Volume Title: Frontiers in the Economics of Aging

Volume Author/Editor: David A. Wise, editor

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-90304-4

Volume URL: http://www.nber.org/books/wise98-1

Publication Date: January 1998

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

Chapter Author: David A. Wise

Chapter URL: http://www.nber.org/chapters/c7295

Chapter pages in book: (p. 1 - 20)

Introduction

David A. Wise

This is the seventh in a series of volumes on the economics of aging. The previous ones were *The Economics of Aging, Issues in the Economics of Aging, Topics in the Economics of Aging, Studies in the Economics of Aging, Advances in the Economics of Aging,* and *Inquiries in the Economics of Aging.* The papers in this volume discuss the implications of the rapid spread of personal retirement saving, discuss several aspects of health care, investigate important methodological advances in studying aging issues, and consider new aspects of inequality. The papers are summarized in this introduction, which draws heavily on the authors' own summaries.

Personal Retirement Plans

Three papers direct attention to important aspects of personal retirement saving plans: their effect on personal saving, the rapid increase in their use in recent years and future projected use, and the taxing of pension saving. In addition, using the analysis of the saving effect of individual retirement account (IRA) and 401(k) plans as a backdrop, there is an extended discussion of how personal retirement saving might be viewed from a behavioral perspective.

Their Effect on Saving

Contributions to personal retirement saving programs now exceed contributions to traditional employer-provided defined benefit and defined contributions plans. In "Personal Retirement Saving Programs and Asset Accumula-

David A. Wise is the John F. Stambaugh Professor of Political Economy at the John F. Kennedy School of Government, Harvard University, and the director for Health and Retirement Programs at the National Bureau of Economic Research.

tion: Reconciling the Evidence," James Poterba, Steven Venti, and I review the series of analyses on the saving effect of IRA and 401(k) plans that we have undertaken over the past several years and present additional results. We also give considerable attention to the analyses undertaken by others. We conclude that the weight of the evidence provides strong support for the view that contributions to these programs represent largely new saving.

A large fraction of American families reach retirement age with virtually no personal financial assets. The median level of all personal financial assets of families with heads aged 55–64 was only $8,300 in 1991; excluding IRA and 401(k) balances the median was only $3,000. Almost 20 percent of families had no financial assets at all. Other than social security and pension benefits, and illiquid housing wealth, the typical family has very limited resources to meet unforeseen expenses.

The IRA and 401(k) programs introduced in the early 1980s were intended to encourage individual saving. Although very small at the beginning of the decade, by 1989 contributions to all personal retirement saving plans exceeded contributions to traditional employer-provided pension plans. Whether these programs increase net saving can be of critical importance to future generations of older Americans and to the health of the economy in general. The issue remains an important question of economic debate. In a series of papers based on very different methods of analysis we have concluded that a large fraction of the contributions to these accounts represents new saving. Our previous research is summarized here, along with several new results.

As interest in the saving effect of these programs evolved, several other investigators also directed attention to the issue. In some instances, alternative analyses came to conclusions that differed dramatically from ours. Thus in describing our results we have tried to point out the differences between our methods and alternative approaches that have been used to address the same questions.

The key impediment to determining the saving effect of IRAs and 401(k)s is saver heterogeneity. Some people save and others do not, and the savers tend to save more in all forms. For example, families with IRAs also have more conventional savings than families without IRAs. Thus a continuing goal of our analyses has been to consider different methods of controlling for heterogeneity. The methods that could be used when each analysis was conducted were largely dependent on the available data. As new data became available we used alternative and possibly more robust methods to control for heterogeneity.

We first present our results, organized by the method used to control for heterogeneity. In each case the question is whether IRA and 401(k) contributions substitute for conventional financial asset saving. We also discuss closely related results reported by others. While early work in this area focused on the potential substitution between IRA assets and liquid financial assets, subsequent analyses considered the potential substitution between personal retirement saving plan assets and employer-provided pension assets and housing

equity. Thus we also consider other margins of substitution, particularly the possibility that saving in these programs is financed by drawing down home equity. Finally, we address the divergence between our conclusions based on Survey of Consumer Finances data and the parametric analysis of the same data by Gale and Scholz (1994).

We emphasize that no single method can provide sure control for all forms of heterogeneity. Taken together, however, we believe that our analyses address the key complications presented by heterogeneity. In our view, the weight of the evidence, based on the many nonparametric approaches discussed here, provides strong support for the view that contributions to both IRA and 401(k) plans represent largely new saving. Some of the evidence is directed to the IRA program, other evidence to the 401(k) plan, and some of the evidence to both plans jointly. We believe that the evidence is strong in all cases. We have devoted particular effort to explaining why different approaches, sometimes based on the same data, have led to different conclusions. In some instances, we believe that the limitations of the methods used by others have undermined the reliability of their results.

Viewed from a Behavioral Perspective

Increasingly, economists are questioning standard analysis based on strong rationality and typical maximization assumptions. Many of the findings on IRA and 401(k) saving behavior and the saving effects of these plans are left unexplained by standard theories of saving. Thus there is considerable motivation to look more broadly for explanations of saving behavior. A realistic explanation of saving must recognize much broader economic and psychological determinants of individual saving decisions. Using the Poterba-Venti-Wise analysis as a backdrop, David Laibson in his comment considers the insights that behavioral economics bring to the discussion of personal retirement saving. He discusses four behavioral phenomena: bounded rationality, self-control problems, peer group sensitivities, and overoptimistic beliefs. He then discusses empirical strategies, suggested by the behavioral phenomena, that might be used to further our understanding of saving behavior. He finds that "most of the behavioral analysis that I have reviewed implies that tax-deferred retirement instruments will raise net national savings." After several cautionary notes, he goes on to conclude: "Nevertheless, the prospects for 401(k)s and similar tax-deferred retirement instruments seem bright. Social psychology research has identified a core set of features of successful behavioral interventions. The 401(k) seems to have been designed by someone who intuitively or formally understood those lessons."

Implications of Their Rapid Spread

About half of families now are eligible for a 401(k) plan, and the use of these plans is spreading rapidly. In "Implications of Rising Personal Retirement Saving" James Poterba, Steven Venti, and I consider the projected use of

these plans and implications for the retirement saving of future retirees. We conclude that the cohort that reaches age 65 between 2025 and 2035 will have 401(k) assets that greatly exceed the personal financial assets of current retirees and that these assets are likely to exceed their social security assets, perhaps by a great deal.

Contributions to IRAs grew rapidly until 1986, when $38 billion was contributed to these accounts. The Tax Reform Act of 1986 curtailed this program, and by 1990 contributions had fallen to only $10 billion. They were only $8 billion in 1994. On the other hand, the 401(k) plan has grown unimpeded since 1982. Now contributions to the 401(k) plan alone are greater than contributions to traditional employer-provided defined benefit and defined contribution plans combined. In 1993, 401(k) plan contributions exceeded $69 billion. Approximately 45 to 50 percent of employees were eligible for 401(k) plans in that year, and over 70 percent of those who were eligible to contribute did in fact make contributions.

The increase in personal retirement saving can have important implications for the accumulation of retirement saving for future generations of retirees. Now, a large fraction of families approach retirement with virtually no personal financial asset saving. The median of personal financial assets of Health and Retirement Survey (HRS) families—whose heads were ages 51–61 in 1992—was approximately $7,000. This includes all financial assets held outside IRAs, 401(k)s, and related retirement saving accounts. Perhaps half of all families rely almost exclusively on social security benefits for support in retirement. The spread of 401(k) plans in particular could change this picture substantially. In this paper we simulate the 401(k) assets of future generations of retirees and compare these assets with the social security and other assets of the households who are approaching retirement now.

Our goal is to project 401(k) assets of households who will retire 35 or 40 years from now. We direct attention in particular to the cohort that was age 33 in 1993, and will be age 65 in 2025. We compare the projected 401(k) assets of this cohort to the assets of the HRS respondents. We first trace backward to obtain approximate lifetime earnings histories of the HRS respondents. Lifetime earnings are grouped into 10 deciles, assuming that over their careers household earnings were in the same decile. Contributions to 401(k) plans are projected for each lifetime earnings decile. Thus we are able to ask what level of 401(k) assets such families would have accumulated, in 1992 dollars, had they had the same earnings histories as the HRS respondents but different amounts of contributions to 401(k) plans. We base the projections on the past growth in 401(k) participation rates and on the fraction of earnings contributed to the plans. The growth in participation since their inception, however, has been enormous and simple projections—based on recent increases in participation—are not very meaningful. Therefore, we make what we believe to be plausible inferences about future participation.

We have projected the accumulation of 401(k) assets at retirement for the cohort that was age 25 in 1984 and the cohort that was age 15 in 1984. The cohort 25 projections are based on what we hope are plausible assumptions about future 401(k) participation rates. Indeed, our intention is that these projections be conservative and thus likely to underestimate realized contributions. The cohort 15 projections are further from historical rates but we hope are also based on plausible assumptions about potential future participation. For comparison, we have also made projections assuming universal 401(k) coverage.

In each case, the accumulation of 401(k) assets is large compared to current wealth at retirement. Because a large fraction of current retirees depend almost entirely on social security benefits for support in retirement, we have compared future 401(k) assets to social security wealth. Our cohort 25 projections suggest that when this cohort reaches retirement age—they will be age 65 in 2025—their average 401(k) assets are likely to exceed their average social security assets. But the projections also suggest that relative to social security wealth, 401(k) assets will vary a great deal with lifetime earnings. While this is surely true, we are uncertain about the exact magnitude of the variation by earnings decile. The projections suggest that the lowest earnings decile may have very little in 401(k) assets. But for families with lifetime earnings above the lowest two or three deciles, 401(k) assets are likely to be a substantial fraction of social security wealth. For families with lifetime earnings above the median, 401(k) assets could exceed social security wealth, and this would almost surely be true for families in the top four earnings deciles.

Universal 401(k) participation would likely yield 401(k) assets at retirement greater than social security wealth for all but the lowest lifetime earnings decile, and possibly for the lowest decile as well. The intermediate cohort 15 projections yield 401(k) accumulations that could represent a substantial fraction of social security wealth for lifetime earnings histories as low as the 2d decile.

Thus we believe that 401(k) assets will almost surely be an important component of the retirement wealth of future generations of retirees and could be the dominant component for a large fraction of them.

The Taxing of Retirement Saving

In "The Taxation of Pensions: A Shelter Can Become a Trap" John Shoven and I consider the greatly increased tax burden imposed on pension saving since 1982. In particular, we address the implications of the excess distribution and excess accumulation taxes that were introduced virtually unnoticed in 1986 legislation. We also consider the important implications of full estate tax imposed on pension assets, which before 1982 were not subject to estate taxes at all. We conclude that these taxes can severely limit the advantage of personal retirement saving and pension assets that pass through an estate can be virtually confiscated by tax rate that are often as high as 95 percent or more. Ironi-

cally, the excess distribution and the excess accumulation taxes were eliminated as part of the Tax Relief Act of July 1997, we believe in no small measure due to the attention that our analysis brought to this issue.

The recent legislation that raised the minimum wage was accompanied by the Small Business Job Protection Act of 1996, which, among other things, temporarily (for years 1997–99) suspends the 15 percent excise tax on "excess distributions" from qualified pension plans. Surely, few people know about the excise tax in the first place, let alone its suspension. In fact, while people are keenly aware that pensions allow them to save before-tax dollars and compound their investment returns without current taxation, it is our impression that very few people know how pension assets are taxed on withdrawal or on the death of the owner of the pension. In this paper we present a comprehensive examination of the taxation of pensions, with particular emphasis on large pension accumulations. The analysis answers a number of questions: (1) How do the excess distribution excise tax and its companion excess accumulation excise tax work? How do these taxes interact with the personal income tax systems and the estate tax? (2) Should only high-income individuals be concerned with these taxes or might they be imposed on people with relatively modest incomes? (3) Are pension plans still attractive saving vehicles once these excise taxes are applicable? For example, should someone whose base pension plan is likely to trigger either the excess distribution tax or the excess accumulation tax participate in a supplemental 401(k) plan? (4) Are pensions equally advantageous for stock investments and bond investments? If not, which assets should be held inside a pension plan and which should be held outside the plan? (5) Does it always make sense to delay distributions from pension plans as long as possible, thereby maximizing the tax deferral advantage that they offer? (6) How are pension accumulations treated when they are part of an estate? We focus on these microeconomic issues without discussing the social desirability of current tax policy. However, we should acknowledge at the outset that we see little economic merit in a penalty excise tax applying to people who save "too much" through the pension system. Most observers of the U.S. economy agree that the country's saving rate is too low. Since we know that savers only capture a fraction of the social return on their investments, it is unclear why the biggest savers in the economy should be penalized.

The Small Business Job Protection Act of 1996 is just the latest in a series of bills over the past 15 years that has changed the way pensions are taxed. To illustrate how radically the rules have changed, we consider four individuals with exactly the same wealth (and composition of wealth) at the time of their deaths, each of whom died at age 70. They differ only in the date of death—1982, 1984, 1988, or 1996. The estates, all valued at $1.9 million in 1996 dollars, are composed of $600,000 in nonpension assets including a house, $1.2 million in a defined contribution pension plan, and $100,000 in a supplemental plan such as an IRA or a Keogh plan. The estates of these individuals faced radically different tax laws. One important difference is that before 1983

pension accumulations were completely exempt from the estate tax. The result is that in 1982 the heir was able to consume more than 60 percent of the value of the inherited supplemental pension plan; the combined tax rate was less than 40 percent. In contrast, by 1996—because of the estate tax and the excess accumulation tax in particular—the heir could spend less than 15 percent of the value of the inherited supplemental plan. This case is far from extreme. We describe cases in which the total marginal tax rate on assets in qualified pension plans passing through an estate ranges from 92 to 96.5 percent. The highest such rate we have seen exceeds 99 percent. Thus we believe that any pension saving strategy adopted more than a few years ago needs to be reviewed, given how drastically the rules have changed. And, because large pension accumulations are taxed so heavily when they pass through an estate, withdrawing pension assets before death, if at all possible, needs to be considered.

The excess distribution tax and the excess accumulation tax were included in the Tax Reform Act of 1986 to prevent people from taking advantage of the favorable tax treatment of pensions to amass wealth beyond what is thought to be reasonably necessary for a comfortable retirement. We believe that the wisdom of this policy is open to question. People who increase saving because of the tax shelter opportunity offered by pension plans, or for other reasons like the payroll deduction feature of many pension plans, do not reduce the resources available to the rest of the population. In fact, individual savers reap only a portion of the social return of the incremental capital. If the extra pension saving results in extra capital for the economy, then this extra capital pays corporation income taxes and the pension saver ultimately pays personal income taxes, which improves the overall budget picture for everyone in the economy. The social return to the capital significantly exceeds the private return received by the pension saver.

Our analysis has led to several striking results:

- The tax rates faced by pensions deemed "too large" can be extraordinarily high. The marginal rate on distributions over $155,000 can be roughly 61.5 percent. The effective marginal tax rate faced by large pension accumulations passing through an estate can dwarf this rate, however, reaching 92 to 96.5 percent.
- These high tax rates, which include the excess distribution and the excess accumulation taxes, can be faced by savers who do not have extraordinarily high incomes. The "success tax" is not limited to the rich but rather primarily affects lifetime savers. For example, someone who works between ages 25 and 70, makes $40,000 at age 50, and contributes 10 percent to a pension plan invested in the S&P 500 will likely be penalized by the success tax for overusing the pension provisions.
- The advantage of pensions relative to conventional saving is greatly reduced and in many cases eliminated once accumulations exceed the amounts that will trigger the success tax. Even in cases where additional

pension saving still provides more resources in retirement than conventional saving, when the plan owner dies, the heirs get less than they would have if the saving had been done outside of a pension plan.

- The advantage of pension saving is reduced by the availability of tax-advantaged investments outside of pension plans. Examples are tax-free municipal bonds and tax-efficient low-dividend stock portfolios.
- Not only is there little, if any, incentive to continue to save via pensions once the excess distribution and excess accumulation taxes become applicable, there is a strong incentive to withdraw money while living rather than risk the nearly confiscatory tax rates faced by pension assets transferred through estates. This means that there is an incentive to consume more in retirement than would otherwise be the case.
- Individuals can realize significant efficiency gains by allocating their investments appropriately between pension accounts and outside holdings. Simply locating assets in their most advantageous environment could improve the net proceeds of saving by almost 25 percent.

Now that it is recognized that pensions are the primary vehicle for personal saving in the economy, we emphasize that a careful reconsideration of the legislation of the early and mid-1980s is called for. Once the excess distribution and excess accumulation taxes are understood, we believe these taxes are likely to become quite effective at discouraging pension saving and hence will reduce economic welfare in an economy that surely has lower than optimal saving now.

Health: Spending Patterns and Implications and Effect on Work

A Forty-Year Perspective

The rising cost of health care has become an important public issue in almost all developed countries. In "The Medical Costs of the Young and Old: A Forty-Year Perspective" David Cutler and Ellen Meara investigate the nature of the rise in health care costs in the United States and the gains in health outcomes associated with this spending. They find that, just as medical spending at a point in time is concentrated among high-cost users, spending growth has also been concentrated among the high-cost users. They find further that the growth in expenditures for high-cost users has been accompanied by gains in health outcomes for these users.

It is widely known that medical costs have increased over time. In the United States, as in most of the developed world, medical spending growth has exceeded income growth by several percentage points per year for three decades or longer. In country after country, the cost of medical care has become a major public sector issue.

But much less is known about what medical spending is buying us. Is medical spending valuable or wasteful? Should we want to limit total spending or increase it? The answer to this question is by no means clear. On the one hand is voluminous evidence that medical spending conveys great value. Randomized clinical trials, for example, routinely document the benefits of new pharmaceuticals and medical devices. Cutler and Meara suggest that most people would prefer today's medical system to the medical system of 30 years ago, even given the much higher cost of medical care today. This suggests that people are on net better off because of the additional medical care spending than they would be without it.

On the other hand is a great sense that medical care often brings little in the way of health benefit. Nearly one-third of Medicare spending occurs in the last six months of life, which has been interpreted as evidence that a lot of medical care is wasted on those who will in any case not survive. Other studies such as the RAND Health Insurance Experiment show that putting people in less generous insurance policies reduces their spending on medical services but does not affect their health. And direct estimates of the value of medical care typically find that, at the margin, a substantial amount of medical care has little or no health benefit.

Cutler and Meara try to understand why medical care has become so expensive over time and what has been its value to society. They focus particularly on medical spending by age. Their analysis is based on periodic surveys of national health expenditures conducted in 1953, 1963, 1970, 1977, and 1987.

Their analysis of age-based spending documents two conclusions. First, there has been a dramatic change in the distribution of medical spending over time. While spending on medical care has increased for all people, it has increased disproportionately for the very young and the old. Over the 24-year period from 1963 through 1987, per person spending on infants increased by 10 percent per year, and per person spending on the elderly increased by 8 percent per year, compared to only 4.7 percent per year for the "middle-aged." The share of medical care spending for infants and the elderly doubled from 15 to over 30 percent.

The authors further show that essentially all of the disproportionate growth of spending for the very young and the old is accounted for by high-cost users within those groups. For infants, 90 percent of the excess spending increase over the middle-aged is accounted for by the top 10 percent of the spending distribution. For the elderly, the equivalent share is 70 percent. Thus, the authors emphasize, to understand the concentration of medical spending by age, we need to understand the concentration of spending among high-cost users.

Cutler and Meara also consider who the high-cost users of medical care are. They show that a substantial amount of high-cost medical use is associated with the increasing technological capability of medicine. Among infants, high-cost users are premature babies with substantial respiratory or other acute conditions. For the elderly, high-cost users are generally patients with severe car-

diovascular problems or cancer. For both infants and the elderly, the capacity to devote many more resources to the most pressing cases has increased over time.

Finally, they consider how health outcomes for premature infants and the sick elderly have changed over time. They find substantial health improvements in most of the categories of high-cost medical care. Infant mortality among very low birth weight infants has fallen substantially at exactly the time when the cost of these infants has risen most rapidly. And mortality improvements among the elderly have been especially prominent in cardiovascular care, where spending increases have been most dramatic. The authors emphasize that their analysis is not causal: "We do not have any direct link between the technologies we discuss . . . and the outcomes we analyze." They conclude, however, that their results suggest such a link is plausible.

End-of-Life Expenditures

Perhaps a quarter of Medicare expenditures each year are for the care of recipients who die in that year. This has led some to conclude that much of end-of-life expenditure must be wasteful. But in their paper "Diagnosis and Medicare Expenditures at the End of Life" Alan Garber, Thomas MaCurdy, and Mark McClellan find that characteristics of high-cost users who die are very similar to those of like patients who live. Thus they conclude that it will be difficult to formulate policies that limit expenditures for persons who will die without at the same time limiting care for those who may live.

Expenditures for health care at the end of life have been the object of considerable policy interest. To many observers, aggressive health care administered shortly before death is wasteful, at least in retrospect, and there is no doubt that elderly Americans use health care heavily before they die. In 1990, the 6.6 percent of Medicare recipients who died accounted for 22 percent of program expenditures. Some question the benefits of such care, especially because severe morbidity and disability may compromise any years of life potentially gained from such treatments in the very old. The increasing use at advanced ages of costly and aggressive health interventions, whose impact on both the quality and length of life is often unknown, exacerbates these concerns. For example, the use of radical prostatectomy for invasive prostate cancer increased threefold between 1983 and 1989, and the rate at which both new and old major operations are performed in the elderly increased dramatically from 1972 to 1981 and from 1984 to 1991.

Nonetheless, the policy significance of end-of-life care is uncertain. Although they consume a large share of health dollars, patients nearing death are not the main source of large Medicare payments. Most high-cost elderly patients are survivors, not decedents. In 1990, for example, 8 percent of Medicare enrollees with over $10,000 in Medicare costs accounted for 65 percent of all Medicare expenditures. Even if every enrollee who died that year generated more than $10,000 in payments, they could have only been responsible for at

most one-third of this total. In addition, some studies suggest that the elderly who are most likely to die are not responsible for the greatest expenditures.

An even more important limitation of an exclusive focus on end-of-life expenditures is its implicitly retrospective nature. If health care providers knew ex ante which patients would die despite their interventions, expenditures on dying patients could potentially be reduced. Yet often it is impossible, even with detailed clinical records, to predict *individual* patient mortality with sufficient specificity to influence medical decision making. Furthermore, specific policies designed to reduce end-of-life expenditures, such as advance directives and "do not resuscitate" orders, seem to have had only a modest impact on overall expenditures.

Although there is growing disenchantment with policies designed to identify and avert "futile care" at the individual patient level, the consequences of population-level policies merit further exploration. Little is known about the impact on the cost and outcomes of elderly populations with major illnesses of policies that would lead to a generally more parsimonious approach to health care delivery. Would efforts designed to limit major operations, aggressive diagnostic approaches, or care delivered in intensive care units lead to an increase in mortality rates, health complications, and other adverse outcomes in the populations with major illness? Would substantial cost savings result from associated reductions in technology use?

As a first step toward addressing these questions, Garber et al. explore the basic characteristics of Medicare recipients during the time approaching their deaths. What characteristics of patients are associated with high-cost care and high probability of death? Can these characteristics be used to identify individuals and populations with major illnesses prospectively, to guide policies on the use of intensive treatments? What is the distribution of costs and survival across high-risk illnesses, and across individuals with a given illness, and how are they correlated? They report the results of analyses of expenditures patterns for Medicare decedents and the associated Medicare expenditures.

Like previous investigators, the authors find that Medicare expenditures for decedents decline with advancing age. Overall, their results confirm that decedents have disproportionately high Medicare expenses. But a comparison with high-cost users suggests that many of the characteristics of decedents are similar to those of high-cost Medicare enrollees who do not die during a given period. The authors conclude that, at least based on the information available in Medicare claims files, it will be difficult to formulate policies to limit expenditures for individuals in the last year of life without simultaneously limiting health care delivered to sick Medicare recipients who have the potential to survive.

Correlation of Expenditures within Families

Work reported in the previous volume on the economics of aging concluded that the potential persistence over time in individual medical expenditures does

not critically limit the feasibility of medical saving accounts, in conjunction with catastrophic health insurance, as a means of helping to control health care costs efficiently (Eichner, McClellan, and Wise 1998). The problem posed by persistence could, however, be worse to the extent that expenditures by family members are correlated, so that persistent expenditures by one member tend to be accompanied by persistent expenditures by other family members as well. In his paper "The Impact of Intrafamily Correlations on the Viability of Catastrophic Insurance" Matthew Eichner concludes that such correlation is not sufficient to importantly limit the feasibility of medical saving accounts.

This paper explores the relationship between health care expenditures of spouses. If expenditures are due largely to random shocks, the household's medical expenditures are smoothed by two or more family members each drawing from the distribution. On the other hand, if the shocks are positively correlated, the potential exists for negative wealth shocks that are greater than those that would be predicted based on studies of the persistence of individual medical expenditures over time. Under traditional systems of insurance, with relatively low coinsurance levels, the consequences of any putative positive correlation across family members in expenditures are relatively mild. But under the sort of high-deductible insurance that is currently attracting interest from the policy community, the wealth and utility effects may be appreciable. And under systems that include both high-deductible insurance and medical saving accounts, intrafamily correlations might dramatically change the accumulation of wealth in such accounts over a working lifetime.

Matthew Eichner seeks to determine the reasons for intrafamily correlations and then to determine where in the distribution of expenditures the correlations appear to be strongest. The empirical work in this paper—based on data from two Fortune 500 firms—suggests that, while correlation of expenditures among married partners is large at the low end of the expenditure distribution, the relation diminishes appreciably in the upper ranges. For example, Eichner estimates a correlation for positive expenditures of 0.41 for firm 1 and 0.69 for firm 2. But the correlations in expenditures above $4,000 are only 0.13 and 0.14, respectively. Thus the correlations are low in the ranges relevant for a discussion of catastrophic health insurance schemes. His analysis therefore suggests that intrafamily correlation in expenditures does not appreciably increase persistence in expenditures made toward a family insurance plan and thus does not appreciably affect the feasibility of medical saving account insurance schemes.

Health Shocks and Labor Supply

It has long been understood that persons in poor health are more likely than otherwise similar persons in better health to leave the labor force. The effect of specific health shocks, however, has not been analyzed. Mark McClellan, in "Health Events, Health Insurance, and Labor Supply: Evidence from the Health and Retirement Survey," analyzes the effect on labor supply of a variety

of specific health events, many of which have very large consequences for labor supply.

The economic consequences of health problems are reported to be enormous. For example, many investigators have concluded that the cost to society of common health problems such as heart disease, diabetes, and cancer is many billions of dollars per year in terms of lost work productivity, intensive medical treatments, and additional supportive care. However, McClellan points out that estimates have several important limitations. Few data sets have incorporated detailed information on health problems and economic circumstances such as retirement. Consequently, most existing studies have had to combine data from different sources, possibly missing important correlations between variables such as insurance availability and the occurrence of health problems.

McClellan uses the first two waves of the HRS to provide insights into how changes in health status affect two issues of considerable policy interest: health insurance coverage and labor supply for middle-aged Americans. His analysis of new health events in older Americans in the HRS suggests several conclusions about the effects of health problems. First, new health events of all types are more prevalent in individuals with lower education, income, and wealth and are more prevalent in individuals with other prior health conditions as well. These relationships persist after adjusting for age. Second, health events may be quite heterogeneous in nature, and thus in their consequences for functional status and expectations about future functional status, consumption, and survival. Only a minority of new health events lead to substantial short-term functional impairments, even for major events such as heart attacks and strokes. Old health problems (or health problems for which information was not obtained in the HRS) are also important in explaining functional declines.

Third, different types of health events have quite different consequences for health insurance coverage and labor supply. Major health events have particularly large effects on retirement decisions, and these effects go well beyond the consequences of the events for functional status. For example, males with major events associated with major functional status declines leave the labor force at rates over 40 percentage points higher than males with major functional status declines in the absence of new health events. New chronic health problems have milder, though significant, effects on increasing rates of labor force exit beyond their association with functional declines alone. In contrast, health problems that are unlikely to have long-term consequences for health (accidents) are not associated with additional labor force departures. Though these health events have enormous significance for labor force departure rates, they have only modest impact on individuals' self-reported retirement status, especially for males in couples and single females. Examining the subsequent labor supply of individuals with these events is thus a question of considerable importance for understanding the long-term impact of health events.

In conjunction with their effects on labor supply, health events also have substantial effects on health insurance coverage, especially for males. Health

events are associated with small increases in the probability of having health insurance, despite the fact that they tend to lead to reductions in private insurance coverage, particularly for males and for individuals without retiree insurance coverage. These reductions in private insurance coverage are offset by increased coverage through government insurance programs, primarily Medicare, as a result of qualification through the disability insurance system. These insurance changes are more related to the actual occurrence of disability than the labor supply changes, though major health events do lead to more switches to government insurance regardless of functional status change.

Methodological Innovations

Measuring Household Consumption and Saving

An important advance in the HRS is the use of bracketing methods to measure household assets, consumption, income, and other values. If when asked, for example, "How much money do you have in bank savings accounts?" the respondent gives a dollar value, the information is clear. But a large fraction of respondents—in some instances as many as 40 percent or more—give no answer or say they do not know. In the HRS, these respondents are pressed further by asking whether the amount is, for example, less than $5,000. If the response is no, the respondent is asked, for example, whether it is between $5,000 and $20,000. And so forth. The sequence is referred to as "unfolding brackets." Following a dollar value question with unfolding bracket questions very substantially reduces the nonresponse rate, perhaps reducing nonresponse from 40 percent to only a few percent in some instances. But the brackets also present a problem called "anchoring"—that is, to the interval at which the sequence of bracket questions starts. In "Consumption and Savings Balances of the Elderly: Experimental Evidence on Survey Response Bias," Michael Hurd, Daniel McFadden, Harish Chand, Li Gan, Angela Merrill, and Michael Roberts consider the effect of anchoring on estimated consumption and saving. They find that variation in bracket starting values can change estimated values enormously—in some examples by as much as 100 percent.

The authors emphasize that a prerequisite for understanding the economic behavior of the elderly, and the impacts of public policy on their health and well-being, is accurate data on key economic variables such as income, consumption, and assets, as well as on expectations regarding future economic and demographic events such as major health costs, disabilities, and death. Standard practice is to elicit such information in economic surveys, relying on respondents' statements regarding the variables in question. The authors emphasize, however, that economic studies are often too sanguine about the reliability of subjects' statements regarding objective economic data. They focus on biases induced by anchoring to prompts presented by questions on eco-

nomic variables and show that anchoring bias is a significant issue in consumption and savings variables of key interest for study of the elderly.

Anchoring describes a family of effects observed in many psychological studies of beliefs about uncertain quantities. A psychological explanation for the phenomenon of anchoring is that a prompt creates in the subject's mind, at least temporarily, the possibility that the uncertain quantity could be either above or below the prompt. This could result from classical psychophysical discrimination errors, or from a cognitive process in which the subject treats the question as a problem-solving task and seeks an appropriate framework for "constructing" a correct solution, utilizing the prompt as a cue. Both formal and informal education train individuals to use problem-solving protocols in which responses to questions are based not only on substantive knowledge but also on contextual cues as to what a correct response might be. Consequently, it should be no surprise if subjects apply these protocols in forming survey responses.

In psychological experiments, anchoring is found even when the gate amount is explicitly random, suggesting that there is more to anchoring than "rational" problem solving. This could happen because subjects are subrational, making cognitive errors and processing information inconsistently, or because they are "superrational," going beyond the substantive question to "model" the mind of the questioner and form superstitious beliefs about the behavior of nature.

The study uses an experimental module in a panel of the survey of Asset and Health Dynamics among the oldest old to establish that anchoring can cause significant biases in unfolding bracket questions on quantitative economic variables. In the case of savings, variation in starting values for unfolding brackets from $5,000 to $200,000 induces a 100 percent difference in estimated median savings. The anchoring is even stronger for consumption: increasing the starting value for unfolding brackets from $500 to $5,000 induces nearly a doubling of estimated median consumption.

The authors find that a simple model in which each gate presented to the subject can induce discrimination errors is successful in explaining much of these anchoring effects. Thus variation in unfolding bracket gates, in tandem with the discrimination model or an alternative model of anchoring, promises to be effective in identifying the effects of anchoring and in undoing most of these effects. The authors recommend that survey researchers who wish to use unfolding bracket elicitations adopt experimental variations in their designs that permit identification and correction of anchoring biases and that they exercise caution in imputing economic variables based on stated brackets.

Accounting for Uncertainty in Social Security Forecasts

Social security forecasts of reserve fund balances typically provide median, high, and low forecasts. Most assessments of the financial status of the fund

are based on the median forecasts and tend to ignore the high and low estimates. But even these high and low estimates provide a grossly inadequate picture of the true uncertainty faced by the system. In their paper "Stochastic Forecasts for Social Security" Ronald Lee and Shripad Tuljapurkar begin to develop a method that does provide a realistic picture of the uncertainty inherent in projections of fund reserves.

Population aging is projected to have a major impact on the federal budget in the next century, in part through its effects on health costs through Medicare and Medicaid, and in part through its effect on the retirement system. Despite the inevitability of the aging of the baby boom, and its dramatic effect on the old age dependency ratio, a great deal of uncertainty remains about the extent of future population aging. On the one hand, we do not know how rapidly mortality will decline and how long people will be living, and on the other hand, we do not know what fertility will be, and therefore we do not know how large the labor force will be in the future. Immigration adds another layer of uncertainty, but it is not considered in this paper. In addition to these demographic sources of uncertainty, there are economic variables with important effects on the future finances of the social security system, notably the rate of growth of productivity or real wages and the level of the real interest rate. Lee and Tuljapurkar emphasize that rational and foresightful planning for the next century must somehow take into account not just our best guesses about the future but also our best assessments of the degree of certainty about the future.

Currently, scenario-based forecasts are in virtually universal use. In these, the forecaster chooses, for each variable, a medium or best-guess trajectory, complemented by high and low trajectories. Then one trajectory for each variable is grouped with others in a scenario, or collection of trajectories. These scenarios may, in turn, be described as "high," "medium," or "low," or by other terms such as "optimistic" and "pessimistic." A high scenario would typically be based on a high trajectory for fertility and migration and a low trajectory for mortality; this combination would yield high population growth. Alternatively, a low-cost scenario for social security would bundle together high fertility and migration with high, rather than low, mortality; this scenario would generate the lowest old age dependency ratio. For the social security forecast, this low-cost demographic scenario would then be combined with high trajectories for productivity growth and (perhaps) for interest rates. The authors emphasize that this approach provides a very inadequate appreciation of the true uncertainty faced by the social security system.

In this paper, the authors build on earlier work to develop stochastic forecasts of the social security reserve fund. Their forecasts are based on stochastic models for fertility, mortality, productivity growth, and interest rates, which are fitted on historical data at the same time that they are constrained (except for mortality) to conform in long-run expected value to the middle assumptions of the Social Security Administration (SSA). The four factors the authors have modeled have the largest effect on the SSA forecast for 2070. Through analysis

of the SSA projections, they find that the four variables they treat as stochastic account for 76 percent of the width of the SSA low-cost–high-cost range in 2070, and 63 and 70 percent in 2020 and 2045, respectively.

The authors emphasize that they are still at early stages of digesting their stochastic forecasts of the finances of the social security system and are still exploring new ways in which these forecasts and experiments might be useful and informative. Perhaps, they say, the most promising use is to test the consequences of a range of strategies for dealing with the uncertainty about the system's finances. Is it better to wait and see, adjusting policy continuously as we gain information? Or is it better to accumulate large reserves early on, to provide a buffer against unlikely but possible transitory insults to the system? Or should policy simply be set to deal reasonably with the mean trajectory, ignoring the uncertainty?

Views of Inequality

Health Inequality over the Life Cycle

In earlier work Angus Deaton and Christina Paxson documented that in several developed and developing countries inequality in income, consumption, and earnings increase with age. They now ask whether this phenomenon is true of other individual attributes as well. In "Health, Income, and Inequality over the Life Cycle" they begin to explore the relationship between age and inequality in health status. They find substantial evidence that, indeed, measures of health status become more widely dispersed within any given birth cohort as that cohort ages.

In their previous work, Deaton and Paxson used data from the United States, Great Britain, Taiwan, and Thailand to document that inequality increases within cohorts with age for consumption, income, and earnings. In this paper, they extend the analysis to two health-relevant measures, the body-mass index (BMI) and self-reported health status (SRHS). They use data on more than 500,000 adults in the United States to track birth cohorts over time and to document the evolution of the two measures with age, looking at both cohort means and within-cohort dispersion. They also consider the life cycle profile of dispersion in income and health jointly, presenting evidence separately for men and women, and for blacks and whites.

Their original work on consumption and income inequality was motivated by the prediction of the standard theory of intertemporal choice that within-cohort inequality in consumption and income should increase with cohort age, at least up to the date of retirement. Although the theory has no immediate extension to processes other than income and consumption, the authors point out that there are a number of reasons to extend the analysis to health status.

First, they wish to investigate the generality of the proposition that dispersion increases with age. For the four countries where they looked earlier, it

is true of income, consumption, and earnings. They are now curious as to whether the proposition is true for other state variables, such as weight, BMI, SRHS, dexterity, intelligence, or ability to complete specified tasks.

Second, while health status is interesting as an example, it is also important in its own right. Inequalities in income and consumption are of concern because they are important components of welfare. But as we move from a narrow, economic measure of well-being toward broader definitions, health status has the most immediate claim on our attention. Nor is health status independent of economic status. Access to health care is expensive in the United States, so that health shocks can have a direct effect on wealth. Independent of this effect, there is a well-documented but poorly understood relationship linking socioeconomic status to a wide range of health outcomes. There is also a literature linking health status to relative deprivation, or to the income distribution. Third, it is plausible that the theoretical reasons that consumption, income, and earnings processes disperse also apply to health status.

Deaton and Paxson present evidence on life cycle patterns of BMI and SRHS, as well as on their relationship with income. They emphasize that it is important to explore differences between people in their health, even in the absence of an agreed methodology for thinking about inequality in health status, or even about health status itself. But by the same token, it is important to be cautious about attributing causality to any of their findings. Income and their measures of health status are linked in many different ways, through ability to pay for health, through education that is correlated with income, through lifestyle choices—such as whether to smoke and what to eat—that are conditioned by income, race, and sex.

From their findings, the authors highlight the following:

- There is ample evidence for the proposition with which they began, that their two measures of health status become more widely dispersed within any given birth cohort as that cohort ages. They view this as evidence in favor of a cumulative random model of health status.
- The rate of dispersion with age of BMI, but not of SRHS, is much more rapid for women than for men. BMI is more variable among women to start with. SRHS is more variable among young women than among young men, possibly reflecting pregnancy.
- Health status (positively measured) is positively correlated with income within cohort-year-sex cells. The correlation is lowest for the young, increases until age 50–60, and then diminishes. BMI is uncorrelated with incomes for men, but negatively correlated with incomes among women. This correlation is highest in middle age.
- The joint distribution of SRHS and income and the joint distribution of BMI and income "fan out" with age.

- Blacks consistently report lower health status than do whites. Some fraction—but not all—of this difference can be attributed to the lower income of blacks. Less of the difference is explained by income among women than among men, a result that is even more pronounced for BMI.

Pensions and Inequality

Quite a different view of inequality is investigated by Kathleen McGarry and Andrew Davenport. In their paper "Pensions and the Distribution of Wealth" they consider whether pensions increase or decrease inequality in the wealth of families approaching retirement. They find that pensions somewhat increase wealth inequality and find no evidence that pensions crowd out private savings.

Over the past few decades, the financial status of the elderly improved dramatically. Poverty rates for those age 65 or over fell from 25 percent in 1970 to 14 percent in 1994. These gains are attributed in large part to increases in the generosity of the social security program. Yet despite the inclusiveness of the social security system and the progressivity of the benefit schedule, some subgroups of the elderly continue to face disproportionately high risks of poverty. Unmarried women, for example, had a poverty rate of 22 percent in 1994, while the poverty rate for married women was 5 percent. Similarly, the poverty rate for elderly blacks is close to three times that for elderly whites (31 vs. 12 percent).

Social security is just one component of retirement income and given the structure of benefits, differences across individuals in the level of social security wealth are likely to be small in comparison to differences in the other components of total wealth. The large differences in economic well-being within the elderly population therefore stem from differences in the other modes of savings. Recently, much has been written about differences in net worth and savings behavior. Less well studied are differences in pension wealth and the interaction of pensions and individual wealth.

In this paper, McGarry and Davenport take advantage of the HRS to focus on differences in pension wealth for various subgroups of the retirement age population. They ask how pensions affect the distribution of wealth in the population. They compare the distribution of net worth to the distribution of private wealth (net worth plus pension wealth) and to the distribution of total wealth (net worth plus pensions wealth plus social security wealth).

They find evidence that pensions somewhat increase the inequality of wealth across the population. They demonstrate the relationship between total wealth and inequality directly and find that single women in particular fare much worse relative to single men and couples when pension wealth is included in the calculation of total wealth. The paucity of pension holdings among women suggests that their eventual well-being as widows will depend

heavily on the resources left after the death of a spouse. Thus the issue of survivorship benefits for pensions will have important consequences for the eventual poverty rates of widows. While the results are purely descriptive, the authors fail to find evidence to support the notion that pensions crowd out private savings.

References

Eichner, Matthew, Mark McClellan, and David Wise. 1998. Insurance or self-insurance? Variation, persistence, and individual health accounts. In *Inquiries in the economics of aging*, ed. David Wise. Chicago: University of Chicago Press.

Gale, William G., and John Karl Scholz. 1994. IRAs and household saving. *American Economic Review* 84:1233–60.