain an equation with $\delta > 0$ and $\omega < 0$. This, by itself, is quite striking confirmation of the theory. I have listed these equations as the "A" components of each system, and I informally interpret them as term life insurance relationships.

The coefficients of SSB in the A equations are not only positive, but also highly statistically significant, and roughly twice as large as the corresponding coefficients obtained through the simple Tobit procedure (Table 2). These estimates suggest that life insurance purchases offset incremental social security benefits at a rate somewhere between twenty and thirty cents on the dollar. The strengthening of this effect (both absolutely and in terms of statistical significance) suggests that the decontamination procedure is at least partially effective.

Decontamination has a much greater impact on the estimated marginal resource effects. Compared with the simple Tobit estimates, the absolute values of the ω's in the A equations increase by almost an order of magnitude. All of these effects are statistically significant at extraordinarily high levels of confidence. Moreover, this component of life insurance now appears to be an inferior good for all population groups, including couples with children.

All other coefficients in the A equations are also consistent with the view that these relationships measure the latent demand for term life insurance. The coefficients of AGE are negative and significant, as predicted (see the appendix). The presence of children raises life insurance holdings, reflecting a stronger bequest motive. One would still expect the coefficients of WID to be negative for the reasons mentioned previously. Finally, the effect of SINGLE is statistically insignificant.

I turn next to the B equations. In systems 1 and 2, the coefficients of SSB are negative and statistically significant. This is consistent with the view that social security depresses private saving through the vehicles of whole and endowment life insurance. However, the coefficients of SSB are positive in systems 3 and 4. The effect in system 3 is quantitatively small, but statistically significant, while the effect in system 4 is insignificant both economically and statistically. This should not be too surprising. Even if social security depresses an individual's total saving dollar for dollar, it would have a very small effect on any given component, such as life insurance. Moreover, contamination might also result from the "residue" effect, in which case one would not expect to discover a systematic relationship between life insurance and social security. Failure to find compelling evidence that $\delta^2 < 0$ should not, therefore, be interpreted as damaging to the bequest motive hypothesis. Indeed, the large body of other evidence indicating that social security depresses an individual's total saving (at least to some extent) strongly supports this component of the overall theory.

As expected, the marginal resource effects are positive and highly statistically significant in the B equations of all four specifications. Given the sources of data contamination discussed in section 4, this is exactly what one would expect. Moreover, this finding helps to explain why single equation procedures (section 5B) yield inconclusive results concerning ω.

The remaining coefficients are also consistent with the view that the B equation measures the latent demand for the non-term component of life insurance. The negative coefficients on AGE presumably capture the effects of dissaving, or cashing out shortly after retirement. The positive coefficients of KIDS could well reflect the "residue" effect discussed previously. As before, one would expect the WID coefficient to be negative for fairly mechanical reasons. Finally, the coefficients of SINGLE are again statistically insignificant.

In summary, an informal inspection of the estimates reveals the following. Every system contains one equation that behaves exactly as term life insurance demand should under the bequest motive hypothesis. In addition, the second equation of each system behaves almost exactly as one would predict given the contaminating factors discussed in section 4. This evidence argues very strongly in favor of the bequest motive hypothesis.

Thus far, I have discussed these estimates informally under the assumption that one can interpret the A equations as the term insurance relationships, and the B equations as the cash value relationships. I now turn to the formal hypothesis tests that justify this interpretation.

The model described in section 6A presupposes the existence of two separate life insurance processes. If there is in fact only one process, then the simple Tobit model is appropriate. Note that the Tobit model is nested within the more general model. One can therefore test for the presence of a second relationship by constructing a likelihood ratio test. I illustrate with system 1. The test statistic ($-2\log\lambda$, where λ is the likelihood ratio) is 7148. Since this statistic has a χ^2 distribution with 5 degrees of freedom, one can conclude with overwhelming confidence that there are at least two separate processes at work. The same conclusion follows for systems 2, 3, and 4.

I now turn to the specific tests described in section 6A. In systems 1 and 2, λ_1, λ_2, and λ_3 are all negative, and each, taken individually, is statistically significant. Of course, I wish to test these three hypotheses jointly. It is possible to calculate a conservative joint confidence region by taking the Cartesian product of confidence intervals for the three individual parameters. Specifically, for any $0 < C < 1$, there is at least a $100C^3$% probability that the true parameter lies within the Cartesian product of the 100C% confidence intervals for the three parameters. In essence, one conducts the joint test by dividing up the significance level three ways. For both systems 1 and 2, if one chooses any reasonable value of C, the 100C%

confidence region lies entirely within the negative orthant. Alternatively, one can reject with high confidence any specific hypothesis in which any of the λ's is positive. These results provide striking statistical confirmation of the underlying theory.

For systems 3 and 4, λ_1 and λ_2 are negative and individually statistically significant (as predicted), but λ_3 is positive. This reflects the fact that we have estimated $\delta^B > 0$, as discussed above. One should not be inclined to discount other evidence favoring the bequest motive hypothesis on this basis.

If we do not wish to insist on testing the somewhat tenuous hypothesis that $\delta^2 < 0$ jointly with the other three inequalities ($\delta^1 > 0$, $\omega^1 < 0$, $\omega^2 < 0$), then one must abandon the formulation described in equation (12) through (15), and proceed differently. Specifically, define the indicator function I(.) as follows:

$$I(x) = \begin{cases} 0 \text{ if } x > 0 \\ 1 \text{ otherwise.} \end{cases}$$

Under the new joint hypothesis, (12) must still be satisfied, along with the following inequalities:

(16) $$\lambda_2^* = I(\omega^1)\delta^1 \geq 0$$

(17) $$\lambda_3^* = I(\omega^2)\delta^2 \geq 0$$

These inequalities essentially say that δ must be positive in any equation for which ω is negative. Together, (12), (16), and (17) form the revised joint hypothesis.

Calculating asymptotic standard errors for the λ_i^* parameters is extremely straightforward. To apply standard asymptotic results, one only requires that I(.) has continuous second order derivatives in some neighborhood of the true parameter (see Theil [1971, p.383]). As long as the true value of ω is non-zero, this condition is satisfied; moreover, derivatives of I are equal to zero. Thus, the asymptotic variance of λ_i^* is simply equal to the asymptotic variance of δ^i when ω^i is positive, and zero otherwise. Accordingly, one tests this weaker formulation of the underlying theory as follows: construct a joint confidence region by taking the Cartesian product of confidence intervals for λ_1 and δ^i, choosing i such that ω^i is positive.

The logic of this test is straightforward. One tests the hypothesis that $\delta^i > 0$ conditional upon having found $\omega^i < 0$. In small samples, conditioning will change the distribution of the estimator. However, when the true value of ω^i is negative, asymptotically almost all probability mass is placed on negative values. Therefore, asymptotically, the fact that $\omega^i < 0$ conveys almost no information, and in the limit the conditional distribution of δ^i is equal to the unconditional distribution.

Recall that I have labelled equations so that $\omega^A < 0$, and $\omega^B > 0$. The parameters of interest are therefore λ_1 and δ^A. The estimates of all four systems are very favorable to the hypothesis that $\lambda_1 < 0$ and $\delta^A > 0$ -- confidence regions for all reasonable levels of statistical significance lie entirely within this quadrant. One can therefore reject with high confidence any other specific hypothesis for which either $\lambda_1 > 0$ or $\delta^A < 0$.

It is useful to summarize these findings in words. The data overwhelmingly support two hypotheses: first, that there are two separate life insurance processes, and second, that one component of life insurance is inferior and rises with social security benefits, while the other component is normal. These findings are inconsistent with all theories of life cycle behavior that omit significant bequest motives.

Assuming that the A relationships measure the latent demand for term life insurance, one can use these relationships to construct indices that measure the strength and extent of bequest motives. I make two illustrative calculations. First, for an average individual (one whose characteristics equal sample averages), I calculate the probability that term life insurance purchases are positive. Second, I estimate the fraction of the sample with positive term life insurance holdings. To obtain this estimate, I calculate the probability that L^A is positive for each individual in the sample, sum these probabilities, and divide by the sample size.

These calculations are of interest for the following reason. The analysis in section 2 suggests that term life insurance holdings will be positive only when an individual's desired bequest substantially exceeds his conventional asset holdings. Since the value of conventional assets typically amounts to about half of total wealth, those who intentionally purchase term life insurance have very strong bequest motives. Note that the lack of term insurance does not imply the absence of a bequest motive. Desired bequests may simply be less than conventional asset holdings. Even when they slightly exceed the value of conventional assets, individuals may still be unwilling to enhance bequests by paying the premium (over and above the actuarially fair rate) on life insurance. Thus, my calculations indicate the number of individuals (as a fraction of the population) who have very strong bequest motives.

Results appear in Table 8. I provide separate calculations for households with and without children. Approximately one quarter of the sample overall has positive discretionary term life insurance holdings. This suggests that approximately two-thirds of all life insurance policies owned by individuals in the RHS sample functioned exclusively as savings accounts, or were maintained because of inertia or irrationality. These calculations document the severity of data contamination. More importantly, they demonstrate that strong bequest motives are quite common. For close to 30% of households with children, desired bequests substantially exceeded the value of conventional assets. A surprising number of childless households (roughly 16 to 18%) also acted to augment their bequests. A strong bequest motive among childless households is consistent with models of intrafamily exchange, where services can be obtained from family members other than children.

APPENDIX

Consider an individual who wishes to allocate consumption in discrete time over a potentially infinite horizon. Let a_t, b_t, and ℓ_t be net purchases of annuities, conventional assets, and term life insurance in period t. One unit of annuities purchased in t yields a stream of income equal to \$1 per year in all years after t, conditional upon survival. The price of the annuity is denoted q_t. A conventional asset valued at \$1 yields a stream of income equal to $\$\rho$ in all years after t -- this stream is not conditional upon survival. One unit of life insurance costs p_t, and pays \$1 should the individual die before the beginning of period t+1.

Total annuity and conventional asset holdings in period t are denoted A_t and B_t, and are given by

(A.1) $$A_t = A_{t-1} + a_t$$

(A.2) $$B_t = B_{t-1} + b_t$$

If the individual should die in period t, his bequest is

(A.3) $$D_t = B_t + \ell_t$$

Initial resources, W_0, must be allocated as follows:

(A.4) $$W_0 = q_0A_0 + B_0 + p_0\ell_0$$

Subsequently, expenditures must equal income in every period:

(A.5) $$C_t + q_ta_t + b_t + p_t\ell_t = A_{t-1} + pB_{t-1}$$

If life insurance and annuities are actuarially fair, then

(A.6) $$p_t = \pi_t,$$

and

(A.7) $$q_t = \sum_{\tau=1}^{\infty} (1+\rho)^{-\tau} \prod_{k=1}^{\tau} (1-\pi_{t+k-1}),$$

where π_t is the probability that the individual dies before t+1, conditional upon reaching t. The assumption that insurance is actuarially fair is inessential. One may think of π_t as the mortality probability implicit in life insurance prices. The preceding pricing relationships indicate that annuity prices are based on the same implicit probabilities. As in section 2A, this simply allows me to abstract from income effects.

Now suppose that the government raises A_0 by one unit. If the individual is prevented from selling annuities, he can still offset the change as follows (Δ denotes a change from the original program):

(A.8) $$\Delta B_0 = -\frac{q_0}{1-\pi_0},$$

and for each t,

(A.9) $$\Delta b_t = \frac{q_{t-1}}{1-\pi_{t-1}} - \frac{q_t}{1-\pi_t}$$

(A.10) $$\Delta \ell_t = \frac{q_t}{1-\pi_t}$$

(A.11) $$\Delta C_t = \Delta a_t = 0.$$

To see that this offsets the annuity, first note that, by (A.2), (A.8), and (A.9),

$$\Delta B_t = - \frac{q_t}{1 - \pi_t}$$

Thus, by (A.10),

$$\Delta \ell_t + \Delta B_t = 0.$$

(A.3) then implies that $\Delta D_t = 0$ for all t.

Next, from equations (A.8) and (A.10), we have

$$q_0 + p_0 \Delta \ell_0 - \Delta B_0 = 0.$$

(A.4) then implies that the portfolio is feasible in period 0. Finally,

$$\Delta b_t + p_0 \Delta \ell_t = \frac{q_{t-1}}{1 - \pi_{t-1}} - q_t$$

$$= 1 - p \frac{q_{t-1}}{1 - \pi_{t-1}}$$

$$= \Delta A_t - p \Delta B_t$$

(this uses the fact that $q_t = q_{t-1}(1+\rho)/(1 - \pi_{t-1}) - 1)$. Thus, the new portfolio is feasible in every period. The sale of an annuity has therefore been accomplished through the purchase of life insurance.

Note finally that if π_t rises with t, then $\Delta \ell_t$ falls with t. I allude to this property at several points in the text.

REFERENCES

Abel, Andrew B. (1985), "Precautionary Saving and Accidental Bequests," American Economic Review 75 (4), pp. 777-91.

Akerlof, George A. (1970), "The Market for 'Lemons': Qualitative Uncertainty and the Market Mechanism," Quarterly Journal of Economics 84 (3), pp. 488-500.

American Council of Life Insurance (1979), 1979 Life Insurance Fact Book, American Council of Life Insurance: Washington, D.C.

Auerbach, Alan J. and Laurence J. Kotlikoff (1985), "Life Insurance of the Elderly: Adequacy and Determinants," mimeo.

Bernheim, B. D., A. Shleifer, and L. Summers (1985), "The Strategic Bequest Motive," Journal of Political Economy 93 (6), pp. 1045-76.

Bernheim, B. Douglas (1987), "The Economic Effects of Social Security: Toward a Reconciliation of Theory and Measurement," Journal of Public Economics 33, pp. 273-304.

Biggs, John H. (1973), "Do Social Security Benefits Crowd Out Private Life Insurance?" mimeo.

Davies, J. (1981), "Uncertain Lifetime Consumption, and Dissaving in Retirement," Journal of Political Economy 89, pp. 561-77.

Diamond, Peter and Jerry Hausman (1984), "Individual Retirement and Savings Behavior," Journal of Public Economics 23, pp. 81-114.

Eckstein, Zvi, Martin S. Eichenbaum, and Dan Peled (1985), "Uncertain Lifetimes and the Welfare Enhancing Properties of Annuity Markets and Social Security," Journal of Public Economics 26(3), pp. 303-26.

Feldstein, Martin (1974), "Social Security, Induced Retirement, and Aggregate Capital Accumulation," Journal of Political Economy 82, pp. 905-26.

Feldstein, Martin and Anthony Pellechio (1979), "Social Security and Household Wealth Accumulation: New Microeconomic Evidence," Review of Economics and Statistics 61 (3), pp. 361-68.

Fitzgerald, John (1987), "The Effects of Social Security on Life Insurance Demand by Married Couples," Journal of Risk and Insurance 54(1), pp. 86-99.

Fox, Alan (1976), "Alternative Measures of Earnings Replacement for Social Security Benefits," in Reaching Retirement Age: Findings of a Survey of Newly Retired Workers, 1968-70, Research Report No. 47, Office of Research and Statistics, Social Security Administration (Government Printing Office).

Heckman, James J. (1976), "The Common Structure of Statistical Models of Truncation, Sample Selection and Limited Dependent Variables and a Sample Estimator for Such Models," Annals of Economic and Social Measurement 5, pp. 475-92.

Hurd, Michael (1987), "Savings of the Elderly and Desired Bequests," American Economic Review 77(3), pp. 298-312.

King, Mervin and Louis Dicks-Mireaux (1984), "Pension Wealth and Household Savings: Tests of Robustness," Journal of Public Economics.

Kotlikoff, Laurence (1979), "Testing the Theory of Social Security and Life Cycle Accumulation," The American Economic Review 69, pp. 396-410.

Kotlikoff, Laurence (1988), "Intergenerational Transfers and Savings," Journal of Economic Perspectives 2, pp. 41-58.

Kotlikoff, Laurence and Daniel Smith (1983), Pensions in the American Economy, NBER - University of Chicago Press: Chicago.

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

Kotlikoff, Laurence and Lawrence Summers (1981), "The Role of Intergenerational Transfers in Aggregate Capital Accumulation," Journal of Political Economy 89, pp. 706-32.

Levin, Laurence (1988), "Precautionary Saving, Health Insurance, and Medical Costs," mimeo, University of Santa Clara.

Modigliani, Franco (1988), "The Role of Intergenerational Transfers and Life Cycle Saving in the Accumulation of Wealth," Journal of Economic Perspectives 2, pp. 15-40.

Munnell, Alicia (1974), "The Impact of Social Security on Personal Savings," National Tax Journal 27 (4), pp. 553-67.

Rothschild, Michael and Joseph E. Stiglitz (1976), "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," Quarterly Journal of Economics 90 (4), pp. 629-49.

Theil, Henri (1971), Principles of Econometrics, John Wiley and Sons, Inc.: New York.

Tomes, Nigel (1981), "The Family, Inheritance, and the Intergenerational Transmission of Inequality," Journal of Political Economy 89, pp. 928-58.

Yaari, M. (1965), "Uncertain Lifetime, Life Insurance, and the Theory of the Consumer," Review of Economic Studies 32, pp. 137-50.

FOOTNOTES

1. Throughout, I use the term "annuity" to describe a contract for survival-contingent income. Yaari [1965] has also referred to this as an "actuarial contract."

2. This observation is occasionally taken to imply that, except under total failure of insurance markets, life cycle savers will completely annuitize. This is not necessarily the case. When the failure arises from asymmetric information, insurance market equilibria may involve quantity constrained contracts (see Rothschild and Stiglitz [1976]). Thus, annuities may be available in the market, yet unavailable on the margin.

3. Fitzgerald [1987] find support for the view that life insurance purchases reflect bequest motives, but, in contrast to this paper, he includes the desire to provide for one's spouse as a bequest motive. Accordingly, his evidence does not shed light on the importance of intergenerational transfers.

4. This does not necessarily imply that the couple names some other individual as beneficiary to a life insurance policy. A husband may name his wife as primary beneficiary, with the understanding that, in the normal course of events, she will eventually bequeath the benefit to their children.

5. Exclusions based upon bequeathable wealth holdings (our third and fifth criteria) may induce sample selection biases, to the extent that bequeathable wealth is determined endogenously. As I have dropped only extreme cases from the sample, this effect is likely to be insignificant.

6. Biggs [1983] found some evidence to support the hypothesis that Social Security survivor benefits substitute for life insurance holdings among individuals with substantial future human capital (his study focused on women who had been widowed between ages 25 and 65). To the extent life insurance holdings among the elderly reflect the unintended residue of decisions made earlier in life, one should find this pattern in the RHS data.

7. In principle, one could solve this problem by adopting different functional forms for the two relationships, or by excluding certain variables from either or both. I regard this potential solution as very unsatisfactory. There is simply no good *a priori* basis for assuming that the two relationships differ in some structural, identifiable way.

Table 1: Life Insurance Probits

Variable	Specification 1	2	3	4
Constants				
$\alpha/10^4$			-6.80 (1.26)	-5.47 (1.30)
β	-0.211 (0.0965)	-0.438 (0.122)	-0.014 (0.105)	-0.281 (0.128)
SSB	13.5 (6.11)	5.92 (6.53)	31.8 (7.05)	22.0 (7.62)
$LR/10^6$	1.27 (0.100)	1.19 (0.109)	0.992 (0.112)	0.994 (0.118)
$AGE/10^2$	-5.91 (1.19)	-5.28 (1.24)	-5.16 (1.20)	-4.73 (1.25)
KID/10		5.99 (0.616)		5.87 (0.616)
WID/10		-6.88 (0.590)		-6.77 (0.590)
SINGLE/10		2.26 (1.00)		2.84 (1.02)
ω	1.32 (0.052)	1.00 (0.075)	1.21 (0.055)	0.923 (0.077)
Log Likelihood	-2097	-1944	-2082	-1935

Table 2: Life Insurance Tobits

Variable	Specification 1	2	3	4
Constants				
$\alpha/10^3$			-1.77	-0.908
			(0.604)	(0.609)
$\beta \times 10^3$	-15.7	-20.8	-11.2	-18.7
	(4.45)	(5.48)	(4.72)	(5.67)
SSB	1.37	1.04	1.88	1.32
	(0.290)	(0.298)	(0.340)	(0.353)
$LR/10^9$	14.2	8.64	8.69	6.22
	(4.26)	(4.48)	(4.66)	(4.77)
$AGE/10^4$	-28.3	-24.6	-26.4	-23.7
	(5.16)	(5.13)	(5.20)	(5.17)
$KID/10^2$		1.92		1.89
		(0.281)		(0.282)
$WID/10^2$		-2.76		-2.75
		(0.273)		(0.274)
$SINGLE/10^2$		0.163		0.261
		(0.465)		(0.470)
$\omega \times 10^3$	1.48	-10.34	-0.64	-11.2
	(2.02)	(3.19)	(2.15)	(3.24)
Log Likelihood	3280	3374	3284	3375
$\sigma \times 10^3$	56.7	55.9	56.7	55.9
	(0.780)	(0.767)	(0.781)	(0.768)

Table 3: Life Insurance Heckits

Variable	Specification 1	2	3	4
Constants				
$\alpha/10^3$			-5.83 (2.52)	3.90 (1.12)
$\beta \times 10^3$	-88.7 (27.6)	3.04 (20.6)	-105 (35.2)	34.5 (20.7)
SSB	2.44 (0.315)	1.76 (0.296)	4.32 (0.974)	0.608 (0.454)
$LR/10^9$	57.0 (18.9)	-8.11 (11.3)	60.2 (20.9)	-19.6 (10.3)
$AGE/10^4$	-54.0 (11.0)	-19.0 (7.13)	-58.9 (13.2)	-11.9 (6.93)
$KID/10^2$		0.791 (0.680)		-0.345 (0.722)
$WID/10^2$		-1.61 (0.775)		-0.239 (0.837)
$SINGLE/10^2$		-0.388 (0.489)		-1.10 (0.551)
$MILLS/10^2$	12.9 (3.35)	1.58 (2.17)	16.7 (5.02)	-2.82 (2.41)
$\omega \times 10^3$	-19.7 (5.28)	-6.78 (7.79)	-32.2 (10.4)	10.8 (9.35)
$\sigma \times 10^3$	52.3	52.3	52.3	52.2

Table 4: Pension Probits

Variable	Specification 1	2	3	4
Constants				
$\alpha/10^4$			-4.52	-6.80
			(1.82)	(1.97)
β	-1.19	-1.41	-1.05	-1.25
	(0.109)	(0.134)	(0.123)	(0.145)
SSB	-25.4	-17.7	-15.3	0.013
	(7.75)	(8.16)	(8.97)	(9.95)
$LR/10^6$	2.02	2.27	1.88	2.09
	(0.106)	(0.116)	(0.122)	(0.128)
$AGE/10^2$	-6.34	-6.45	-6.00	-6.03
	(1.19)	(1.21)	(1.20)	(1.22)
$KID/10^2$		5.62		4.16
		(6.48)		(6.50)
WID		-0.442		-0.433
		(0.0627)		(0.0628)
SINGLE		0.628		0.702
		(0.105)		(0.108)
ω	1.26	1.35	1.22	1.29
	(0.051)	(0.078)	(0.054)	(0.079)
Log Likelihood	-2094	-2037	-2091	-2031

Table 5: Pension Tobits

Variable	Specification 1	2	3	4
Constants				
α			-6.78 (85.8)	-93.8 (88.8)
$\beta \times 10^3$	-4.44 (0.583)	-5.28 (0.694)	-4.39 (0.621)	-5.09 (0.721)
SSB	-0.174 (0.0417)	-0.130 (0.0424)	-0.172 (0.0471)	-0.103 (0.0499)
$LR/10^9$	6.54 (0.530)	7.53 (0.560)	6.52 (0.581)	7.32 (0.596)
$AGE/10^4$	-4.05 (0.639)	-4.07 (0.632)	-4.05 (0.643)	-3.99 (0.637)
$KID/10^4$		1.76 (3.26)		1.51 (3.27)
$WID/10^3$		-2.43 (0.333)		-2.42 (0.334)
$SINGLE/10^3$		3.44 (0.535)		3.54 (0.544)
$\omega \times 10^3$	3.50 (0.24)	3.83 (0.37)	3.50 (0.25)	3.77 (0.37)
Log Likelihood	5298	5361	5298	5362
$\sigma \times 10^3$	6.17 (0.114)	6.05 (0.112)	6.17 (0.114)	6.06 (0.112)

Table 6: Pension Heckits

Variable	Specification 1	2	3	4
Constants				
$\alpha/10^2$			16.1 (1.40)	17.4 (1.41)
$\beta \times 10^3$	-21.5 (2.47)	-13.8 (1.86)	-6.97 (2.28)	-3.92 (1.94)
SSB	-0.378 (0.0812)	-0.0890 (0.0509)	-0.376 (0.0496)	-0.323 (0.0447)
$LR/10^9$	18.2 (1.94)	12.7 (1.48)	7.43 (1.75)	5.32 (1.50)
$AGE/10^4$	-9.60 (0.978)	-6.77 (0.724)	-5.23 (0.845)	-4.12 (0.715)
$KID/10^4$		2.76 (2.61)		1.37 (2.45)
$WID/10^3$		-4.06 (0.409)		-2.04 (0.431)
$SINGLE/10^3$		5.53 (0.605)		1.53 (0.729)
$MILLS/10^2$	1.91 (0.183)	1.22 (0.121)	0.628 (0.181)	0.370 (0.140)
$\omega \times 10^3$	0.54 (0.20)	1.53 (0.29)	2.03 (0.25)	2.52 (0.30)
$\sigma \times 10^3$	4.31	4.34	4.11	4.10

Table 7: Maximum Likelihood Estimates

Variable	1A	1B	2A	2B	3A	3B	4A	4B
Constants								
$\alpha/10^3$					-1.07 (0.588)	-2.92 (0.426)	0.145 (0.607)	-2.45 (0.388)
$\beta \times 10^3$	-39.2 (5.59)	-1.34 (0.805)	-27.6 (6.23)	-1.46 (0.899)	-37.2 (6.75)	4.75 (1.37)	-40.6 (8.11)	4.01 (1.32)
SSB	2.73 (0.292)	-0.417 (0.084)	1.95 (0.277)	-0.505 (0.086)	3.06 (0.312)	0.235 (0.099)	2.31 (0.320)	0.0814 (0.102)
$LR/10^9$	-22.6 (6.14)	7.06 (0.631)	-32.3 (5.63)	6.12 (0.645)	-34.4 (7.44)	1.23 (1.07)	-42.3 (7.55)	1.27 (0.990)
$AGE/10^4$	-30.6 (6.88)	-4.94 (0.857)	-19.2 (6.36)	-5.02 (0.828)	-32.2 (7.45)	-3.40 (0.871)	-27.6 (7.75)	-3.53 (0.833)
$KID/10^2$			1.52 (0.402)	0.237 (0.040)			2.08 (0.509)	0.194 (0.039)
$WID/10^2$			-5.46 (0.520)	-0.279 (0.039)			-5.97 (0.584)	-0.319 (0.039)
$SINGLE/10^2$			-0.836 (0.577)	0.0487 (0.0699)			-0.477 (0.742)	0.130 (0.071)
$\omega \times 10^3$	-66.5 (3.85)	7.20 (0.379)	-66.7 (5.11)	5.95 (0.476)	-78.7 (4.24)	6.23 (0.378)	-91.8 (6.32)	5.55 (0.466)
$\sigma \times 10^3$	60.4 (0.792)	5.76 (0.154)	53.3 (0.701)	5.39 (0.142)	65.2 (0.872)	5.70 (0.143)	63.9 (0.864)	5.52 (0.132)
Log Likelihood	6854		7067		6919		7116	
$\lambda_1 \times 10^4$	-4.79 (0.409)		-3.97 (0.478)		-4.91 (0.424)		-5.10 (0.591)	
$\lambda_2 \times 10$	-1.81 (0.251)		-1.30 (0.204)		-2.41 (0.281)		-2.12 (0.328)	
$\lambda_3 \times 10^3$	-3.01 (0.719)		-3.01 (0.661)		1.47 (0.587)		0.452 (0.555)	

Table 8: True Life Insurance Holdings

Equation System	Children	Probability of positive holdings at mean values	Estimated fraction with positive holdings
3	No	0.133	0.164
3	Yes	0.259	0.284
4	No	0.150	0.182
4	Yes	0.275	0.299

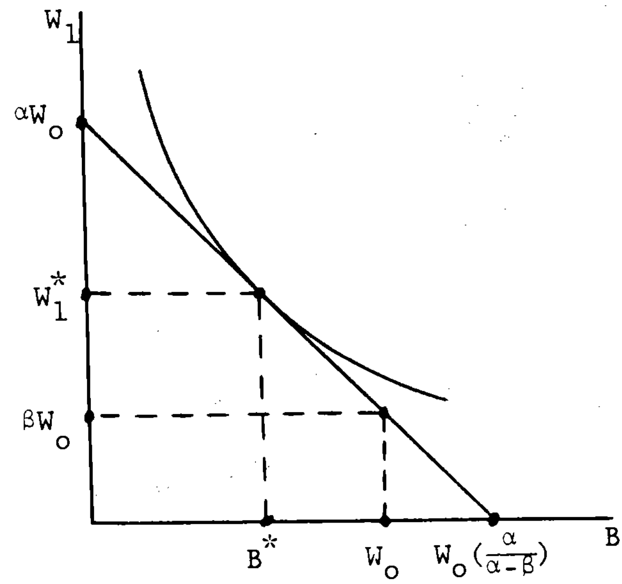

Figure 1

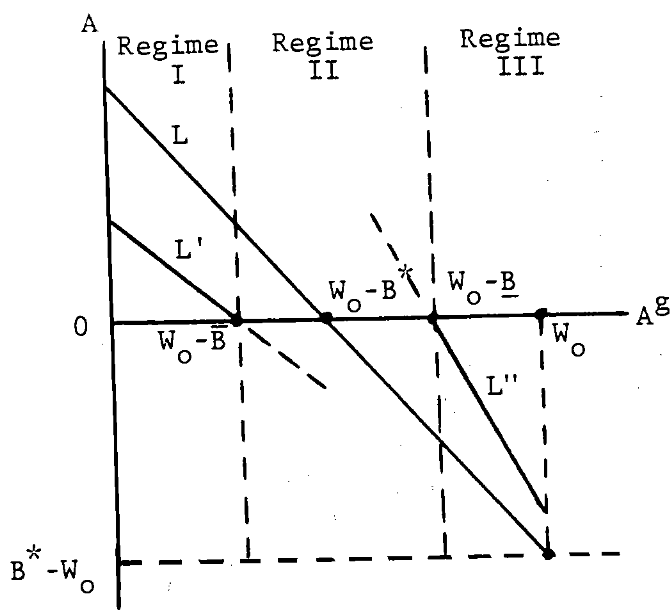

Figure 2

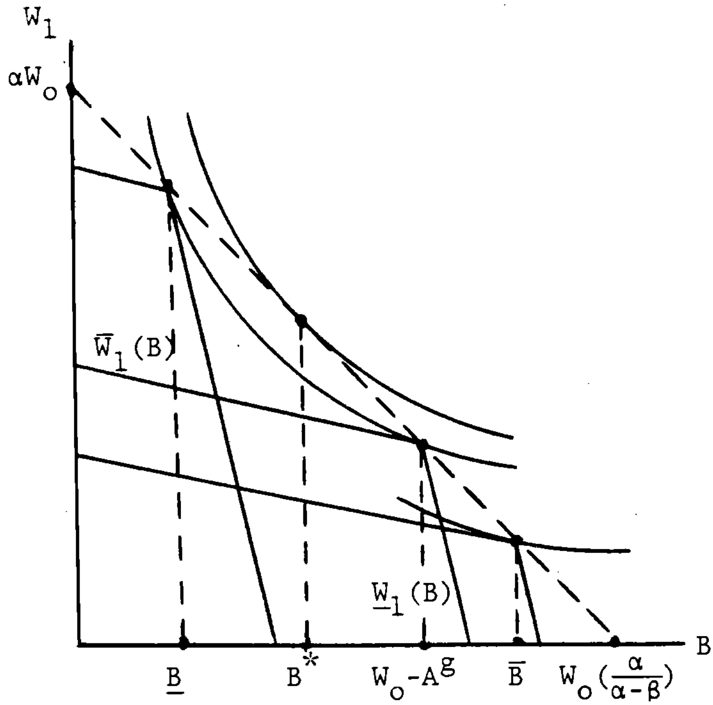

Figure 3

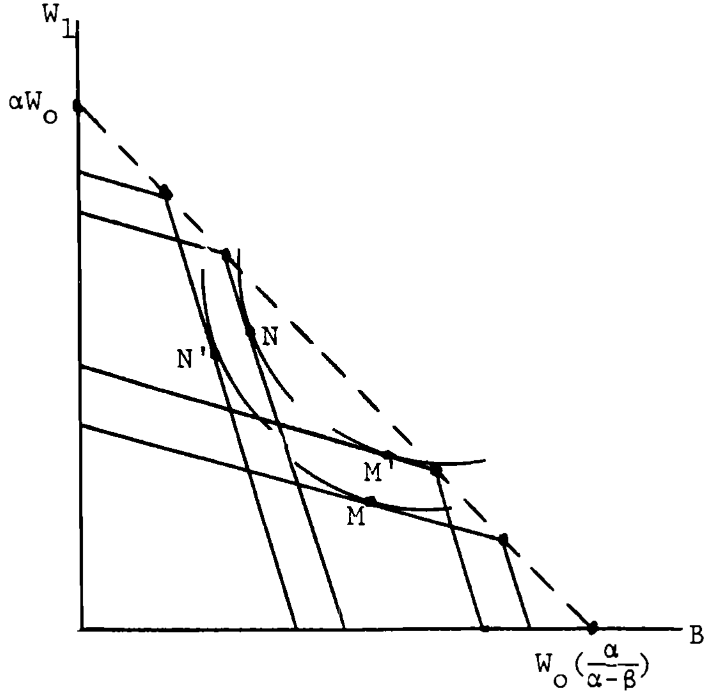

Figure 4