Volume Title: Issues in Pension Economics

Volume Author/Editor: Zvi Bodie, John B. Shoven, and David A. Wise, eds.

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-06284-8

Volume URL: http://www.nber.org/books/bodi87-1

Publication Date: 1987

Chapter Title: Annuity Markets, Savings, and the Capital Stock

Chapter Author: Laurence J. Kotlikoff, John B. Shoven, Avia Spivak

Chapter URL: http://www.nber.org/chapters/c6860

Chapter pages in book: (p. 211 - 236)

8 Annuity Markets, Savings, and the Capital Stock

Laurence J. Kotlikoff, John B. Shoven, and Avia Spivak

8.1 Introduction

This paper examines how the availability of annuities affects savings and inequality in economies in which neither private nor public pensions exist initially. The absence of widespread market or government annuity insurance clearly describes many less developed countries in the world today; it was also characteristic of virtually all countries prior to World War II. While there is now a considerable body of literature addressing the savings impact of funding or not funding government pensions (Barro 1974; Feldstein 1974; and numerous others), the effect of the insurance provision per se has received less attention.

Sheshinski and Weiss (1981) is the first analysis of the pure insurance effects of social security on national saving. They demonstrate that when private arrangements are unavailable, the government's provision of fully funded old age annuities alters household consumption possibilities. In their model in which agents have a bequest motive, the short-run saving impact of such provision is ambiguous. Hubbard (1983) points out that this provision unambiguously reduces national saving if agents have no bequest motive. Fuller descriptions of life cycle (zero

Laurence J. Kotlikoff is professor of economics, Boston University, and research associate, National Bureau of Economic Research. John B. Shoven is professor of economics, Stanford University, and research associate, National Bureau of Economic Research. Avia Spivak is senior lecturer, Department of Economics, Ben Gurion University, and visiting professor, Department of Economics, University of Pennsylvania.

This paper was originally prepared for presentation at the Econometric Society Winter Meetings, December 28–30, 1983, San Francisco. We thank Zvi Eckstein, Truman Bewley, Linus Yamane, and Ludo Van der Heyden for very helpful discussions and Linus Yamane for excellent research assistance. The research reported here is part of the NBER's research program in pensions. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

bequest motive) economies in the absence of annuity insurance are presented in Eckstein et al. (1985) and Abel (1985). Both papers independently derived the stochastic steady state properties of economies in which agents involuntarily leave bequests to their children. Abel also considers the effects of introducing a fully funded social security system into such an economy; his chief finding is that such a policy reduces savings.[1]

The assumption entertained by Eckstein et al. (1985) and Abel (1985) that completely selfish parents with no interest in their children leave involuntary bequests to their own children seems rather arbitrary. Clearly parents have the option to bequeath their wealth to surviving spouses, friends, other relatives, or charitable organizations. In addition, the notion that bequests are completely involuntary seems implausible. An alternative assumption is that selfish parents and selfish children collectively pool the risks of the parents' date of death in a manner that is mutually advantageous. There are three reasons why cooperative (voluntary) risk pooling seems a more realistic assumption. First, cooperative risk pooling Pareto dominates noncooperative behavior. Second, as described in Kotlikoff and Spivak (1981), the risks of uncertain longevity appear to be very large; the amount of resources that mildly risk-averse, selfish individuals would surrender to have access to fair annuity insurance is potentially quite sizable. This suggests a very substantial demand for market insurance if selfish parents cannot make comparable risk-pooling arrangements with their children, friends, or other relatives. Third, pooling longevity risk with even a single child can capture a large fraction of the gains from perfect insurance (Kotlikoff and Spivak 1981); hence, such risk pooling with children appears well "worth the trouble," with the gains far exceeding any reasonable transaction costs.

This paper models cooperative risk pooling of selfish parents and children taking into account the arrival of future selfish family members, namely, unborn grandchildren, great grandchildren, great great grandchildren, and so on. At each point in time the anticipated arrival of additional agents with whom young family members can share risks influences the set of current risk-sharing arrangements that are of mutual advantage to young and old family members. As a consequence the solution to the bargaining problem between currently living family members takes account of the infinite sequence of bargains struck by family descendants.

In addition to modeling the process of sequential generational risk sharing, we calculate, for the CES utility function, the stochastic steady state level and distribution of wealth. These calculations suggest that perfecting annuity insurance can have major impacts on national savings. For our preferred set of parameter values, the introduction of perfect annuity insurance reduces wealth by 35%–60% in the long run. The exact percentage reduction in savings within this range depends on assumptions about the cooperative bargaining solution. These figures are large, and larger still if one assumes a greater degree of risk aversion.

Given our parameterization of preferences, the 35%–60% range should, however, be viewed as an upper bound for the impact of introducing what amounts to a fully funded social security system in an economy with family risk sharing. There are two reasons why these figures are likely to considerably overestimate the actual outcome. First, they are partial equilibrium estimates, that is, they do not take account of potential changes in factor prices (wages and interest rates) that would arise, in a closed economy, from a major reduction in national wealth. Such price changes can significantly dampen savings reductions in models of this kind. Second, in order to highlight the impact of insurance provision, we assume that at most two family members are alive simultaneously. This generates the smallest possible risk sharing within families. Obviously, a sufficiently large number of family members is capable of pooling virtually all risks of uncertain longevity. With large enough families sharing mortality risks, the effect on aggregate wealth of perfecting insurance provision could be quite small.

While these numbers are partial equilibrium estimates and intentionally biased upward by our modeling of family size, they are surprisingly large relative to our prior beliefs. They suggest that the insurance aspects of social security are potentially as important in altering national savings as is the method of social security finance. It is also worth pointing out that the transition to the full annuity insurance equilibrium is completed once the initial generation of young family members reach old age. In real time, this is 40–50 years, but one would expect to see most of the ultimate change in savings occurring within the first 20 years. A final point that aids in evaluating these findings is that full insurance, while generating a Pareto-efficient steady state, may involve a steady state level of welfare that is lower than the minimum level of welfare in the family insurance stochastic steady state. This somewhat paradoxical result is explained as follows: the provision of full insurance transfers resources to the first cohort of elderly at the expense of initial young and future generations. While the new steady state is efficient, it has a smaller stock of resources, in this case capital, because of the initial transfer. This transfer to the initial elderly is not effected by explicit redistribution across age groups. It arises more subtly, namely, from the inability of young family members to continue selling insurance to their parents in exchange for their parents' potential bequests. Rather than bargain at less than fair insurance terms with children, provision of perfect annuities, which involves each cohort's pooling risk with its own members, permits the initial generation of elderly to consume at a higher rate. The initial set of children as well as all future generations are better off because of the perfection of the insurance market, but worse off because they no longer receive inheritances. Since all children in this paper are born with identical endowments, eliminating inheritances by providing perfect insurance also eliminates inequality.

The next section presents the infinite-horizon bargaining model; the zero bargaining, involuntary bequests model is also presented for purposes of comparison. This section also describes the algorithm used to solve the bargaining problem. Section 8.3 discusses the process of wealth transmission in the stochastic steady state. Section 8.4 compares long-run stocks of wealth under (1) perfect annuity markets, (2) three alternative parent-child bargaining solutions, and (3) no-insurance arrangements with involuntary transfers made to children. This section also considers how the presence of additional children would alter the findings. Section 8.5 summarizes the paper and discusses ideas for additional research.

8.2 The Model

As a prelude to presenting the selfish family, infinite-horizon bargaining problem, this section briefly reviews wealth accumulation under perfect annuity markets. In the subsequent modeling of family risk sharing, each selfish parent reaches a bargain with a single selfish child regarding the risk of long life. This is the simplest of family structures, but the associated intergenerational bargaining problem remains moderately complicated. The final part of this section describes how our stylized economy operates when family bequest-annuity agreements do not exist, but where involuntary bequests are made to children as in Eckstein et al. (1985) and Abel (1985). In this case it is everyone for himself; that is, there are no risk-pooling opportunities to ameliorate the risk of long life.

In comparing the economy under these three insurance arrangements—perfect insurance, self-insurance between parent and child, and no insurance—it is important to distinguish between transition effects and steady state comparisons. Clearly, if we move from no insurance to a family deal or from a family bargain to perfect insurance, the first generation gains. These gains are due to the fact that the generation alive during the switch received an inheritance from its parent but gives none or one of smaller expected value to its children. Kotlikoff and Spivak (1981) estimated that these gains to the first generation could be very substantial. For instance, consider a completely selfish 55-

year-old male who gains no pleasure from leaving bequests and whose time-separable consumption preferences are isoelastic, with a relative risk aversion coefficient of .75. This individual would consider the introduction of a perfect annuities market equivalent to an increase in his (her) wealth of 47%; with perfect annuities, there is no need to maintain precautionary balances to provide for an extraordinarily long life, and the individual can, therefore, enjoy a higher consumption stream for the remainder of his (her) life. The gains to those who first get access to a perfect annuities market increase with the age and degree of risk aversion of the individual. For uninsured individuals the gains to deals within the family are also large. With two participants the gain is roughly half that offered by perfect insurance, and with three it is roughly 70% as great. Hence, one would also expect significant start-up gains in moving from zero to family insurance.

This paper, in contrast to Kotlikoff and Spivak (1981), concentrates on steady state comparisons of the three insurance environments. In the case of family insurance we look at situations where a parent is insuring with a child, the child later makes a deal with his child, and so on. The analysis of aggregate wealth requires consideration of the entire family history of insurance arrangements and mortality experience. Obviously, the consumption and saving of current family members depends on their inherited wealth, which depends on the sequence of wealth and death dates of all previous ancestors.

There are 4 periods of life in this model. People live with certainty for the first 3 periods and survive to the fourth with probability P. So, the fraction $(1 - P)$ of the population live only 3 periods, while P live 4 periods. Children are 1 period when their parents are 3. Any negotiation or deal, explicit or implicit, between parent and child takes place before the parent and child engage in their respective third- and first-period consumption.

Individuals are exogenously endowed with earnings. The time pattern of the receipt of these earnings greatly influences saving and wealth in the economy. We assume that no earnings are received in the fourth period of life and examine a number of patterns of income receipt in the first 3 periods. Consumers are modeled as maximizing expected lifetime utility subject to one or more budget constraints. Utility is taken as separable in consumption (C_t) over time.

The perfect annuities case is by far the simplest to analyze since an individual's choice problem is separate from that of his parents and children. In this case each individual at age 1 maximizes

(1) $$EU = \sum_{t=1}^{4} P_t U(C_t) \alpha^{(t-1)}$$

subject to

$$\sum_{t=1}^{4} P_t C_t R^{(t-1)} = W_1,$$

where P_t is the probability of surviving to period t ($P_1 = P_2 = P_3 = 1$, and $0 < P_4 < 1$), C_t is consumption in period t, R is the discount factor (one divided by one plus the interest rate), α is the pure time discount factor, and W_1 is the present value of earnings. Throughout this paper we use the isoelastic form for $U(C_t)$,

$$U(C_t) = \frac{C_t^{1-\gamma}}{1-\gamma}, \tag{2}$$

where $1 - \gamma$ is the elasticity of utility with respect to consumption. The parameter γ measures the (constant) degree of relative risk aversion.

The solution to the consumer's problem in the case of perfect annuities takes the form

$$C_t = \frac{W_1 \, (R\alpha)^{(t-1)\gamma}}{\sum_{j=1}^{4} R^{(j-1)(1-\gamma)/\gamma} \alpha^{(j-1)/\gamma} \, P_j}. \tag{3}$$

Knowing C_t and the time pattern of earnings one can derive the accumulated wealth of each cohort. Total wealth in the economy equals the sum of each cohort's wealth holdings.

The family insurance solution where each member acts solely out of self-interest is much more complicated. When the bargaining takes place the parent is age 3 with one more period of certain life followed by one period of uncertain life. The agreement reached by parent and child can be thought of as the parent's buying an annuity from the child. In return for some money in period 3 (the price of the annuity) the child promises to offer a specified level of support for the parent in period 4 in the event that the parent lives that long. Equivalently, the deal can be arranged such that the child gives the parent some money before period 3 in return for being made beneficiary of the will of the parent. The equivalence can be seen in the following example which assumes a zero rate of interest for simplicity: say the parent pays \$1 for an annuity that gives him \$2 in period 4 of his life should he live. In the equivalent support-for-bequest arrangement the child gives the parent \$1 in period 3 in return for the parent's agreeing to save \$2 for this fourth period and makes the child his beneficiary should he die at the end of period 3. In both of these arrangements the child makes a net transfer of \$1 to the parent if the parent lives to old age and receives \$1 if the parent does not. Regardless of how the bargain is explicitly or implicitly specified, the parent and child share the risk of the parent's life span. Perhaps the simplest way to think about these deals is the first way, the purchase of annuity insurance by the parent from the child. The next issue to address is what is the price of this insurance.

Both the parent and the child can be made better off by striking a bargain. However, there is some indeterminacy as to how the surplus will be divided. One can imagine the price of the annuity being set sufficiently high that the parent's utility is just the same as if no deal had been struck, and, therefore, all of the gains from trade go to the child. At some low price, all of the gains from trade would go to the parent. An additional complication is that the child, in striking an arrangement with the parent, considers the third-period bargain he will make with his own child. The expected utility from that future bargain is denoted $\hat{V}$ and depends on the child's level of third-period wealth, W_{s3}, that is, $\hat{V} = \hat{V}(W_{s3})$. Throughout the paper we assume that successive children all earn identical amounts with certainty in the first three periods of their lives. Hence, the resources of the grandchild, with whom the child will bargain, is suppressed as an argument of $\hat{V}$.

The frontier of the utility possibilities space with intergenerational bargaining is located by solving the following problem:

Maximize

$$(4)\quad \frac{C_{f3}^{1-\gamma}}{1-\gamma} + \frac{\alpha P C_{f4}^{1-\gamma}}{1-\gamma} + \frac{\theta C_{s1}^{1-\gamma}}{1-\gamma} + \theta P\left[\frac{\alpha C_{s2,a}^{1-\gamma}}{1-\gamma} + \alpha^2 \hat{V}(W_{s3,a})\right] + \theta(1-P)\left[\frac{\alpha C_{s2,d}^{1-\gamma}}{1-\gamma} + \alpha^2 \hat{V}(W_{s3,d})\right]$$

subject to

$$C_{f3} + C_{s1} + R(C_{f4} + C_{s2,a}) + R^2 W_{s3,a} = W_{s1} + W_{f3}/R$$

and

$$C_{f3} + C_{s1} + RC_{s2,d} + R^2 W_{s3,d} = W_{s1} + W_{f3}/R,$$

where C_{f3} and C_{f4} are the parent's certain and contingent consumption in periods 3 and 4, respectively; C_{s1} is the child's first-period consumption, and $C_{s2,a}$ and $C_{s2,d}$ are the child's second-period consumption contingent upon the parent being alive or dead in period 4, respectively. The child's certain present value of resources is W_{s1}, and his (her) parent's third-period wealth is W_{f3}. Finally, $W_{s3,a}$ and $W_{s3,d}$ are the third-period levels of wealth of the child, that he or she uses in bargaining with the grandchild, contingent upon the parent being alive or dead in period 4.

Problem (4) involves maximizing a weighted sum of the two participants' expected utility where the weight θ, applied to the child's utility, potentially ranges from zero to infinity. The child considers his consumption in periods 3 and 4 under two eventualities: either his parent dies early, and he, therefore, does not have to pay off on the annuity insurance agreement (this is reflected in the final term of eq. [4] which is weighted by the $[1 - P]$ possibility of its occurrence), or the parent dies late and, hence, the child does have to pay off on the annuity insurance (the fourth term in eq. [4]). As stated, the $\hat{V}(W)$ function gives the expected utility the child experiences from his third- and fourth-period consumption discounted to period 3 of his life as a function of his wealth in period 3.

Equation (4) has two budget constraints because total consumption plus savings for the child's third period equals total initial wealth of the parent and child under both lifetime possibilities for the parent. The weight θ reflects the terms of trade in this bargaining problem. In general one would expect θ to be a function of the resources of both the parent and the child, W_{f3} and W_{s1}, respectively. However, since W_{s1} is constant in our analysis, we express $\theta = \theta(W_{f3})$.

Solving problem (4) for different values of θ traces out the utility possibility frontier for family deals shown in figure 8.1. Obviously, not all values of θ will generate outcomes that are in the core. We have labeled as θ_s the critical value for θ for which the parent receives none of the gains from trade (i.e., the child receives all gains from trade).

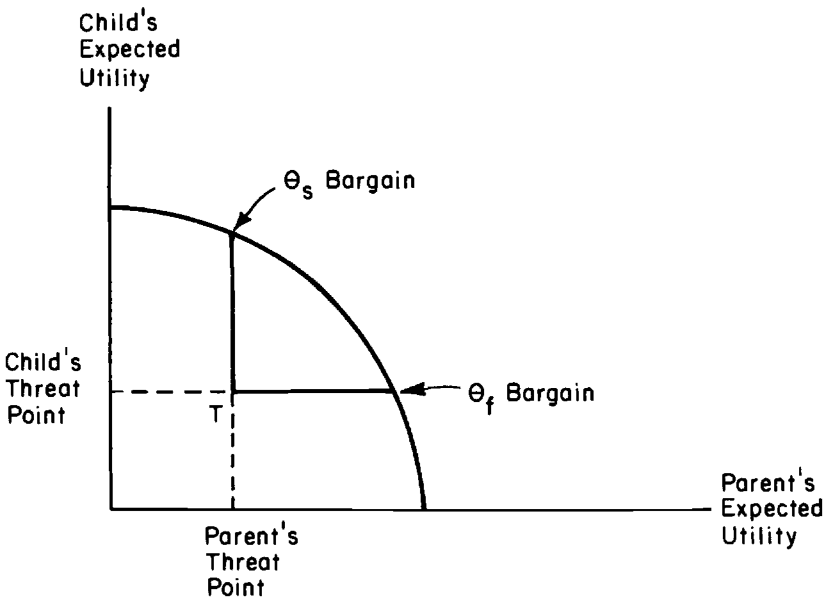

Fig. 8.1 Utility possibilities frontier

We define θ_f symmetrically with the parent getting all of the surplus. The point T is the threat point, indicating the parent's and child's expected utility levels if they fail to bargain with each other. As is clear from problem (4), figure 8.1 depends on the respective resources of the father and the son, W_{s1} and W_{f3}, and on the function $\hat{V}_s(\bullet)$.

Since we consider a stationary environment in which tastes and endowments of children remain unchanged, we will limit ourselves to stationary bargaining solutions. That is, we assume that the $\hat{V}$ function will be the same for the bargaining of each successive pair of generations. An implication of stationarity is that the parent's expected utility in (4) expressed as a function of his wealth, W_{f3}, equals the child's expected utility function, $\hat{V}$, when the child becomes a father. An immediate property of stationarity is that the child reaches the same deal with his child as his parent did with him if respective resources are the same. More formally, a stationary solution is defined as a bargaining function $\theta(W_{f3})$ and an expected utility function $V(W)$ such that if C^*_{f3}, C^*_{f4} are optimal values of consumption derived from solving problem (4), where $V(W_{s3})$ is substituted for $\hat{V}(W_{s3})$, then

$$V(W_{f3}) = \frac{1}{1-\gamma}C^{*1-\gamma}_{f3} + \alpha P\frac{1}{1-\gamma}C^{*1-\gamma}_{f4}.$$

Solving problem (4) involves searching for a fixed-point function V and an associated $\theta(W_{f3})$ function that produces outcomes that are in the core. We consider and compute three solutions to problem (4). In the first solution, denoted θ_s, the child receives all the gains from trade; furthermore, all successive bargains involve children receiving all gains from trade. In the second, θ_f solution, the initial and all successive fathers receive all gains from trade. In the third solution the gains from trade are always divided between child and son according to John Nash's (1953) two-person bargaining solution.

In the θ_s solution parents receive their threat-point level of expected utility. This is the expected utility received by the parent if he acts on his own and is given by the solution to (5). Maximize

$$\text{(5)} \qquad \frac{C^{1-\gamma}_{f3}}{1-\gamma} + \frac{\alpha P C^{1-\gamma}_{f4}}{1-\gamma}$$

subject to

$$C_{f3} + R\,C_{f4} = W_{f3}/R.$$

The structure of the problem is very much like that with perfect annuities, except that providing for C_{f4} costs R instead of only PR. The advantage of annuity markets is precisely this reduced cost of consumption in periods where survival is uncertain. Denote $V_s(W_{f3})$ as the maximum utility that the parent with wealth W_{f3} can achieve on his own by solving (5). Thus, $V_s(W_{f3})$ is the indirect utility function when no deal is struck and is given by

$$V_s(W_{f3}) = k\frac{W_{f3}^{1-\gamma}}{1-\gamma},$$

where

$$k = R^{\gamma-1}[(1 + (\alpha P)^{1/\gamma}R^{\frac{(\gamma-1)}{\gamma}}]^{\gamma}.$$

Naturally, $V_s(W_{f3})$ is the minimum the parent is willing to accept in an annuity bargain with his child. In addition, V_s is the expected utility function of the child in the θ_s bargain with his own child. Replacing V_s for $\hat{V}$ in (4) and choosing θ_s for each value of W_{f3} such that

$$V_s(W_{f3}) = \frac{C_{f3}^{*1-\gamma}}{1-\gamma} + \alpha\frac{PC_{f4}^{*1-\gamma}}{1-\gamma}$$

provides a proof by construction that V_s is a fixed-point function for the θ_s problem. In addition the computed values of θ_s for different values of W_{f3} determine the function $\theta_s(W_{f3})$. While parents, in this θ_s bargain, receive their threat-point levels of expected utility, their actual pattern of consumption differs from what they would choose on their own. As described below, C_{f3}^* is smaller and C_{f4}^* greater than the respective solution values to problem (5).

Although the V_s function was obtained analytically, this is not generally possible. For the θ_f and Nash (denoted θ_n) solutions an iterative technique described below is used to find fixed-point functions and their associated θ functions. Both the θ_f and θ_n solutions require specifying the child's threat point. Given our assumption of a cooperative, efficient solution to father-son bargaining, the child, if he fails to bargain with his father, can credibly assert to his father that he will be able to reach a deal with his child. The child's threat point, EU_s^T, is the solution to problem (6); it involves the child's consuming C_{s1} and C_{s2} in his first two periods, respectively, and bargaining with his child in period 3 based on third-period wealth, w_{s3}.

Maximize

$$(6) \qquad EU_s^T = \frac{C_{s1}^{1-\gamma}}{1-\gamma} + \frac{\alpha C_{s2}^{1-\gamma}}{1-\gamma} + \alpha^2\hat{V}(W_{s3})$$

subject to

$$C_{s1} + RC_{s2} + R^2W_{s3} = W_{s1}.$$

In the case of $\hat{\theta}_f$ bargaining, $\hat{V}$ is replaced by V_f in (6) as well as (4). The θ_f solution proceeds by first guessing a function V_f. Next we solve (6) to determine the son's threat-point utility EU_s^T. Given the guess of V_f and the derived value of EU_s^T, θ is chosen in (4) such that the son's expected utility in the solution to (4) equals EU_s^T. This last calculation is repeated for different values of W_{f3}, thereby generating a function $\theta_f(W_{f3})$. In addition to computing a θ_f function based on the initial guess of V_f, the solution to (4) based on $\theta_f(W_{f3})$ determines the father's expected utility in the bargain. The maximizing values of

$$\frac{C_{f3}^{*1-\gamma}}{1-\gamma} + P\alpha\frac{C_{f4}^{*1-\gamma}}{1-\gamma}$$

for different values of W_{f3} provide an expected utility function for the parent in his θ_f bargain with his child. This function is used as the next guess of the V_f function, and the calculations are repeated. The iteration proceeds until the guess of the V_f function equals the father's expected utility as a function of W_{f3}, that is, until we have found a function V_f, which is a fixed point of the mapping described.

In the Nash bargaining case a very similar solution technique is applied. The Nash solution involves choosing θ in (4) to maximize the quantity $(EU_f - EU_f^T)(EU_s - EU_s^T)$, where EU_f and EU_s are the expected utilities obtained by the parent and child, respectively, and EU_f^T equals V_s, the parent's threat point. To find V_n, the Nash fixed-point function, we again choose an initial guess of V_n and solve (6) to find EU_s^T. We also solve (5) to find EU_f^T. Next the guessed value of V_n is substituted for $\hat{V}$ in (4), and θ_n is chosen to maximize $(EU_f - EU_f^T)(EU_s - EU_s^T)$. Repeating this last step for alternative values of W_{f3} generates a function $\theta_n(W_{f3})$ as well as an expected utility function of the father arising from Nash bargaining. This latter function is used as the second guess of the V_n function. The iteration continues until we find a fixed-point function V_n. In this bargaining solution as in the previous θ_s solution, the $\theta_f(W_{f3})$ and $\theta_n(W_{f3})$ functions calculated in the last round of the iteration correspond to the correct bargaining functions for the functions V_f and V_n, respectively.

The V_s function is used as the initial guess of the V function for the θ_f and Nash bargaining solutions. In each iteration we computed the solution to (4) for 80 different values of W_{f3}. We then fit a fifth-order polynomial in W_{f3} to these points and used the resulting regression as the guess of V in the next iteration. The iterative procedure for determining V converged roughly by the eighth iteration; 12 iterations each were used for the θ_f and Nash cases. By "rough convergence" we mean that economic choice variables were identical to at least the second digit between iterations. For a range of intermediate values of

W_{f3} the calculated consumption terms are identical to five digits between iterations. While we believe more accurate values of the V_s and V_n functions could be obtained, the computation costs of achieving the additional accuracy is considerable; solving (4) for any one of the 80 values of W_{f3} in any one of the 12 iterations requires rather extensive computation.

8.2.1 The Involuntary Bequest Model

The next case we examine is the situation in which there are no insurance arrangements but unintentional bequests are made to children. This case has been examined in 2-period models by Eckstein et al. (1985) and Abel (1985). The solution differs from that of the threat points because the child inherits money unspent by the parent. The child in period 1 of his life can observe the wealth of his parent and can calculate the potential inheritance, I, he will receive should his parent die young. The child is assumed to solve the following problem.

Maximize

$$(7)\quad \frac{C_{s1}^{1-\gamma}}{1-\gamma} + P\left[\frac{\alpha C_{s2,a}^{1-\gamma}}{1-\gamma} + \alpha^2 V_s(W_{s3,a})\right] + (1-P)\left[\frac{\alpha C_{s2,d}^{1-\gamma}}{1-\gamma} + \alpha^2 V_s(W_{s3,d})\right]$$

subject to

$$C_{s1} + RC_{s2,a} + R^2 W_{s3,a} = W_{s1}$$

and

$$C_{s1} + R\,C_{s2,d} + R^2 W_{s3,d} = W_{s1} + I,$$

where

$$I = W_{f3}/R - C_{f3}.$$

The child maximizes his welfare subject to the certain earnings endowment, W_{s1}, and the inheritance I left by the parent if he dies young. The V_s function gives the level of expected utility the child can receive in periods 3 and 4 with no deal with his child, that is, the solution to problem (5) above.

8.3 The Transmission of Wealth in the Stochastic Steady State

Figure 8.2 graphs the wealth of children in their third period (when they are parents) against their parents' wealth, W_{f3}, for the case of family insurance bargains. The amount of wealth the child brings into

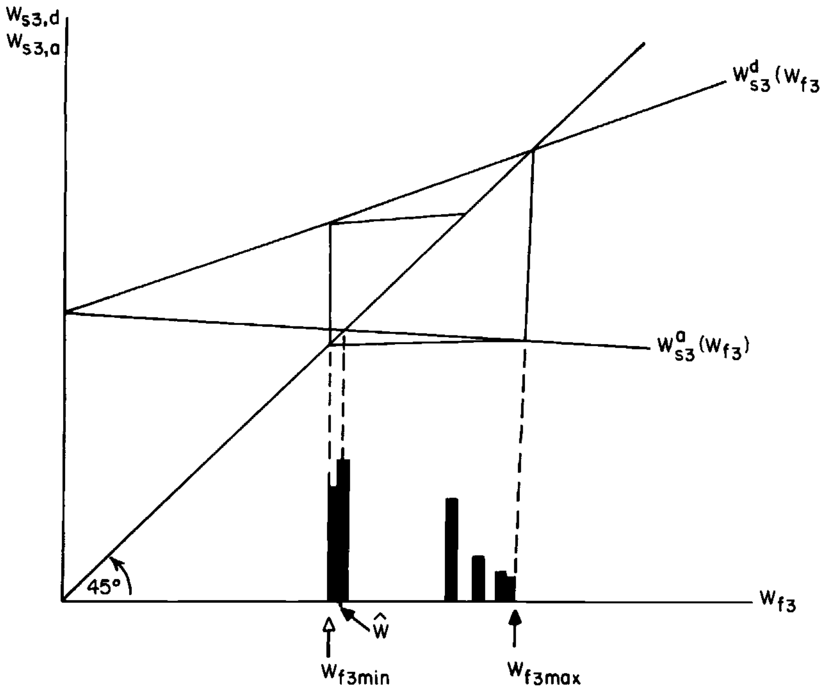

Fig. 8.2 Wealth transmission functions and the steady state distribution of parents' wealth

his third period depends, of course, on the age at which his parent dies. The curves $W_{s3}^d(W_{f3})$ and $W_{s3}^a(W_{f3})$ indicate the third-period wealth of the child if his own parent lives for 3 periods and 4 periods, respectively. Note that the two curves intersect on the vertical axis, since a child whose parent has no wealth engages in the same consumption regardless of the date of his parent's death.

The exact position and shapes of these curves depend on the specification of the utility function as well as the parent-child bargaining solution. For the examples we describe here, the curves were constructed by fitting fifth-order polynomials to the values of $W_{s3}^d(W_{f3})$ and $W_{s3}^a(W_{f3})$ calculated for 80 different values of W_{f3}. The intercepts in each regression were constrained to equal the amount of resources a child would save for period 3 assuming he engages in no bargain with his parent. In each calculation, the estimated curves were essentially straight lines, with $W_{s3}^d(W_{f3})$ and $W_{a3}^a(W_{f3})$ monotonically increasing and decreasing W_{f3}, respectively.

Intuitively, $W_{s3}^d(W_{f3})$ rises with W_{f3} because a fraction of the parent's increased resources will be allocated to the parent's contingent fourth-period consumption, C_{f4}. If the parent dies after period 3, the additional

C_{f4} is passed on to the child. For the child the inheritance is allocated to larger second-period consumption as well as larger third-period savings, $W^d_{s3}(W_{f3})$, that is, used in the bargain with his own child. The decline in $W^a_{s3}(W_{f3})$ as W_{f3} increases is explained as follows: regardless of the bargaining solution between the parent and child, the parent's contingent bequest rises with W_{f3}. Part of the price the child pays for the larger contingent bequest is somewhat lower values of second-period consumption and third-period wealth in the case the parent does not die young. This permits the parent to consume more in period 3 and, potentially, in period 4.

Assuming, as is verified in our actual calculation, that the slope of W^d_{s3} (W_{f3}) is everywhere positive and less than unity, W_{f3max} is the unique limiting value of a parent's third-period wealth when *all* his forefathers have died early. For values of W_{f3} above W_{f3max}, successive early deaths of parents lead to smaller values of W_{f3} for each successive parent until the sequence converges to W_{f3max}. Similarly, starting with a value for W_{f3} below W_{f3max} and assuming that all successive parents die early leads to successively larger values of W_{f3} until W_{f3max} is reached.

We next turn to the minimum bound on the stochastic steady-state distribution of a parent's wealth. If the slope of $W^a_{s3}(W_{f3})$ is between 0 and -1, which is the case in the examples presented below, then $\hat{W}_{f3}$ is the unique limit of the value of a parent's wealth as successive parents in a family continue to live through period 4. In this case the sequence of W_{f3}s, starting at any particular value, converges as a "Cobb-web" to $\hat{W}_{f3}$; that is, each successive parent with more wealth than $\hat{W}_{f3}$, who lives to period 4, has a child who has less than $\hat{W}_{f3}$ when the child becomes a parent.

In the stochastic steady state W_{f3min} is the lower bound on a parent's third-period wealth. Values below W_{f3min} cannot arise in the stochastic steady state; any parent with W_{f3} below W_{f3min} will have a child whose wealth as a parent is between W_{f3min} and W_{f3max}. Once the W_{f3} for a particular family falls within W_{f3min} and W_{f3max}, no parent in the family will ever appear with wealth outside this range. Values of W_{f3} below W_{f3min} and above W_{f3max} are nonrecurrent states in the Markov process that maps W_{f3} into $W^d_{s3}(W_{f3})$ with probability $1 - p$ and into $W^a_{s3}(W_{f3})$ with probability p. As can readily be seen by tracing out alternative p and $(1 - p)$ sequences, starting with values of W_f between W_{f3min} and W_{f3max}, the larger the value of W_f in the preceding generation, the smaller will be the W_{f3} in the next generation if the parent dies late. W_{f3min}, therefore, corresponds to the value of $W^a_{s3}(W_{f3})$ for W_{f3max}, that is, $W_{f3min} = W^a_{s3}(W_{f3max})$. Hence, if the richest parent survives to period 4, his child is the poorest parent when he reaches period 3. This extreme "riches to rags" result is quite intuitive. A parent with the largest possible wealth, W_{f3max}, provides the largest estate if he dies early but no estate if he dies late. In order to "purchase" the right to this largest potential estate the child pays the largest price in terms of reduced consumption and third-period wealth if his parent lives.

Since the Markov process described in figure 8.2 is nonrecurrent, there are large regions between W_{f3min} and W_{f3max} that have zero mass with respect to the steady state distribution of wealth. The shaded areas in figure 8.2 chart this distribution for the case θ's in which parents receive none of the gains from bargaining with their children. This distribution was constructed by giving 100,000 families the same initial value of W_{f3} and then simulating 25 successive generations using a 0.6 probability of a 4-period lifetime. The distribution of W_{f3} stabilized after roughly eight generations. Since we assume that a new generation is born every period, rather than every other period, there are also orphaned 2-period-old children as well as 2-period-old children with surviving parents who hold wealth at any point in time. Calculating the stochastic steady state's stock of wealth requires simply summing the wealth holdings of all age 3 parents, the wealth of orphaned children, and the wealth of 2-period-old children and their surviving 4-period-old parents. The wealth holdings of these latter two groups are derived from the distribution of wealth holdings of 3-period-old parents; the consumption of each of the 100,000 parents and their children, when these parents are age 3, is subtracted from the income of these families to compute their combined saving. This saving plus each parent's initial wealth represents the next-period wealth holdings of families consisting either of orphaned children or of children with surviving parents. Since this wealth distribution is stationary in the stochastic steady state, next period's wealth holdings of these groups is identical to this period's wealth holdings of such groups. Similar calculations are made for the case in which there are no insurance bargains between parents and children, but children nonetheless inherit their parents' estates.

Parameter values were chosen as follows: the time preference factor, α, and the discount factor, R, both equal .86. The coefficient of risk aversion, γ, equals 4, and the fourth-period survival probability, ρ, equals 0.6. If one thinks of each period as consisting of 15 years, then a discount factor of .86 corresponds to a 1% annual real rate of return. In addition, if we view parents as being age 50 and children age 20 when the bargains are struck, the 0.6 fourth-period survival probability is roughly equivalent to assuming an expected age of death of 74.

Table 8.1 presents the calculated values for a parent's third- and fourth-period consumption at alternative levels of W_{f3} under perfect insurance, the three alternative parent-child bargains (the θ_f, Nash, and θ_s solutions to [6]), and the case of no-insurance arrangements. In each of these cases, the parent's consumption increases with his third-period

Table 8.1 **Parent's Consumption Under Alternative Insurance Arrangements**

Parent's Third-Period Wealth ($WF3$)	Perfect Insurance		θf Bargain (Parents Receive All Gains from Trade)		Nash Bargaining Solution		θs Bargain (Children Receive All Gains from Trade		No Insurance Arrangements	
	$CF3$	$CF4$	$CF3$	$CF4$	$CF3$	$CF4$	$CF3$	$CF4$	$CF3$	$CF4$
9.0	6.9	6.9	6.4	5.8	6.2	5.6	5.9	5.3	5.0	5.3
8.0	6.7	6.7	5.7	5.2	5.5	4.9	5.2	4.7	5.3	4.6
7.0	5.4	5.4	5.0	4.5	4.89	4.3	4.5	4.1	4.6	4.0
6.0	4.6	4.6	4.4	4.0	4.2	3.8	3.9	3.6	4.0	3.5
5.0	3.8	3.8	3.6	3.3	3.5	3.2	3.2	3.0	3.3	2.9
4.4	3.4	3.4	3.2	3.0	3.0	2.8	2.8	2.6	2.9	2.5

NOTE: Table assumes $\gamma = 4$, $P = .6$, $\alpha = R = .86$.

wealth. Access to perfect insurance results, for this parameterization of utility, in higher levels of consumption for the parent in both periods 3 and 4 relative to the other cases of partial or zero insurance. For example, if the parent's wealth is 4.4 at the beginning of period 3, he consumes 2.9 and 2.5 in periods 3 and 4, respectively, with no insurance, and 3.4 in both periods with perfect insurance. The present value difference in these consumption paths is 25%.

A parent-and-child bargain in which successive parents receive all gains from trade with successive children, the θ_f bargain, provides parents with consumption values that are roughly midway between those of perfect and zero insurance. Consumption values for the parent under the Nash bargaining solution lie between the θ_f and θ_s deals. This is the expected result since the Nash solution divides the gains from trade between parents and children. The θ_s bargain, in which the parent receives no benefits from dealing with his child, involves slightly less third-period consumption and slightly more fourth-period consumption when old than in the case of zero insurance.

Table 8.2 shows consumption and third-period wealth values of children in different insurance regimes. Under perfect insurance the child's consumption is 3.4 in each period; with no insurance arrangements and no involuntary bequests the child consumes 3.2 during the first 3 periods and 2.8 in the last period. Depending on the parent's wealth and longevity and the bargain struck between the two, the child can potentially consume well in excess of the perfect insurance values. As an example, take the case of a parent with wealth of 6.0 who agrees to a θ_s bargain with his child. The child's first-period consumption is 3.4, the same as under perfect insurance. If the parent dies after his third period, the child consumes 4.5 in period 2 rather than 3.4, the perfect insurance amount. Furthermore, the child's third-period wealth in this case is 7.9, substantially in excess of 4.4, the third-period wealth of a son under perfect insurance. With third-period wealth of 7.9, the child's third- and contingent fourth-period consumption values are, from table 8.1, roughly 5.5 and 4.9. For this child the total potential realized present value of consumption is 14.4, although the present value of his earnings is only 10.

8.4 The Savings Impact of Alternative Insurance Arrangements

Table 8.3 compares steady state per capita wealth stocks in the different insurance regimes under alternative assumptions about age-earnings profiles. Each of the age-earnings profiles has a present value of 10, which is received with certainty over the course of the first 3 periods. Since the child's resources are identical in each of these cases, the consumption decisions of the child and parent are the same for

Table 8.2 **Child's Consumption Under Alternative Insurance Arrangements Consumption and Wealth Values: Father-Son Bargains**

Father's Wealth	θ_s Bargain (Children Get All Gains from Trade)					Nash Bargaining Solution					θ_f Bargain (Children Get No Gains from Trade)				
	*CS*1	*CS*2*A*	*CS*2*D*	*WS*3,*A*	*WS*3,*D*	*CS*1	*CS*2*A*	*CS*2*D*	*WS*3,*A*	*WS*3,*D*	*CS*1	*CS*2*A*	*CS*2*D*	*WS*3,*A*	*WS*3,*D*
9.0	3.4	3.1	5.2	5.4	9.1	3.4	3.1	5.5	4.6	8.3	3.3	2.9	5.1	4.4	8.5
8.0	3.4	3.1	4.9	5.4	8.7	3.5	3.1	5.2	4.6	7.9	3.3	3.0	4.9	4.4	8.1
7.0	3.4	3.1	4.7	5.4	8.3	3.5	3.1	5.0	4.6	7.5	3.3	3.0	4.7	4.5	7.7
6.0	3.4	3.1	4.5	5.4	7.9	3.5	3.1	5.0	4.6	7.1	3.3	3.0	4.5	4.5	7.3
5.0	3.3	3.1	4.3	5.4	7.5	3.4	3.1	4.5	4.7	6.8	3.2	3.0	4.3	4.5	6.9
4.4	3.3	3.1	4.1	5.4	7.3	3.4	3.2	4.4	4.7	6.6	3.3	3.0	4.2	4.6	6.7

Consumption Values

Perfect Insurance				No Insurance Arrangements and No Involuntary Bequests			
*CS*1	*CS*2	*CF*3	*CF*4	*CS*1	*CS*2	*CF*3	*CF*4
3.4	3.4	3.4	3.4	3.2	3.2	3.2	2.8

NOTE: Table assumes $\gamma = 4$, $P = .6$, $\alpha = R = .86$.

Table 8.3 **Wealth Per Capita and Percentage Long-Run Decline in Wealth from Switch to Perfect Insurance**

Age-Earnings Profile	Perfect Insurance	θ_f Bargain (Parents Receive All Gains from Trade)		Nash Bargaining Solution		θ_s Bargain (Children Receive All Gains from Trade)		No Bargain Involuntary Bequests	
	Wealth	Wealth	Percentage Wealth Decline	Wealth	Percentage Wealth Decline	Wealth	Percentage Wealth Decline	Wealth	Percentage Wealth Decline
(10,0,0,0)	12.7	13.9	8.6	14.2	10.6	15.9	20.1	15.5	18.1
(5.0,5.8,0,0)	7.7	8.9	13.5	9.2	16.3	10.9	29.4	10.5	26.7
(3.3,3.9,4.5,0)	2.2	3.4	35.3	3.7	40.5	5.4	59.3	5.0	56.0
(3.0,5.8,2.7,0)	3.4	4.6	26.1	4.9	30.6	6.6	48.5	6.2	45.2
(2.0,5.8,3.5,0)	0.7	2.9	171.4	3.2	78.1	3.9	82.0	3.5	80.0

NOTE: Table assumes $\gamma = 4$, $P = .6$, $\alpha = R = .86$.

each of these earnings paths. Hence, the difference in stocks of wealth by row in table 8.3 are simply a function of the timing of the receipt of labor income.

The absolute size of these economies' wealth stocks may appear small in comparison to the level of earnings or income in a particular period. However, such stock-flow ratios must be adjusted for the fact that flows in this model are received over a period that corresponds to roughly 15 years. In the case of the third and probably the most realistic earnings profile in table 8.3, the ratio of wealth to one-fifteenth of a period's labor earnings is 6.9 in the case of the θ_s bargain. A wealth/earnings ratio of 6.9 is somewhat greater than that observed in the United States.

The percentage reductions in wealth from moving to perfect insurance reported in table 8.3 are very large. For the earnings profile in the third row the long-run wealth reduction is 59% starting from the θ_s (children take all) stochastic steady state. It is 41% in the case of an initial Nash bargaining equilibrium and 35% when the initial equilibrium involves θ_f (parents take all) bargain.

The values in table 8.3 are highly sensitive to the shape of the age-earnings profile. The smallest percentage wealth reduction arises when all earnings are received in the first period; in this case wealth falls by 20.1% starting from the θ_s bargain and by 13.9% starting from the θ_f bargain.

The percentage change in wealth appears relatively insensitive to variations in the degree of relative risk aversion, γ. For example, reducing γ from 4 to 1.5 lowers the percentage decline in wealth under row 3's earnings profile and initial θ_s bargaining from 59.3% to 50.7%. Raising γ to 8 increases the value to 63.2%. Under table 8.3's first age-earnings profile the percentage wealth reductions starting from θ_s economies are 15.1, 20.1, and 22.9 for values of γ equal to 1.5, 4, and 8, respectively.

There is considerably more sensitivity to changes in the fourth-period survival probability P; however, the sensitivity depends on the choice of earnings profile. For example, lower P from 0.6 to 0.3, which reduces the expected age of death from roughly 74 to roughly 69, converts the 59.3% θ_s reduction (row 3, table 8.3) to 83.6%. The same reduction in P raises table 8.3's row 1, θ_s value from 20.1% to only 23.4%.

The large differences in wealth stocks between the perfect insurance and family insurance regimes suggests that steady state welfare could actually be lower in the case of perfect insurance. This is indeed possible. Under θ_f (children take all) bargaining and assuming γ equals 1.5, the expected utility of even the child of the poorest parent exceeds the uniform, steady state expected utility under perfect insurance. Starting from a situation of zero insurance, achieving the perfect insurance expected utility level requires a 7% increase in resources; achieving the expected utility of the child with the poorest parent in the θ_s stochasic steady state requires an 8% increase in lifetime resources, starting from this benchmark regime. Attaining the level of welfare of the child whose parent in the θ_f steady state has the maximum potential wealth, W_{f3max}, requires a corresponding 12% increase in resources.

The steady state stocks of wealth in the case of no family arrangements, but involuntary bequests to children, are slightly smaller than those under θ_s bargaining. This is not surprising since in both cases parents receive their threat-point levels of utility and consume roughly similar amounts. In the θ_s deal, however, the child's insurance provision leads to a somewhat lower level of the parent's consumption in period 3 and a somewhat higher level in period 4 (see table 8.1). In addition, given W_{f3}, the child consumes slightly less in period 1 in the θ_s deal than in the involuntary bequest setting. This consumption pattern explains the larger wealth stock in the θ_s insurance regime.

Another question raised by table 8.3 is the extent to which imperfections in annuity markets can fully explain observed intergenerational transfers. Kotlikoff and Summers (1981) invoked the assumption of perfect insurance arrangements in estimating that roughly 80% of private U.S. wealth corresponds to accumulated inheritances of those currently alive. This assumption that annuity insurance is fairly well developed in the United States can be defended by pointing to social security and other government annuities, private pensions, old age labor earnings that are partly contingent on survival, and the potential for family risk sharing involving multiple members. Still, it is interesting to ask how their calculation turns out when it is applied to the two-member family insurance economy described above. Their technique involves subtracting accumulated consumption from accumulated earnings for each cohort and then summing across cohorts to get a total wealth stock. This "life-cycle" wealth is then compared with actual wealth holdings. If agents in the economy are selfish and annuity arrangements are perfect or very close to perfect, computed and actual aggregate wealth will be identical or extremely close to one another.

The two-person family regime is, however, quite far from that of perfect insurance. As described here, this imperfection produces a stochastic steady state in which observed consumption profiles often exceed what could be financed from one's own labor earnings even under perfect insurance. Hence, in this economy, subtracting, for all cohorts, accumulated consumption, part of which is financed by past intergenerational transfers, from accumulated earnings produces an underestimate of the economy's actual wealth. For the θ_s bargain, with γ equals 4 and with table 8.3's row 3 earnings profile, 1.5, the underestimate is close to 90% of actual wealth. Since Kotlikoff and Summers's (1981) calculation understates United States wealth by 80%, imperfections in annuity markets appear potentially capable of fully explaining actual intergenerational transfers in the United States.

8.5 Conclusion

The preceding calculations as well as the figures presented in table 8.3 must be viewed cautiously. They embed rather extreme assumptions concerning the size of the risk-sharing pool and, in the θ_s case, the nature of risk sharing. A more realistic model would contain two parents pooling risk with two or more children. Since the parents, by themselves, can provide each other with considerable insurance protection, their threat-point values of expected utility are greater in collective bargaining with their children. As a consequence one would expect parents, in such a model, to have an expected utility level considerably greater than that described by the θ_s solution. In addition, if they can extract most of the gains from trade from dealing with their children, they will end up with close to perfect insurance. In that case the impact of improving annuity arrangements on savings would be minor.

Extending the analysis to different configurations of families is an area for future research. To date we have considered the simplest case of multiple children dealing with a single parent under the θ_s bargain. For table 8.3's third earnings profile the percentage reduction in wealth is quite similar to the 50% figure in table 8.3 over a range of children numbering as great as 5 per parent. Since their earnings profile implies very little saving in period 1, the change in the earnings' age structure from a 1/5 ratio of children to parents has little impact on accumulated earnings of particular cohorts at a point in time. In addition, the consumption patterns of children and the parent are not greatly altered in moving from one to five children under the θ_s bargain. This would not, of course, be the case in the θ_f bargain. A θ_f bargain with five children would provide parents with close to the consumption levels available with perfect insurance.

While the findings should be viewed cautiously, they do suggest that the manner in which annuity markets function can significantly affect saving, wealth, and welfare in an economy. That each generation has large incentives to improve annuity arrangements was demonstrated in Kotlikoff and Spivak (1981). Here we find that the steady state welfare gains are significantly smaller and, in fact, may be negative. The first generations' gain results in a smaller inheritance and capital stock for future generations. This lower wealth may more than offset the welfare gains that each generation receives from the availability of long-life insurance.

We should reemphasize that we are addressing a different question from that of Feldstein (1974) and others who are largely concerned with the funding status of social security. While that line of research attempts to estimate the substitutability of social security wealth for private capital, we are here concerned with the insurance aspects of pensions and social security. It is our feeling, buttressed by the results of this paper, that a considerable amount of saving is potentially done for what could be loosely termed precautionary motives. In addition, the exact manner in which families self-insure can have major consequences for wealth accumulation. When more perfect insurance policies are made available, whether funded or not, less aggregate saving occurs. While we have focused on annuity insurance, the paper's findings suggest that the availability of unemployment insurance, disability insurance, and health insurance could also significantly affect national saving. In addition, the government's pooling of human capital risks through progressive income taxation may also be having a major impact. In general, the study of savings and government insurance provision is an important area for additional research.

Note

1. This paper reaches a similar conclusion about the savings impact of perfecting insurance arrangements, although we model the initial, no-market/government annuity economy quite differently. Abel's research and ours were conducted independently.

References

Abel, Andrew B. 1985. Precautionary saving and accidental bequests. *American Economic Review* 75:777–91.

Barro, Robert J. 1974. Are government bonds net wealth? *Journal of Political Economy* 82:1095–118.

Davies, J. B. 1981. Uncertain lifetime consumption and dissaving in retirement. *Journal of Political Economy* 89:561–78.

Eckstein, Zvi; Eichenbaum, Martin S.; and Peled, Dan. 1985. Uncertain lifetimes and welfare enhancing properties of annuity markets and social security. *Journal of Public Economics* 26:303–26.

Feldstein, Martin S. 1974. Social security, induced retirement, and aggregate capital accumulation. *Journal of Political Economy* 82:905–26.

Hubbard, Glenn R. 1983. Uncertain lifetimes, pensions, and individual saving. In this volume.

Kotlikoff, Laurence J., and Spivak, Avia. 1981. The family as an incomplete annuities market. *Journal of Political Economy* 89:372–91.
Kotlikoff, Laurence J., and Summers, Lawrence H. 1981. The role of intergenerational transfers in aggregate capital accumulation. *Journal of Political Economy* 90:706–32.
Nash, John F. 1953. Two-person cooperative games. *Econometrica* 21:128–40.
Sheshinski, E., and Weiss, Y. 1981. Uncertainty and optimal social security systems. *Quarterly Journal of Economics* 96:189–206.
Yaari, M. E. 1965. Uncertain lifetime, life insurance, and the theory of the consumer. *Review of Economic Studies* 32:137–50.

Comment Michael Rothschild

Models like the one developed in this paper shed light on two issues. They can be used to assess how well different institutions work to share the risks of uncertain length of life. Because in the absence of perfect insurance people accumulate private wealth to insure themselves against poverty in their old age, these models are also used to analyze the effects of these different institutions on capital formation. My comments concern the first issue; I will discuss how I think the welfare consequences of different methods of intergenerational risk sharing ought to be measured. I will also indicate briefly how in one variant of the model analyzed in this paper, taxation can increase welfare.

One of the several virtues of this paper is that it explicitly calculates the distribution of wealth which results from the inheritance process. If there were perfect annuity markets, no one in this economy would leave an estate. Since everyone has the same ability to earn income, all people face the same lifetime budget constraint; all have, at birth, the same expected utility. When there are imperfect annuity markets, people leave estates. How much they leave depends on how long they live and how much they inherited from their parents. The bequest process thus induces a distribution of wealth. The characteristics of this distribution depend on the institutional structure; that is, it depends on the particular contract which fathers and sons make with one another. Because there is a distribution of wealth in societies with imperfect annuity markets, different individuals have different expected utility at birth. Expected utility is determined by one's father's wealth, and this varies from person to person.

This suggests using the following standard to compare welfare under different annuity arrangements: Suppose you were going to be born

Michael Rothschild is professor of economics, University of California, San Diego, and research associate, National Bureau of Economic Research.

into a society with a particular annuity structure. What is your expected utility if you assume a position in the wealth distribution according to the steady state distribution of wealth? This is, of course, the same thing as total utility in steady state using a utilitarian measure of welfare. This seems an appropriate criterion for this model as everyone is the same except for accidents of birth and length of life. Since they calculated the steady state distribution of wealth, Kotlikoff et al. could easily calculate this measure of welfare. It would be interesting to use this standard to compare the different imperfect annuity agreements discussed in the paper.

From this perspective there are three kinds of uncertainty in the model. The first is uncertainty about how long one will live. Call this length-of-life uncertainty; it is the primary source of uncertainty. Other kinds of uncertainty arise because people cannot completely insure themselves against a long life. The second kind of uncertainty is uncertainty about how wealthy one's father will be. ("Will be" because we are considering the thought experiment of being born into a random family.) Call this wealth uncertainty. Finally there is uncertainty as to how long one's father will live and thus what bequest one will actually get. Call this bequest uncertainty. Bequest uncertainty causes real misallocation of resources. Because I do not know what my wealth will be until bequest uncertainty is resolved, I cannot hope to allocate consumption over my lifetime as well as I could if I knew what my lifetime budget constraint would be before I started consuming.

With perfect annuity markets none of these kinds of uncertainty exist—or at least they can be perfectly insured against. If there are no annuity markets, then all three kinds of uncertainty exist. Intergenerational risk sharing mitigates length-of-life uncertainty; it does this at the expense of increasing bequest uncertainty. The size of bequest risk is determined by the difference between the amount the son gets if his father lives 3 periods and the amount (possibly negative) the son gets if his father lives 4 periods. This difference is larger if the son partially insures his father than if he doesn't. The larger, in the sense of second-degree stochastic dominance, are bequests, the greater is bequest uncertainty. I suspect that welfare is larger the smaller is bequest uncertainty. Bequests are smallest when fathers appropriate most of the gains from the annuity bargain. Thus, I think it likely that welfare or expected utility in steady state is highest when $\theta = \theta_f$. What makes this a hunch rather than a conjecture is my inability to speculate about the effect of different values of θ on wealth uncertainty.

One institution which would increase welfare from the no-insurance situation is a 100% inheritance tax, with proceeds distributed in a lump sum fashion. Such a tax would have almost the opposite effect of the imperfect annuities studied in this paper. It would do nothing to mitigate life uncertainty, but it would do away completely with both wealth uncertainty and bequest uncertainty. Such an inheritance tax would substitute for a market on which one could insure perfectly against bequest and wealth uncertainty. It would be interesting to compare expected welfare in a society which had no annuities but which did have inheritance taxes with expected welfare in a society with the imperfect annuities created by intergenerational risk sharing. An inheritance tax would make the imperfect annuity arrangements studied in this paper impossible. It would also entail more capital in steady state than would any intergenerational risk-sharing agreement.