Volume Title: Topics in the Economics of Aging

Volume Author/Editor: David A. Wise, editor

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-90298-6

Volume URL: http://www.nber.org/books/wise92-1

Conference Date: April 5-7, 1990

Publication Date: January 1992

Chapter Title: Introduction to "Topics in the Economics of Aging"

Chapter Author: David A. Wise

Chapter URL: http://www.nber.org/chapters/c7096

Chapter pages in book: (p. 1 - 19)

Introduction

David A. Wise

This volume contains papers presented at the National Bureau of Economic Research conference on the economics of aging in Carefree, Arizona, in April 1990. It is the third in a series of conference volumes associated with the NBER's ongoing project on the economics of aging. The first volume was *The Economics of Aging* (1989); the second was *Issues in the Economics of Aging* (1990). (Both were published by the University of Chicago Press.) The goal of the Economics of Aging Project is to further our understanding of the consequences for older people and for the population at large of an aging population. This volume pursued analysis of several of the issues addressed in the previous volumes and considers several new topics as well. This introduction summarizes the motivation for the papers and their principal findings. Much of the text is abstracted from the papers themselves.

Retirement

A great deal of effort has been devoted in recent years to the development of models of retirement that emphasize the need to compare present with future circumstances in deciding whether to retire. Some versions of these models are numerically complex, while others are much simpler. In "Three Models of Retirement: Computational Complexity versus Predictive Validity," Robin L. Lumsdaine, James H. Stock, and David A. Wise consider whether more complex models predict the actual retirement behavior of individuals better than less complex models. They conclude that retirement from the firm that they consider is poorly predicted by the simplest model but that the most complex model does no better than one that is numerically much simpler to implement. They also show that estimates of the effects on retirement of changes in Social Security provisions are very inaccurate if firm pension plan

provisions are not also accounted for—as has been the case in previous analyses of Social Security.

Empirical analysis often raises questions of approximation to underlying individual behavior. Closer approximation may require more complex statistical specifications. On the other hand, more complex specifications may presume computational facility that is beyond the grasp of most real people and therefore less consistent with the actual rules that govern their behavior, even though economic theory may lead analysts to increasingly complex specifications. Thus, the issue is not only whether more complex models are worth the effort but also whether they are better. The answer must necessarily depend on the behavior that the analysis is intended to predict. In their paper, Lumsdaine, Stock, and Wise consider the relation between computational complexity and the predictive validity of three models of retirement behavior.

Retirement has been the subject of a large number of studies over the past decade. Most have emphasized the effect of Social Security provisions on retirement age, but a wide range of methods has been employed. The earlier studies in this time period were based on regression or multinomial logit analysis. Subsequent analysis relied on nonlinear budget constraint formulations of the retirement decision and on proportional hazard model formulations. More recently, several authors have developed models that focus on forward-looking comparison of the advantages of retirement at alternative ages in the future and on the updating of information as persons age. Although the spirit of these latter models is basically the same, they vary widely in computational complexity. The potential advantages in predictive validity of the computationally more complex versions of these models are the primary motivation for this study, although to broaden the scope of the comparison the authors consider a much simpler model as well.

They compare the predictive validity of three models of retirement. The first is a simple probit model. The second is the "option value" model developed in earlier papers by Stock and Wise. The third is a stochastic dynamic programming model.

The analysis is guided by several key ideas. First, all the models are theoretical abstractions; none of them can be reasonably thought of as "true." The important consideration is which decision rule is the best approximation to the calculations that govern actual individual behavior. In this paper, judgments on which rule is best are based on empirical evidence on the relation between model specification and predictive validity.

Second, the models vary substantially in the computational complexity of the decision rules that they attribute to individual decision makers. The option value and the dynamic programming rules are both intended to capture the same underlying idea, but implementation of dynamic programming rules typically implies considerably more computational complexity than implementation of the option value rule. The option value model makes a simplify-

ing assumption that substantially reduces complexity. The probit model is much simpler than either of these.

Third, although the mathematically correct implementation of some decision rules requires dynamic programming, there is no single dynamic programming rule. The implied computational complexity depends in important ways on specific assumptions, in particular the disturbance term correlation structure. It is easier to incorporate more flexible correlation assumptions in the option value than in the dynamic programming formulations. Thus, for example, the option value specification may be a suboptimal solution to a dynamic programming rule that implies computational complexity difficult to implement even with a computer.

A question of practical importance is therefore whether different decision rules yield significantly different results.

The comparisons in this paper are made by estimating the models on the same data. The data, which pertain to the retirement decisions in a large *Fortune* 500 firm, have two important advantages. First, the retirement decisions can be related to the provisions of the firm's pension plan, and it is therefore possible to simulate the effect of changes in the pension plan provisions. Second, the firm offered an unanticipated "window" plan in one of the years covered by the data.

The principal measure of the predictive validity of the models is how well they predict the effects of the window plan. Like the typical defined benefit pension plan, this firm's plan provides substantial incentives to retire early. In addition, the window plan provided further incentive to retire early. Window plans, which have been offered by many firms in recent years, provide special bonuses to workers in a specific group—often defined by age, occupational group, or even a division within the firm—if the worker retires within a specified period of time, typically a year or less. The window plan allows a unique external test of the predictive validity of the models; it is possible to compare model predictions with actual retirement rates under the window plan.

The authors begin by obtaining model parameter estimates based on retirement decisions in a year (1980) prior to the window plan. They then use these estimates to predict retirement in a later year (1982) under the window plan. The estimates and predictions are based on male nonmanagerial employees.

Lumsdaine, Stock, and Wise find that the option value and dynamic programming models are considerably more successful than the less complex probit model in approximating the rules that individuals use to make retirement decisions but that the more complex dynamic programming rule approximates behavior no better than the simpler option value rule. The authors caution that definitive conclusions will have to await accumulated evidence based on additional comparisons using different data sets and with respect to different pension plan provisions.

The authors also use their analysis to show that predictions of the retirement

effects of changes in Social Security provisions are likely to be grossly inaccurate if the analysis does not consider both Social Security and firm pension plan provisions. Firm plan provisions have been unavailable to researchers analyzing the effects of changes in the Social Security system.

Saving for Retirement

In "Stocks, Bonds, and Pension Wealth," Thomas E. MaCurdy and John B. Shoven consider past returns to stocks versus bonds. They present strong evidence that systematic contributions proportional to earnings over a career have always led to more wealth at the time of retirement if the investment is in stocks rather than bonds.

For many people, the present value of their future pension annuity is their largest financial asset. The retirement income may come from a variety of pension accumulations, including defined contribution plans, defined benefit plans, individual retirement accounts, Keogh plans, and tax-deferred annuity plans. With many of these accumulation vehicles, the individual participant bears the responsibility of determining the assets in which the funds are invested and bears any uncertainty about the rate of return that will be realized on those assets. In choosing between stocks and bonds for their pension accumulation vehicle, most people probably know that bonds have a lower average return and a lower variance in return; bonds offer additional "safety" at the expense of a lower expected outcome. While this risk-return trade-off is both correct and well understood for short-term investment horizons, the extent to which it applies for long holding periods is not clear. For many workers, the time between the current contribution to the retirement account and the purchase of an annuity is thirty years or more. What is the relative risk and return on stocks versus bonds for such a long horizon? The pension participant typically not only has a long horizon but also makes many contributions throughout his or her career. For example, faculty at Stanford University make payments to their retirement accounts twice each month over their term of employment. How does such a pattern of purchase affect the relative desirability of stocks versus bonds as pension accumulation assets? Finally, most individual retirement accounts, Keogh plans, and defined contribution plans allow the participant not only to choose which assets are purchased with new contributions but also to move existing accumulations between asset categories. This raises the question of the desirability of gradually moving stock accumulations into bonds late in one's career. Such an option offers the potential advantage that one's retirement annuity would depend on the value of the stock portfolio at several selling dates rather than just its value on the date of purchase of the annuity.

MaCurdy and Shoven examine how some naive investment strategies for pension accumulations would have performed for employment careers of varying length between 1926 and 1989. Given a strategy, they calculate the

implied value for the pension account at the time of retirement for all possible completed careers of a specified horizon within the sixty-four-year period. They consider only strategies in which investors allocate their pension contributions either entirely into stocks (with all dividends and other returns reinvested in stocks) or entirely into bonds (with interest reinvested in bonds). These strategies are not optimal in any sense since they ignore any market timing issues as well as standard portfolio theory. They then consider some strategies for converting from stocks to bonds as a worker approaches retirement, but they do not attempt to determine the optimal portfolio composition as a function of years until retirement. Despite these limitations, they find that an "all stocks" strategy dominates all other investment policies considered for all career lengths of twenty-five years or longer. By "domination," they mean that an all stocks allocation would have generated a larger pension accumulation for every career that ended in retirement over the period 1926–89.

MaCurdy and Shoven's findings have important implications for pension investment policies, and they suggest that the vast majority of people choose the wrong accumulation strategies. Not only are their results applicable to defined contribution plans, but they are also relevant for defined benefit pension programs and for other long-horizon saving targets.

All the material presented in the paper is from the point of view of a participant in a defined contribution pension system. However, it is applicable to a wider class of problems, including the funding of defined benefit retirement plans by corporations. The findings simply say that systematic contributions proportional to earnings over a career have always led to more wealth at the time of retirement if the investments are in stocks rather than bonds. This information seems completely relevant to an employer who has promised retirement benefits based on final salary and years of service. The defined benefits can be funded with smaller cash contributions owing to the higher rates of return earned on stocks over long horizons.

Not only has an all stocks strategy always bested an all bonds one for all careers exceeding twenty-five years, but it has also always yielded more than the popular fifty-fifty allocation or any other constant mix of stock and bond purchases. While it is impossible to predict the likelihood that this dominance will continue, the evidence favoring stocks for long horizons is overwhelming.

Living Arrangements and Family Support

1. In "Health, Children, and Elderly Living Arrangements: A Multiperiod-Multinomial Probit Model with Unobserved Heterogeneity and Autocorrelated Errors," Axel Börsch-Supan, Vassilis Hajivassiliou, Laurence J. Kotlikoff, and John N. Morris develop a model of living arrangements that should have wide applications in future research in this area. The model overcomes major technical problems that have prevented implementation of the most re-

alistic models in past research. The paper also confirms that living arrangements are predominantly governed by functional ability and to a lesser degree by age. A surprising result is that changes in marital status have little effect on the choice of living arrangement, after controlling for other family characteristics.

Decisions by the elderly regarding their living arrangements (e.g., living alone, living with children, or living in a nursing home) seem best modeled as a discrete choice problem in which the elderly view certain choices as closer substitutes than others. For example, living with children may more closely substitute for living independently than living in an institution does. Unobserved determinants of living arrangements at a point in time are, therefore, quite likely to be correlated. In the parlance of discrete choice models, this means that the assumption of the independence of irrelevant alternatives (IIA) will be violated. Indeed, a number of recent studies of living arrangements of the elderly document the violation of this assumption.

In addition to relaxing the independence assumption of no intratemporal correlation between unobserved determnants of competing living arrangements, the authors also relax the assumption of no intertemporal correlation of such determinants. The assumption of no intertemporal correlation underlies most studies of living arrangements, particularly those estimated with cross-sectional data. While cross-sectional variation in household characteristics can provide important insights into the determinants of living arrangements, the living arrangement decision is clearly an intertemporal one. Because of moving and associated transactions costs, elderly households may stay longer in inappropriate living arrangements than they would in the absence of such costs. In turn, households may prospectively move into an institution "before it is too late to cope with this change." That is, households may be substantially out of long-run equilibrium if a cross-sectional survey interviews them shortly before or after a move. Moreover, persons may acquire a taste for certain types of living arrangements. Such habit formation introduces state dependence. Ideally, therefore, living arrangement choices should be estimated with panel data, with an appropriate econometric specification of intertemporal linkages.

These intertemporal linkages include two components. The first component is the linkage through unobserved person-specific attributes, that is, unobserved heterogeneity through time-invariant error components. An important example is health status, information on which is often missing or unsatisfactory in household surveys. Health status varies over time but has an important person-specific, time-invariant component that influences housing and living arrangement choices of the elderly.

However, not all intertemporal correlation patterns in unobservables can be captured by time-invariant error components. A second error component should, therefore, be included to control for time-varying disturbances, for

example, an autoregressive error structure. Examples of the source of error components that taper off over time are the cases of prospective moves and habit formation mentioned above. Similar effects on the error structure arise when, owing to unmeasured transactions costs, an elderly person stays longer in a dwelling than he or she would in the absence of such costs.

While researchers have recognized the need to estimate choice models with unobserved determinants that are correlated across alternatives and over time, they have been daunted by the high dimensional integration of the associated likelihood functions. This paper uses a new simulation method developed by Börsch-Supan and Hajivassiliou to estimate the likelihood functions of living arrangement choice models that range, in their error structure, from the very simple to the highly complex.

The simulated likelihood method works well and requires a very small number of replications. It easily accommodates highly complex error structures and can handle different error structures without major programming effort

Two main conclusions follow from the estimation results. First, a careful specification of the temporal error process dramatically improves the model fit. It also appears that ignoring intertemporal linkages biases some parameter estimates, although the different specifications produce qualitatively similar coefficients on the substantive parameters.

Second, living arrangement choices are governed predominantly by functional ability and to a lesser degree by age. The analysis confirms that institutions are an inferior living arrangement as measured by a the willingness to spend more in order to avoid entering one. A somewhat surprising result is that changes in marital status do not appear to matter a great deal.

2. In "The Provision of Time to the Elderly by Their Children," Axel Börsch-Supan, Jagadeesh Gokhale, Laurence J. Kotlikoff, and John N. Morris consider the time that children spend with their parents. They conclude that time spent with parents is determined primarily by demographic factors, such as the age of the parent, but that economic factors such as income and wealth play an insignificant role.

There is substantial evidence that support of parents by children has declined in the postwar period. Over 60 percent of the elderly (those over 60) now live alone, compared with only 25 percent in the 1940s. For the old old (those over 85), the fraction living alone has increased from 13 to 57 percent. At the same time, there has been more than a tripling of the rate of institutionalization; today almost one-quarter of the old old live in institutions, compared with only 7 percent in the 1940s. In addition to not living with the elderly, the children of the elderly rarely provide financial transfers to the elderly, and when they do, the amounts are typically quite meager.

One defense of the children's behavior is demographic; the current number of children per elderly parent totals about half the number observed in the

1940s. Since the elderly of today had fewer children than did their parents and have, in some cases, succeeded in outliving their children, the current situation may be of their own making. A second defense is that the relative income position of the elderly has improved, permitting them to live alone and obviating the need for financial transfers from their children. A variety of studies have demonstrated that current poverty rates of the elderly are close to, if not below, those of the nonelderly. Much of the improvement in the relative incomes of the elderly is due to increases in real Social Security benefits legislated in the 1970s. A third point to consider in assessing child support of the elderly involves payment for nursing home care. A good fraction of the elderly in nursing homes are private pay patients. Some of these payments are being made directly by children. It seems plausible that such payments per child measured at constant dollars have increased over time.

The authors emphasize that, while the elderly may need and appear to be receiving less financial help from their children, their needs for companionship and physical assistance may well have increased in the postwar period; the increased longevity of the elderly often means living for years in poor states of health. In addition, those elderly who continue to live will lose a large fraction of their old friends and even some of their children along the way. Most studies of the increasingly separate living arrangements of the elderly conclude that these arrangements reflect the preferences and improved financial means of the elderly. In contrast, Kotlikoff and Morris suggested in prior research that about half the elderly would prefer to live with their children but continue to live apart because of their children's preferences coupled with their children's financial abilities to live apart from their parents.

One reason the jury remains out on family support of the aged involves the issue of time spent by children with their elderly parents. Children's provision of time to their elderly parents is an important, if not the most important, form of economic transfer to the elderly by their children. This papers studies the provision of time by children to their elderly parents. The authors use the 1986 Hebrew Rehabilitation Center for the Aged (HRCA) follow-up survey of Massachusetts elderly and the 1986 HRC-NBER survey of the children of these Massachusetts elderly.

They use these data to answer a number of questions about the provision of time by children to their parents. These questions include, How does the health status of the elderly influence the amount of time given by children? How does the health status of the children influence their provision of time to their parents? Do parents with more income and wealth receive more time from their children? How do the employment status and wage rates of children affect their provision of time? Do children free ride on their siblings' provision of time? Are home care corporations used by children as a substitute for their own time? Do the institutionalized elderly receive more or less time? Are daughters more or less likely to provide time?

The data reveal some clear patterns of time transfers from children to their elderly parents. Children appear to use institutions and home care as a substitute for their own provision of time. Parents who reside in nursing homes or are enrolled in home care programs receive, ceteris paribus, less than half the amount of time received by those in the community. The provision of time is strongly correlated with the age of the elderly parent; other things being equal, the old old receive over twice the time of the young old.

The sex, age, and health status of children are additional important determinants of time provided to the elderly. Male children and younger children spend relatively little time with their parents. Children with poor health spend almost no time with their parents. If the spouse of a child is in poor health, the child also gives very little time, at least according to the model results.

Other things being equal, those elderly who report their health to be "poor" receive over twice the amount of time received by elderly with better self-reports of health. Surprisingly, the degree of elderly disability does not appear to affect the amount of time provided to those elderly not living with their children, although it is a significant determinant in the larger sample that includes elderly living with their children.

The results for the entire sample of children, including those living with their elderly parents, indicate that more time is provided by single children and more time is received by single elderly, at least those who are widowed. There is strong evidence that widowed children spend substantially more time with their elderly parents.

The estimates indicate a small effect associated with higher children's wage rates; children with higher wage rates provide somewhat less time to their elderly parents than other children. In contrast to the modest effect of higher wage rates, the effect of larger values of children's wealth is quite sizable. Wealthier children and children with higher incomes appear to provide less time than poorer children, but the effects are not precisely measured.

To summarize, the results indicate that the main determinants of the amount of time given to parents are demographic. Economic variables, such as wage rate and income levels, appear to play an insignificant role in the provision of time.

On The Life-Cycle Model

In "Wealth Depletion and Life-Cycle Consumption by the Elderly," Michael D. Hurd considers evidence in the Retirement History Survey (RHS) with respect to life-cycle theory. He argues that the data are consistent with life-cycle theory but that they offer no support for a bequest motive as an important determinant of the consumption and saving behavior of the elderly.

Although the life-cycle hypothesis (LCH) of consumption has been the

most important theory for the study of saving behavior, interest in the bequest motive for saving has grown considerably. This has been stimulated by three kinds of empirical results. (1) In simulations of lifetime earnings and consumption trajectories, "reasonable" utility function parameter values lead to savings that are considerably smaller than observed household wealth. This implies that a good deal of household wealth has been inherited. Although, when the date of death is unknown, large inheritances are not necessarily inconsistent with the life-cycle hypothesis, many people believe they indicate that at least part of the bequests are intentional. (2) From estimated earnings and consumption paths it is found that as much as 80 percent of household wealth is inherited. (3) The elderly do not seem to dissave as they age. Because this contradicts a prediction of the life-cycle hypothesis, it has been taken to be particularly damaging to the hypothesis.

In this paper, Hurd first reviews some evidence on how wealth changes as the elderly age. He argues that the best evidence is that the elderly do dissave as required by the life-cycle hypothesis. He then presents findings based on consumption data in the RHS. Hurd concludes that, as measured in the RHS, consumption declines as households age, which is in accordance with the life-cycle hypothesis. If a bequest motive for saving is an important determinant of consumption, the consumption paths of parents and nonparents should differ, but no systematic difference between their consumption paths is found. Hurd's overall conclusion is that the wealth and consumption data in the RHS are consistent with the life-cycle hypothesis; they do not support a role for a bequest motive as a determinant of consumption behavior. Hurd also concludes that in the RHS, observations on both consumption and wealth are consistent with the life-cycle hypothesis of consumption in that both are observed to decline after retirement.

While the findings that consumption and wealth decline with age are consistent with the LCH, they are not inconsistent with a bequest motive for saving: the bequest motive will change the shape and level of the consumption and wealth paths, but they will not necessarily rise. A test for the importance of the bequest motive is based on the assumption that the marginal utility of bequests of a parent is greater than the marginal utility of bequests of a nonparent. This assumption implies that, ceteris paribus, the wealth and consumption paths of a parent should decline more slowly than the wealth and consumption paths of a nonparent. In the RHS, the wealth paths decline at the same rate.

In short, Hurd concludes that the RHS data on wealth and consumption are consistent with the life-cycle hypothesis of consumption but that they offer no support for a bequest motive for saving as an important determinant of consumption behavior. Hurd also gives considerable attention to the substantial measurement problems that plague analysis based on consumption data.

Aging Issues in Developing Countries

In "Patterns of Aging in Thailand and Côte d'Ivoire," Angus Deaton and Christina H. Paxson consider aging issues in less developed countries (LDCs). They emphasize the important differences between developed and less developed countries and the implications of these differences for understanding aging issues and for appropriate directions for research. They draw attention to several issues and questions, emphasizing in particular that the life-cycle model of saving and capital accumulation cannot be applied without modification to economies where the functions of households are very different from those in developed countries.

The authors present and discuss facts about older people in two contrasting developing countries, Côte d'Ivoire and Thailand. They are concerned with standard questions in the aging literature, namely, demographic structure, living arrangements, urbanization, illness, labor force behavior, and economic status. In this paper, they do not attempt to go beyond the presentation of data from a series of household surveys from the two countries. Although recent years have seen increased attention in the demographic and sociological literatures to questions of aging in LDCs, data are still relatively scarce, particularly for Africa, and the authors see their current task as providing stylized facts to help focus further discussion.

There are two research issues that provide the structure for their discussion: household saving behavior and, more broadly, the economics of aging in countries with low living standards but with rapidly expanding shares of old people in the population.

Research on saving behavior in the United States, Japan, and Western Europe has been dominated by permanent income and life-cycle models since their introduction in the 1950s. There has been a good deal less work done on household saving behavior in LDCs, and much of the work that has been done has simply transferred the analytic framework from the more to the less developed context. It is not clear that this is the best way of proceeding. While it makes sense to work with the same basic ideas—that saving can smooth consumption over time and that assets provide a measure of insurance against an uncertain future—there are important differences in environment and in mechanisms, and the same aims may therefore be achieved in very different ways. A much larger share of the population in developing countries is engaged in agriculture, where incomes are very variable, and there are many poor people living close to the subsistence level, so consumption insurance may be of the greatest importance.

The authors find that their exploration of the data allows them to summarize some of what is known and what might usefully be learned:

1. Questions regarding the economic status of the old in LDCs cannot be answered and must be rethought. In more developed countries, where perhaps

nine-tenths of the elderly live by themselves or with elderly spouses, household surveys can tell us a great deal about their living standards. In LDCs, to a greater or lesser degree, older people do not live by themselves, and until a method can be found for measuring intrahousehold allocations, we have no method of assigning welfare levels to them or indeed to other members of the households in which they live.

2. More work needs to be done on the question of whether the source of income (i.e., who earns it) affects what individual members of the household receive. This cannot be done directly, but if the earnings of the elderly are spent differently than other household income, one should be able to detect that fact from consumption data. Data such as those from Thailand show considerable variation in source of income with age, although the patterns are quite different from those in the United States or Western Europe.

3. In the United States and other developed countries, where many elderly people live alone, there has been concern about the possible abandonment of the old. However, such cases seem to be rare; most old people live alone because they want to, and frequency of contract with children is generally high. In Côte d'Ivoire, under current living conditions, abandonment does not seem common because very few old people live alone. There are perhaps more grounds for concern in Thailand, but the population at risk is still small and is probably overstated by the survey results quoted in this paper. However, there is evidence from elsewhere that suggests that these results should not be generalized to all poor countries. In many areas of India, living arrangements for newlyweds are strictly patrilocal, with the result that, after marriage, women are effectively cut off from their parents' family. In turn, they will be looked after in old age by their sons, their daughters having themselves moved to their husbands' families. In consequence, women who fail to produce sons, or fail to produce surviving sons, are likely to fall into destitution as widows.

4. The living arrangements of the elderly will vary from place to place according to marriage arrangements, agroclimatic conditions, and the availability of labor and land. In Côte d'Ivoire, living patterns have been changing in response to the increasing scarcity of land since sons, who were previously guaranteed land nearby, are now often required to set up households at considerable distances. The shortage of land itself reflects a great deal of immigration to the cocoa and coffee areas, an immigration that responded originally to *labor* shortage and that contributed to the destruction of the original lineage system of cocoa and coffee production. One may also wonder whether the pattern of inheritance in northern Thailand—whereby, as a result of the residual stem family system, the youngest daughter typically inherits the land—will continue unmodified into an era where land is increasingly scarce.

5. Individual participation and earnings patterns show the standard life-cycle hump shapes in Côte d'Ivoire and Thailand and presumably do so more widely. However, households act so as to make average living standards

within households much less variable over the life cycle than are the individual patterns. The degree to which this happens in the data is different between the two counties and depends on how household size is measured. Even so, sharing resources between household members is presumably one of the main economic functions of the household. What needs a great deal more research is the extent to which household size and composition adapt to facilitate sharing and to guarantee the best possible living standards to household members. In both Thailand and Côte d'Ivoire, there is a great deal of migration, both seasonal and nonseasonal. In Thailand, the process of household formation is explicitly tied to the pressure on resources within the compound; the departure of a previously married child on the marriage of a younger sibling is therefore as much a matter of economics as of immutable custom. In the panel households in Côte d'Ivoire, there are major differences in membership between 1985 and 1986. The authors emphasize that there is scope for more modeling here, particularly for a simple unifying theory that explains how potential household members decide how to form household groups given the economic opportunities available to them.

6. There are a number of interactions between urbanization and age distributions. Migration tends to lead to young cities and an older countryside, as is the case in Côte d'Ivoire, but much urban growth in LDCs comes from reproductive behavior as well as from migration. The fall in fertility in the demographic transition often begins in the cities, with the result that cities are likely to age more rapidly than more rural areas. The balances between these forces will produce different age distributions in different countries, for example, young cities in Africa and older cities in Asia, and these have a number of repercussions for policy, for example, in the provision of services as well as in the likely effectiveness of older people as a political force.

7. Many LDCs are in a state of transition, not only demographic, but also educational. In both countries examined here, there are very large differences between the educational attainments of the different generations. The *consequences* of these differences are much less clear, and the authors do not wish to subscribe to the view that they always and everywhere undermine the status of the old. But Deaton and Paxson emphasize that models that provide a theoretical framework for the role of the elderly would do well to bear these facts in mind.

8. The life-cycle model of saving and capital accumulation, which has brought so many insights in developed countries, cannot be applied without modification to economies where the functions of households are different. Asset accumulation for old age, with a large share of the capital stock being accounted for (or not accounted for) by life-cycle saving, is not likely to be a very useful model for savings in LDCs. Households can and do provide old-age insurance without an obvious need to accumulate and decumulate assets. The authors' data do not suggest any run down of assets with the age of the household head. Of course, as in more developed economies, heads

have a range of other motives for keeping control of assets for as long as possible.

9. As in developed countries, there is a pronounced household life cycle, with a hump-shaped income, peaking much earlier in Côte d'Ivoire than in Thailand. However, the authors doubt that there is much long-term consumption smoothing associated with these humps, and they tend to attach more importance to saving as a means of smoothing consumption over short-term fluctuations in income that are typically associated with agricultural activities. Indeed, the authors say, it is possible that variations in household structure contribute more to long-term smoothing than do variations in assets.

Social Security Reform

In "Changing the Japanese Social Security System from Pay as You Go to Actuarially Fair," Tatsuo Hatta and Noriyoshi Oguchi consider the implications of changing from the current pay-as-you-go system to one that is actuarially fair. The authors note that, while the merits of a fully funded system are well known to economists, there is a reluctance to move from a pay-as-you-go to a fully funded system because the transition creates instability in benefit distribution and in the macro budget. The authors argue, however, that it is not necessary to accumulate a budget surplus to eliminate the distributional and efficiency problems associated with a pay-as-you-go system. They show that switching to an actuarially fair but unfunded system attains this objective.

The current Japanese public pension system is essentially pay as you go; hence, its rate of return is not actuarially fair for each participant. This is the root of the three problems that the Japanese public pension system faces.

First, the system transfers income intergenerationally. In particular, the generation following the baby boomers is expected to make a large transfer to the baby boomer generation. By the year 2025, the average Japanese worker will have to support three times as many retirees as in 1990. This period, which is characterized by a higher percentage of retirees, is referred as the high-average-age period (HAAP). The arrival of the HAAP will increase the required social security contributions to maintain the promised benefits, resulting in significant income redistributions among different generations. It may even make the very existence of the public pension system uncertain, the authors say.

Second, the system also transfers income within each generation in a way that is difficult to justify. For example, the nonworking wife of a corporate president typically gets a much higher rate of return on her pension benefits than a woman of that company who never marries.

Third, since the social security contribution is not directly linked to the future benefit payments, the current system distorts labor supply.

Had the system been actuarially fair from the beginning, these problems would not have arisen. Once a pay-as-you-go system is in place, however, making it actuarially fair may create new problems. The principal aim of the paper is to evaluate the economic effects of various reform plans that would eventually make the system actuarially fair.

Specifically, Hatta and Oguchi examine the following three plans:

1. *Switch to the Fully Funded System.* This quickly increases the government budget surplus to the level of social security wealth before the arrival of the HAAP.
2. *Switch to the Actuarially Fair System.* This switches the system over to an actuarially fair one before the HAAP. People in the baby boom and subsequent generations will contribute the amount that exactly matches benefits received. The budget surplus never reaches the level of social security wealth; the system never becomes fully funded.
3. *Gradual Shift to the Fully Funded System.* After an actuarially fair system is established as in plan 2, several generations pay taxes at levels greater than the actuarially fair amount until the system is eventually made fully funded. The burdens of building up the fund are shared by several generations.

The authors emphasize first that the Japanese social security system places a heavy burden on the post–baby boom generation by transferring income from it to the baby boom generation. Switching the system to a fully funded one in one generation shifts the heavy burden to the baby boom generation. It also creates national saving in the switching period larger than what would be attained if the system were fully funded from the beginning. Switching to an actuarially fair but unfunded system eliminates the microeconomic problems of the Japanese social security system without causing instability in the transition phase, the authors show.

If accumulation of a budget surplus is necessary to make the system fully funded, it can be done by first changing the system into an actuarially fair but unfunded one and then gradually building up the fund by taxing several generations. Economic effects of such a gradual shift are analyzed.

Nursing Home Stays

1. In "Payment Source and Episodes of Institutionalization," Alan M. Garber and Thomas E. MaCurdy explore the relation between the duration of nursing home stays and the source of payment for nursing home care. The authors conclude that the incentive effects of the subsidies of nursing home care—associated with types of payers—may play an important role in nursing home utilization.

This subject has assumed critical importance as a growing number of pri-

vate insurers begin to offer long-term care insurance, millions of middle-aged and elderly Americans plan for future long-term care needs, and policymakers debate the role that government should play in financing, delivering, and regulating long-term care.

Both private and public initiatives for financing long-term care need accurate projections of utilization, but few studies have examined the effects of insurance on utilization. The size of the insurance subsidy effect on utilization, or moral hazard, is not readily inferred from observed price variation. It is notoriously difficult to gauge the price of nursing home care faced by consumers of this service, in part because price variation reflects differences in the characteristics of nursing homes (e.g., the quality of nursing services, meals, and housing amenities). In the absence of comprehensive, reliable price data or of direct measures of the effects of alternate financing mechanisms on long-term care utilization, studies of the relation between payment source and utilization provide important clues to the likely consequences of changing long-term care insurance benefits.

The measure of utilization that Garber and MaCurdy examine is the length of each nursing home stay. While information about the duration of spells is a key component, additional information is needed to complete any comprehensive picture of nursing home utilization. Comprehensive measures of utilization (or cumulative duration) also require information about both the likelihood that a spell will occur at all and the frequency of readmission. The length of an individual spell has an entirely different interpretation if it is only one of a series of admissions rather than a unique occurrence.

Nursing home utilization also depends on the mode of exit, which in most duration analyses is of little concern. In other medical contexts, the nature of exit may have minimal significance because there is only one way a spell (of an illness, e.g.) may terminate (in death); in economic contexts, even if there is more than one way to end a spell (of unemployment, e.g.), the nature of the exit may be of secondary interest. The type of exit from nursing homes, in contrast, has substantive economic and welfare implications. Nursing home admissions terminate in return to the community, transfer to a hospital, or death. Although the type of discharge clearly matters to the patient, it also affects future long-term care utilization and overall health expenditures. Transfer to a hospital, for example, is often a costly interruption in a lengthy nursing home stay, while discharge to the community may signal resumption of independent living. The length of a nursing home spell, if it terminates in hospital admission, may be short in relation to overall utilization. To accommodate these phenomena, the authors complement the analysis of duration distributions with an investigation of the association between the probabilities of alternate modes of exit and several other factors, including personal characteristics, payment source, and length of the nursing home admission.

Garber and MaCurdy investigate these issues by analyzing data on a sample

of frail, disabled, and otherwise vulnerable elderly men and women who were believed to be likely to enter a nursing home. They were enrolled in the National Long-Term Care (Channeling) Demonstration, a randomized controlled trial of case management as a deterrent to institutionalization. While this sample is not representative of elderly Americans generally, it represents a group of particular interest: persons who are expected to consume a disproportionate share of long-term care and who are likely to be excluded from the purchase of private long-term care insurance. If associations between payment source and duration patterns are significant in this sample, the relation in the general population might be stronger, particularly if the demand for nursing home are among Channeling participants is inelastic.

As planning for and financing long-term care have achieved new prominence in policy circles, there is an urgent need for reliable estimates of the effect of insurance on long-term care utilization. The authors attempt to take a step toward understanding the effect of insurance by measuring the association between payer type and utilization within a high-risk population of older Americans. One might argue that the absence of an association between type of coverage and utilization suggests that the demand for nursing home care is inelastic, an assumption implicit in much of the policy discussion regarding long-term care insurance. Many advocates of broader long-term care insurance coverage believe that widespread adoption of long-term care insurance would not increase the utilization of nursing homes.

Garber and MaCurdy's analysis finds that the distribution of the length of nursing home stays differs substantially among payer types, in ways that may not simply reflect selection. These differences are apparent even in a population of frail elderly individuals who lack social supports and are felt to have "unmet needs." The differences also persist despite the control of the additional covariates incorporated into their model. Although their study was not designed to assess whether the differences in nursing home duration by payer are causal relations, the persistence of strong relations between payer type and duration of nursing home admission despite the selection of the population and the control for additional covariates suggests that the incentive effects of the subsidy of nursing home care may play an important role in nursing home utilization.

The results of the duration analyses reported here suggest that the payment source is strongly associated with the length of nursing home admissions. The covariates have a weak independent association with duration, at least within this population and in the time period studied, but some of them, such as the presence of a severe ADL impairment, are associated with the type of exit. Medicare-financed admissions are much shorter than admissions funded by either Medicaid or some other payment source, and there is a striking early peak in the hazard rate for Medicare admissions.

The type of exit from the nursing home is also highly associated with the

payment source. Exit probabilities reflect the "success" of a nursing home admission, and they also give clues to future utilization of long-term care. Nursing home spells financed by "other," primarily private, payers last nearly as long as Medicaid admissions but are much more likely to end with return to the community. The length of admission and payment source interact, at least for Medicare spells; the longer a Medicare patient is in a nursing home, the more likely is discharge to home. Even long Medicare admissions seem short in comparison to Medicaid admissions and are much more likely to end in return to the community. For the Medicaid admissions, the high rate of discharge to hospitals and the high rate of death are discouraging signs for return to independent living.

The authors emphasize that their results to date should be viewed as suggestive findings, not as definitive answers about insurance effects. But the results make it clear that a complete understanding of nursing home utilization must be based on an adequate characterization of paths leading to and from nursing homes and that it must account for multiple admissions.

2. In "Incentive Regulation of Nursing Homes: Specification Tests of the Markov Model," Edward C. Norton considers whether Markov models can be appropriately used to model transitions in and out of nursing homes. The analysis rests on data from an experiment that tested the effects of performance-based reimbursement on the quality and cost of nursing home care. In prior work, Norton showed that the reimbursement system had a positive effect on both quality and cost. This paper, more narrowly focused on the use of the Markov model, concludes that, while the model may provide a reasonable tool for analysis of these data, it is an imperfect representation of transitions in nursing homes.

The analysis in Norton's previous work used a simple Markov model to estimate transition probabilities between states of health in the nursing home. A comparison of the probabilities for the control group (no incentives) and the experimental group (positive incentives) found them to be different. People in the experimental group stayed for a shorter time and had better outcomes.

The simple Markov model, however, maintains several strong assumptions. For example, it assumes that the transition probabilities are constant over time, independent of past states, and the same for all people. If any of these assumptions are false, the conclusions of the previous paper may be ill founded. This paper extends the analysis to more general models and in doing so subjects the simple Markov model to a series of specification tests. Most of the tests are done on data from the control group nursing homes only so that the effects of the experiment are not mixed with those of the assumptions. The paper tests the following series of assumptions. (1) The probability of being in state j next period depends only on the current state, not on past states. (2) The probabilities are independent of personal characteristics, such as age, sex, race, and marital status. (3) The probabilities are constant over time. (4) The probabilities are independent of how long a person has been in the nursing

home. (5) Nursing homes in the experimental group instantly switched to optimize under the new reimbursement system with no learning period. (6) $P(T) = P(1)^T$ (the basic Markov assumption). (7) Reporting errors by nurses have no effect on the estimated transition probabilities.

This paper contains a summary of the experiment done by the National Center for Health Services Research and the data used in the analysis, followed by a brief review of Norton's earlier results. The remainder of the paper is an extension of the previous analysis.

Norton finds that the Markov model should be viewed as a reasonable but imperfect model of transitions in nursing homes. He concludes that research in this area could benefit from trying other kinds of duration models, such as competing hazard and semiparametric models. These models may have advantages in speed of computation, a more flexible form, and an emphasis on duration and outcome that are important for public policy.