Volume Title: Themes in the Economics of Aging

Volume Author/Editor: David A. Wise, editor

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-90284-6

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

Publication Date: January 2001

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

Chapter Author: David A. Wise

Chapter URL: http://www.nber.org/chapters/c10318

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

Introduction

David A. Wise

This is the eighth in a series of volumes on the economics of aging. The previous volumes 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, Inquiries in the Economics of Aging,* and *Frontiers in the Economics of Aging.* The papers in this volume discuss important implications of private and (potential) public personal retirement plans, discuss aspects of the health and wealth relationship, consider several aspects of health care in the United States, analyze the retirement effects of social security provisions in the United States and Germany, and consider new evidence on bequests and dissaving at older ages. The papers are summarized in this introduction, which draws heavily on the authors' own summaries.

Personal Retirement Plans

Three papers direct attention to personal retirement plans. The first considers private 401(k) plans and the accumulation of retirement assets. The second paper considers the implications of personal retirement saving plans as part of a possible Social Security reform. The third considers annuitization, which will become increasingly important with the rapid expansion of private personal retirement saving plans.

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 of health and retirement programs at the National Bureau of Economic Research.

Private Personal Accounts: Preretirement Cashouts and 401(k) Assets

About half of U.S. families are now eligible for a 401(k) plan, and the use of these plans is spreading rapidly. In the previous volume of this series, James Poterba, Steven Venti, and I considered the "Implications of Rising Personal Retirement Saving." We concluded that the cohorts that reach age sixty-five between 2024 and 2034 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. In this volume we consider "Preretirement Cashouts and Foregone Retirement Saving: Implications for 401(k) Asset Accumulation." Although many analysts have emphasized the prevalence of cashouts and their presumed effect on the accumulation of retirement assets, we find that the importance of cashouts has been greatly exaggerated and that, in fact, cashouts have only a minor effect on retirement asset accumulation.

The way households support themselves in retirement is changing rapidly. Historically, households in the United States have relied on a combination of Social Security, employer-provided defined benefit pensions, and personal saving to support their retirement years. In the last fifteen years, however, retirement saving programs such as 401(k) plans have become an increasingly common component of household retirement planning. Today, more than 35 million workers participate in 401(k) saving plans, and the annual contribution flow to these plans exceeds $100 billion. The tax-deferred nature of wealth accumulation in 401(k)-type plans, coupled with often generous employer matching contributions that enhance the value of employee contributions, make these plans a powerful vehicle for accumulating retirement wealth.

In Poterba, Venti, and Wise (PVW) (1998a), we showed that even with conservative assumptions about the future growth of 401(k) contributions, the average 401(k) balance for households reaching retirement between 2024 and 2034 would likely exceed average actuarial present value of Social Security benefits.

Although 401(k) plan accumulations are likely to account for a very substantial share of the net worth of future retirees, unlike Social Security benefits, they can be affected by a number of individual decisions. Individuals who work at firms that offer 401(k) plans must decide whether to participate in their employers' plans. Those who do not participate forego the opportunity to accumulate retirement wealth in this tax-deferred form. Conditional on participating, individuals must decide how much of their earnings to contribute to these plans. In particular, participants who leave their jobs can also choose to leave their 401(k) accumulations in their former employers' plans, or to roll over their assets either into an individual retirement arrangement (IRA) or into the 401(k) plan of a new employer. The flexibility afforded by these three options enhances the porta-

bility of 401(k) benefits. It reduces the risk, not uncommon in defined benefit pension plans, of forfeiting pension benefits as a result of job change. However, the flexibility associated with the 401(k) withdrawal option raises the possibility that 401(k) participants may draw down their account balances before retirement, and thereby reach retirement without assets in a 401(k) account.

A number of recent studies have noted that conditional on choosing to withdraw assets from the 401(k) system—i.e., conditional on receiving a "lump-sum distribution"—many individuals use their withdrawals in ways that do not preserve retirement saving. In PVW (1998b) we showed, however, that older workers, and those who receive larger lump-sum distributions, are much more likely to preserve the retirement benefits of their lump-sum distributions through IRA rollovers or other forms of saving.

In the current paper, we use data on past 401(k) participation rates by age and income decile, along with information on average 401(k) contribution rates, to project the future 401(k) contribution trajectories of households that are currently headed by individuals between the ages of twenty-nine and thirty-nine. We allow for the possibility of preretirement withdrawal of 401(k) assets when individuals experience employment transitions. By combining data from the Health and Retirement Survey (HRS) on the likelihood of cashing out a 401(k) account conditional on a job change with data from other sources on the probability of job change, it is possible to estimate the prospective preretirement "leakage" from 401(k) accounts. We confirm that for households reaching retirement age between 2024 and 2034, 401(k) balances are likely to be a much more important factor in financial preparation for retirement than they are today. We estimate that average 401(k) balances in 2024 will be between five and ten times the size they are today, and would represent 50 to 200 percent of a retiree's Social Security wealth (depending on investment allocation and based on current Social Security provisions). For persons retiring in 2034 we estimate that 401(k) balances will be 75 to 250 percent of their Social Security wealth. Moreover, we find that preretirement withdrawals have a small effect on the balance in 401(k) accounts. We estimate that these withdrawals typically reduce average 401(k) assets at age sixty-five by about 5 percent. This is largely because most households whose members are eligible for lump-sum distributions when they change jobs choose to keep their accumulated 401(k) assets in the retirement saving system. These households either leave their assets in their previous employers' 401(k) plans, or they roll the assets over to another retirement saving account, such as a new 401(k) or an IRA. Most of those who do withdraw assets have very small accumulated balances. By comparison, the expense ratio charged by the financial institutions administering 401(k) accounts has a larger effect on retirement resources than does the possibility of preretirement withdrawal.

We also evaluated the possible variation in average returns due to variation in market returns over many years. To do this we made projections based on random draws of returns from the empirical distribution of historical returns. The results show that the *median* 401(k) balance at retirement, especially when a substantial share of the 401(k) portfolio is invested in equities, is below the *mean.* In the case of a 50–50 bond-stock portfolio, for example, the mean 401(k) and rollover balance is $98,800, while the median value is $94,600. The mean in this case lies between the 50th and 60th percentiles of the distribution. For the all-stock case, the mean is between the 60th and 70th percentiles of the distribution of realized outcomes. The most appropriate single measure is unclear. The results also draw attention to the great differences between the bond and stock distributions. For example, 95 percent of bond returns are below $85,800, but only slightly more than 20 percent of stock returns are below this level.

We plan further work in the future on random asset returns and the growth of 401(k) balances. The results above, however, make clear the wide variation in potential system-wide returns, especially stock market returns.

Public Personal Accounts: Market Outcomes and Risk

In the past several years there has been a great deal of analysis of the implications of Social Security reform, in particular of the incorporation of personal retirement accounts into the Social Security system. Although the vast recent literature on personal Social Security accounts makes virtually no reference to the rapidly expanding and overriding importance of private personal accounts, both public and private accounts have many features in common, including market risk. In "The Personal Security Account 2000 Plan, Market Outcomes, and Risk," Sylvester J. Schieber and John B. Shoven consider the market risk implications of the plan. They conclude that an individual who chooses to invest personal accounts in equities would run only a small risk of accruing benefits lower than those provided by the current Social Security system.

In early 1997, the 1994–96 Advisory Council on Social Security released its final report, which remarkably altered the nature of the debate in the United States about the reform of our national retirement system. It did so by giving legitimacy to recommendations that some element of Social Security reform should include individual accounts held by workers. The majority of the council's members actually advocated such reform. To be sure, there had been other people and groups who previously had advocated these types of Social Security reform in this country—but never before had a group of individuals assembled under an official charter by a presidential administration come close to such a recommendation. Since the council's report was released there have been several serious proposals put forward for reforming Social Security that include some element of

individual accounts. There have also been numerous criticisms of this approach to Social Security reform.

In their paper, Schieber and Shoven present a framework for assessing Social Security reform proposals by evaluating a specific reform plan. This plan is one derived from the original Personal Security Account (PSA) developed by the 1994–96 Advisory Council on Social Security (1997). This plan has been dubbed PSA 2000 and its full elaboration is presented in Schieber and Shoven (1999). In part, the PSA 2000 plan was developed to respond to some of the criticisms of the original PSA plan.

The authors' key finding is that the PSA 2000 plan allows risk-averse individuals to retain benefits at least as high as current-law benefits. Those who choose to take the risks inherent in stocks bear some chance of having to live on lower than current-law benefits in retirement. These risks are modest, however, and the poor are significantly protected by the presence of the tier-one benefits.

The authors describe a particular partial privatization plan that relies more heavily on individual accounts than do most other proposals, and evaluate its overall actuarial soundness and the outcomes that individuals would face if it were adopted. The plan passes the actuarial soundness test and would permit individuals to enjoy safe benefits approximately equal to current-law benefits if U.S. government inflation-indexed bonds were offered and invested in. If participants invested their tier-two accounts in common stocks, they would face a small probability of having significantly less in retirement than current-law benefits. However, these risks are reduced by the presence of the flat tier-one benefits. This first tier is relatively more important for low-income households, who would enjoy benefits at least as great as current benefits with a high degree of certainty.

It is not surprising that the PSA 2000 plan performs well with respect to the principles set out by the authors. The first-tier defined benefits feature provides an important safety net against poor investment returns and permits the retention of the basic progressive structure of the current program (Principle 1). A primary feature of the program is the mandatory contribution of 2.5 percent of covered payroll. Although these additional contributions would be partially offset by the actions of individuals, there would be a significant net increase in national saving (Principle 2). The disability and early survivor programs would be retained, and if the proportion of the projected federal government surpluses suggested by President Clinton were allocated to the program, there would be enough money to cover the long-run deficit of the disability program (Principle 3). Under all of the scenarios examined, the PSA 2000 plan would be in balance or surplus after seventy-five years and would offer the prospects of payroll tax reductions (Principle 4). Most retiring couples would be treated as two single individuals, thereby improving the equity between these participant classes (Principle 5). The tier-two contributions and payouts would be directly

connected. In fact, the tax element of payroll deductions would be reduced by the 2.5 percent rebate in the form of a 1:1 match of tier-two contributions (Principle 6).

The authors examine the risks borne by individuals and judge them to be tolerable. In particular, the amount of risk one would bear would be a matter of personal choice. Furthermore, the risks are least for low-income households (Principle 7). The PSA 2000 plan has relatively low administrative cost partly because it has relatively large (5 percent) individual accounts (Principle 8). The authors do not advocate in this paper the particulars of the PSA 2000 plan; they do, however, advocate that the riskiness of all serious proposals be evaluated in a manner similar to the procedure they follow in their paper. The authors emphasize that what is "heartening" about the findings is that a plan that relies heavily on individual accounts can still be relatively safe for individual participants.

Annuitization and Retirement Benefits

Mandatory annuitization is an important feature of the current Social Security system and will pose an ever more important question with respect to the increasingly common and likely future accumulation of large private retirement accounts. In "Are the Elderly Really Over-Annuitized? New Evidence on Life Insurance and Bequests," Jeffrey R. Brown focuses on this issue. He does so by considering whether individuals with bequest motives purchase life insurance to offset mandatory annuitization of Social Security benefits. He concludes that the answer is no, and that the evidence is not sufficient to argue against mandatory annuitization of current Social Security benefits.

It is well established in the economics literature that annuities ought to be of substantial value to life-cycle consumers, who face an uncertain dates of death. Buying a life insurance contract is analogous to selling an annuity. It is generally viewed as an appropriate product for working-age individuals who seek to protect their families against the loss of future labor earning. However, it appears to serve little purpose in the portfolio of a retired life-cycle consumer who is concerned only with self-financing retirement out of accumulated wealth. With no labor earnings to insure, an elderly individual should be purchasing annuities in order to provide a certain consumption stream in retirement, not selling annuities through the purchase of life insurance. Even if the individual wishes to leave a portion of wealth to heirs in the form of gifts or bequests, this can be achieved by investing this portion of wealth in ordinary bonds or other nonannuitized assets. In fact, if life insurance premiums are higher than actuarially fair, holding riskless bonds would strictly dominate life insurance as a form of wealth transfer.

Yet elderly households in the United States overwhelmingly hold life insurance, while only a small fraction hold privately purchased annuity

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

One suggestion is that life insurance is being held by elder households to offset an excessive level of mandated annuitization in the form of Social Security. To the extent that this "annuity offset model" is true, it has at least two important implications. First, it would be indicative of very strong bequest motives, which constitute an issue of perennial controversy in the economics literature. Second, if individuals are over-annuitized due to these strong bequest motives, it would indicate a potential welfare gain from lessening the extent of mandated annuitization.

This paper reexamines the annuity offset model using more recent and better data than have previously been available. The paper presents substantial evidence that the reason the elderly hold life insurance is not to offset mandated annuitization in the form of Social Security in order to leave a bequest. This finding is relevant to the current debate over the future of the Social Security system because it bears upon the question of whether mandatory annuitization is desirable. Were it the case that a substantial fraction of elderly households were over-annuitized by the existing Social Security system due to the existence of strong bequest motives, it would be evidence in favor of allowing choice over the annuitization decision. The results of this paper suggest that households are not over-annuitized by Social Security for bequest reasons. Therefore, Brown concludes, the simple fact that many elderly households own term life insurance is not a sufficient argument against mandatory annuitization of retirement resources. This finding is consistent with the idea that annuities are of substantial value in the retirement portfolios of elderly individuals. As a result, mandatory annuitization may be desirable to overcome adverse selection in the annuity market. However, Brown emphasizes, this conclusion should be tempered by the acknowledgment that individuals may be over-annuitized for reasons other than bequest motives.

Wealth and Health

The relationship between wealth and health has received increasing attention in recent years. A key issue has been the direction of causality:

Does wealth affect health, or does health determine wealth? Two papers in this volume address this relationship, but from different perspectives. The first focuses on the relationship between income and health; the second focuses on the predictors of mortality.

Income and Health

Individuals for whom family income was less than $5,000 in 1980 could expect to live about 25 percent fewer years than people whose family income was greater than $50,000. This finding is explored by Angus Deaton and Christina Paxson in their paper on "Mortality, Education, Income, and Inequality among American Cohorts." A key finding is that greater income reduces the risk of death, even after controlling for education.

The strong relationship between mortality and socioeconomic status (SES) has been a major concern in demography, epidemiology, and public health for many years, and is beginning to attract the attention of economists. The concept of SES is used more widely outside the field of economics than within it and one of the issues that remains to be settled is the extent to which these differences are caused by income or by other factors correlated with income (education being the most obvious). Many writers believe that there is at least some direct protective effect of income, and in a recent body of literature much identified with the work of Richard Wilkinson, it is argued that, while higher income is protective (at least at the individual level), *income inequality* is a health hazard that raises mortality, if not at the individual level at least in the general population or larger subpopulations. Wilkinson postulates that *inequality itself* is a health hazard and that it is less healthy both for rich and for poor to live in a more unequal society. It is hardly necessary to emphasize the importance of such a link, if it indeed exists. The proponents of some changes (e.g., improvement in school quality, or raising the return on Social Security) make a plausible case that such changes will make everyone better off, although some will be more so than others. If such changes increase inequality, as almost certainly they would, the cost of lives lost would have to be offset against the economic benefits.

The authors' main purpose in this paper is not to try to come to judgment based on the review of the evidence, but to offer some new evidence based on income, income inequality, and mortality data for birth cohorts of Americans observed over the two decades from 1975 to 1995. It seems that birth cohort data—as opposed to individual data, state data, or country data—have not previously been used in this context; and unlike other sources, birth cohorts offer both a cross-sectional and a time series dimension to the same data. The model developed in the paper is designed as a framework for empirical application, and provides a way of thinking about the effects of income and income inequality in a context in which, although causality runs from income to health, it is not absolute income that matters for health but income relative to that average of an (unobservable) refer-

ence group. Although inequality has no direct effect on health, the fact that reference groups are not observed means that the slope of the observed relationship between health and income varies with the ratio of between- to within-group inequality. The model can be readily extended readily to incorporate a direct effect of inequality by making health depend on the absolute size of income differences within the reference group; but equally plausible specifications give different results so that, according to the theory, income inequality can be either protective or hazardous. Deaton and Paxson give detailed consideration to the aggregation of the relationship between health and income, and to how it can be expected to change as it is examined with different sources of data, such as individual records, averages of states or countries, or averages of birth cohorts. The authors document the strongly protective effects of income and examine how those effects vary at different points in the life cycle. As to inequality, the authors fail to find not only that it increases the risk of mortality, but that there is actually a protective effect, in apparent contradiction with not only the Wilkinson hypothesis but with much of the theory developed in this paper.

Deaton and Paxson also give a good deal of attention to the role of education: whether income is a mask for education, how income and education affect mortality in the cross-section and over time, and whether the treatment of income and education affects our results on the role of inequality. In a cross-section of birth cohorts, income and education are closely correlated so that, in order to disentangle their effects, the authors rely on the time series dimension of the cohort data, supplemented by individual-level data from the National Longitudinal Mortality Study (NLMS). The individual-level data show that both income and education are separately protective against mortality and that only some of the effect of income is removed when we attempt to allow for reverse causality from nearing death to income. In the cohort data, by contrast, income appears to increase the risk of mortality conditional on education, a result that the authors tentatively ascribe to the short-run or business-cycle effects of income on mortality. In concluding, the authors emphasize:

> Our original purpose was to use birth-cohort data to examine the links between mortality and inequality. Controlling for income, we find that higher inequality is associated with lower mortality, a conclusion that comes from negative association of mortality and inequality in the United States in the late 1970s and early 1980s. While it is possible that such a result has some real basis—and there are theoretical mechanisms that could produce it—it is hardly established by these results. In particular, the sign of the effect is implausible, if only because of the expected operation of Jensen's inequality, and the magnitude of the effect is quite sensitive to the way in which other variables are introduced, particularly income and education. Indeed, we suspect that the current priority should not be the investigation of the effects of inequality, but the un-

> packing of "socioeconomic status" into its components . . . so as to allow them to respond to income and education in different ways. The results reported here make it clear that this is no easy task; the way in which education and income affect mortality is not the same for men as for women, nor for young adults as for older adults; it is different over long time periods and over the business cycle, and it is different in the cross-section from over time. We find evidence that short-term increases of income may raise the risk of mortality, particularly for young men. In the cohort data, however, the longer-term effects of income, or of income linked to education, are protective. Yet this evidence needs to be reconciled with the individual-level data from the follow-up studies, which show that, especially for men, income plays a role as large as or larger than that of education. Work on these issues has hardly begun (162).

Socioeconomic Status and Death Rates

In contrast to Deaton and Paxson, Michael D. Hurd, Daniel McFadden, and Angela Merrill conclude in their paper on "Predictors of Mortality among the Elderly" that their findings are consistent with the view that the primary cause of the relationship between wealth and health is unobserved individual characteristics that cause both early death on the one hand and lower wealth and less education on the other.

The authors point out that differential mortality by (SES) has been observed over a wide range of data on population, but because of data limitations, the measures of SES have typically been occupation or education. In the Health and Retirement Survey (HRS) and the AHEAD survey there is scope for expanded studies of differential mortality (the mortality gradient), because these panel surveys follow a large number of older persons and obtain extensive data on income, wealth, and health conditions as well as on occupation and education. Furthermore, the fact that the AHEAD population is almost completely retired means that a very strong confounding effect of health on income via work status is eliminated.

The goal of the paper is to study the predictors of mortality in the AHEAD population. The authors focus on the mortality gradient as a function of wealth. They emphasize that this gradient is important because it causes difficulty in understanding life-cycle behavior from cross-sectional variation in wealth: Besides cohort effects that would, by themselves, cause wealth to decline with age in cross-section, the mortality gradient will cause wealth to increase both in cross-section and in panel. As a cohort ages, those with less wealth die, leaving survivors from the upper part of the wealth distribution. Thus, even if no couple or single person dissaved after retirement, the wealth of the cohort would increase with age. This makes it difficult to study life cycle wealth paths based on synthetic cohorts, which will eliminate cohort differences in lifetime resources but not differential mortality. These difficulties carry over to studies of income and consumption in synthetic cohorts.

A second focus of this paper is the determinants and predictive power of the subjective probability of survival. Respondents were asked to give an estimate of their survival chances to a target age, which was approximately twelve years in the future. In the HRS this variable has been shown to be a significant predictor of mortality between waves 1 and 2. A goal of this paper is to determine whether it has similar predictive power over the AHEAD population.

The authors conclude that the mortality gradient, whether a function of wealth, income, or education, apparently decreases with age. The authors say that any explanation at this point would be rather speculative, but the finding is consistent with the view that the primary cause of the gradient is unobserved individual characteristics that cause both bad health (and therefore early death) and lower earnings (and therefore lower wealth and less education). Were the causality to run primarily from economic resources to health and mortality, the authors argue, we should see a persistent difference in mortality outcomes in very old age between those with substantial resources and those with few. The authors say they do not see this, but they also say this should be confirmed by further analysis. If the differential is due to unobserved individual differences, the mortality gradient operating at younger ages will have truncated the distribution, so that in extreme old age the variation in individual characteristics would be greatly reduced. Therefore, classifying people by SES would not produce any substantial differences in mortality.

The subjective survival probability predicts actual mortality as in the HRS, and the authors say this should increase our confidence that it can be used to construct individualized lifetables for models of life cycle saving behavior as proposed by Hurd, McFadden, and Gan (1998). Whether such lifetables will have substantial explanatory power for saving remains to be determined as more waves of AHEAD become available.

Health Care

Three papers consider different aspects of health care in the United States. Medicare spending in real terms has doubled in the past two decades, even though the health of the older population has improved; two papers aim to reconcile this apparent contradiction. A third paper develops a method for understanding the reasons for the wide differences in health care expenditures among firms in the United States.

End-of-Life Spending, Declining Mortality, and Potential Health Care Saving

Lower mortality, and thus better health at a given age, might reduce the cost of medical care; but this potential reduction can be offset by increasing expenditure given health status. In their paper on "Trends in Medicare

Spending Near the End of Life," Jeffrey Geppert and Mark McClellan find that increasing intensity of health care for both survivors and decedents has far outweighed reductions in cost due to decreasing mortality—and thus the continuing rise in health care expenditures.

Recent decades have witnessed dramatic improvements in health at older ages, including reductions in both mortality and morbidity. Although real growth in health care costs has accompanied improvements in health for the past fifty years, improvements in health give hope that avoided medical utilization due to better health may lower health care costs, or at least significantly reduce the rate of growth. This might occur as a result of mortality improvements, because decedents on average have much higher health care costs than survivors, because average spending in the last year of life declines with age, and because people are dying at older ages. Spending in the last year of life accounts for nearly 30 percent of total Medicare program payments, so a reduction in average spending per decedent might significantly influence total Medicare outlays. A similar argument applies to the decline in the prevalence of chronic disability. Elderly patients with chronic disabilities cost more than the nondisabled elderly, so reducing disability prevalence also holds the promise of lowering Medicare cost.

However, the importance of such program savings from the shift of beneficiaries to lower-cost, more healthy states over time also depends on the changes in expenditures, given health status, that occur at the same time. Expenditures per capita are rising, and much evidence suggests that rising intensity of treatment is the principal cause. Changes in expenditures given treatment may be responsible for some of the improvements in health, and in any event changes in health and in expenditures will continue to occur together. Despite these facts, prior studies generally have not considered both factors jointly, to determine whether the savings that result from improving health over time are likely to have a quantitative impact on expenditure growth. In this paper, Geppert and McClellan determine the importance of changes in Medicare cost that resulted from declines in age-specific mortality between 1988 and 1995, and calculate how much Medicare expenditure trends would have differed in the absence of the mortality improvement.

The authors reach several conclusions. First, despite the fact that mortality has continued to fall, the share of Medicare program payments accounted for by persons in the last year of life has remained relatively constant, declining only slightly over the past two decades.

Second, the rate of spending growth was similar for survivors and decedents—actually, slightly larger for the oldest male survivors than for other demographic groups—as a result of relatively greater growth in spending for nonacute services. Thus, spending growth for survivors continues to account for most of the growth in Medicare cost.

Third, for both decedents and survivors, the composition of spending growth has changed in recent years. In the most recent decade, growth in spending for nonacute services has accounted for half of overall spending growth. Thus, spending growth for decedents was not primarily the result of increasing the heroic, intensive measures near the end of life. In addition, although greater coverage of nonacute alternatives might be expected to affect end-of-life costs disproportionately, growth rates for nonacute services were even greater for older survivors than for decedents. Large utilization effects for both acute and nonacute services occurred in both groups.

Fourth, although improvements in mortality have been a factor leading to lower Medicare spending over time, this effect has been swamped by the much larger increases in expenditures given survival status for decedents and survivors alike. Without the survival improvements, the authors' estimates suggest that Medicare spending would have grown only 4 percent more than it did.

Thus, the authors conclude that increasing utilization for both survivors and decendents has been a far more important determinant of Medicare spending over time than have improvements in mortality.

The Concentration of Medical Spending

In "The Concentration of Medical Spending: An Update," David M. Cutler and Ellen Meara conclude that the explanation for rising health care cost in the face of declining disability is the increase in the use of "postacute service," in particularly home health care and skilled nursing home care. That is, Cutler and Meara try to point to the specific factors that have led to rising costs, given age, that Geppert and McClellan also emphasize.

Cutler and Meara point out that during the last two decades, the number of Medicare beneficiaries has increased by 50 percent, and Medicare spending per beneficiary has doubled in real terms. These findings are difficult to understand, however, in light of changes in the health of the elderly. Disability rates are falling among the elderly by about 1.5 percent per year. Since the disabled spend much more than the nondisabled on medical care, it seems that in relative (if not absolute) terms, spending on the elderly should be falling over time. The combination of large increases in per-person spending and the reduction in disability leads to a paradoxical situation, in which policy analysts call simultaneously for reforms to control Medicare cost growth (to bring spending for the elderly in line with that for other age groups) and for Medicare to cover currently uncovered services such as prescription drugs (to promote further health improvements).

The goals of this paper are to document how trends in spending by age have changed among elderly Medicare beneficiaries in the last decade, and

to reconcile the decline in disability rates with rapid increases in spending among the elderly. In particular, the authors consider what has happened to age-specific spending since 1987. The authors then attempt to reconcile increased spending with sharply declining disability. In particular, these relate medical spending by age to six factors: demographics, disability, time until death, intensity of treatment, prices, and changes in the nature of care.

Cutler and Meara reach two central conclusions. First, they find that the trend of disproportionate spending growth among the oldest old has continued during the decade between 1985 and 1995. Between 1985 and 1995, spending for the younger elderly (ages sixty-five to sixty-nine) rose by 2 percent annually in real, per-person terms, while spending for the older elderly (ages eighty-five and up) rose by 4 percent. This is similar to the differential increase in spending by age over the 1953–87 period, the authors say.

Second, Cutler and Meara conclude that the reason for the large increase in spending on the oldest elderly in comparison to the younger elderly is the rapid increase in use of postacute services—home health care and skilled nursing care in particular—among the oldest old. People aged eighty-five and older used, on average, $241 in postacute services in 1985 and $1,887 in 1995, a 20 percent annual increase. The younger elderly, in contrast, increased their use of postacute services from $49 to $257, a 15 percent annual increase. Use of acute-care services, in contrast, grew relatively evenly by age, 1.2 percent annually for the younger elderly and 0.7 percent annually for the older elderly.

The increase in the use of postacute service, the authors say, is the explanation for the discrepancy between rising medical spending and falling disability. Lower disability by itself contributes to lower spending than we would otherwise observe; but the increase in use of nontraditional services more than offsets the effects of improved health. The increase in postacute service use is also a major difference between the pre- and post-1987 trends. In earlier work these authors found that rising expenditures on the older population were a result of increased intensity of acute-care services for that age group. In the post-1987 period, intensity changes in acute-care treatments do not account for a substantial discrepancy by age.

The authors suggest that increase in postacute service use may reflect several factors: true increased service use for people who were not receiving care in the past; "gaming" of the Medicare system, whereby providers now use out-of-hospital services instead of in-hospital services; or outright fraud. They are unable to discriminate among these explanations, although they suspect each is important.

The Reasons for Expenditure Differences among Firms

Medical expenditure per employee varies enormously across firms in the United States. In our paper on "The Sources of Cost Difference in Health

Insurance Plans: A Decomposition Analysis," Matthew Eichner, Mark McClellan, and I develop a method to identify and quantify the importance of the factors contributing to the wide differences across employer-provided health plans. Our results suggest that further efforts to understand the differences should focus on the relatively intensive inpatient care, and must address both the variation in admission for intensive treatment and the variation in cost given treatment.

We are engaged in a long-term project to analyze the determinants of cost differences across firms. In particular, we look forward to an estimation that can predict the effect on medical expenditures of specific changes in medical insurance plan provisions. The project is based on insurance claims records from a large number of employers. The vast amount of information in insurance claims records is both a blessing and a curse. A key advantage of claims data is the detail they provide, but the detail also poses a challenge: how best to summarize and convey the information contained in the millions of claims filed each year under a typical employer-provided plan.

Our goal in this paper is to present a method that allows us conveniently to summarize information contained in the claims data. In particular, we want to describe the sources of cost differences across plans. We consider eight plans that vary in average expenditure for those filing claims, from a low of $1,645 to a high of $2,484. We then propose a method to decompose these differences into their component parts. The goal is to quantify the contribution of each component to total cost differences across firms. We believe that this method allows us to point directly to the sources of cost difference and thus will help us to focus subsequent analysis where it is most likely to make a difference.

We first present a statistically consistent method for decomposing the cost differences across plans into component parts due to demographic characteristics of plan participants, the mix of diagnoses for which participants are treated, and the cost of treatment given particular diagnoses. The goal is to quantify the contribution of each component to the difference between average cost and the cost in a given firm. The demographic mix of plan enrollees accounts for wide differences in cost ($649). Perhaps the most noticeable feature of the results is that, after adjusting for demographic mix, the difference in expenditures accounted for by the treatment costs given diagnosis ($807) is almost as wide as the unadjusted range in expenditures ($838). Differences in cost due to the various mixes of illness that are treated, after adjusting for demographic mix, also accounts for large differences in cost ($626). These components of cost do not move together; for example, demographic mix may decrease expenditure under a particular plan, whereas the diagnosis mix may increase costs.

We also provide an approximate decomposition of the "variation" in expenditures across firms. Although outpatient care accounts for almost 50 percent of expenditures on average, it accounts for only about 20 per-

cent of the variation in expenditures across firms. Inpatient care accounts for about 34 percent of expenditures on average, but almost 59 percent of the variation in expenditures. Thus, one can conclude that variations in high-cost inpatient treatments are a principal cause of the substantial cost variation across firms. (A "residual" group accounts for about 16 percent of expenditure and about 20 percent of variation in expenditure across firms.) The most important component of variation is the diagnosis rate, which accounts for about 52 percent of variation across firms. Treatment cost differences, given treatment, account for about 40 percent, with the remainder accounted for by the interaction between the two.

Our results suggest that efforts to understand the substantial differences in private insurance expenditures should focus on the relatively intensive inpatient care, and that further analysis should address the sources of variation in both admissions for intensive treatment (the diagnosis rate) and the cost given treatment.

Social Security Provisions and Retirement

Two papers consider the relationship between retirement and the provisions of social security plans. The first, based on German data, emphasizes a method to estimate the effect of plan provisions when it is uncertain under which of several programs a person could retire. The second paper directs attention to Social Security retirement incentives in the United States.

Uncertain Eligibility and Retirement in Germany

Analyses of the retirement incentive effects of public and private pension plan provisions typically have assumed that the eligibility for a particular plan is known. In many countries, however, this is not the case. There are often several pathways to retirement, with disability and unemployment insurance programs essentially serving as social security early retirement programs. This the case in many European countries, for example. In his paper on "Incentive Effects of Social Security under an Uncertain Disability Option," Axel Börsch-Supan considers methods to estimate retirement incentive effects when the pathways available to an individual are uncertain (to the analyst). He finds that an instrumental variable estimate that uses expected incentive effects in an option value retirement model provides reliable estimates. His judgment is confirmed in the discussion of the paper by Daniel McFadden.

In most industrialized countries, old-age labor force participation declined dramatically during the last decades. Together with population aging, this puts the pay-as-you-go social security systems of the industrialized countries under a double threat: Retirees receive pensions for more years while there are fewer workers per retiree to shoulder the financial

burden of the pension system. The decline of old-age labor force participation has turned attention to the incentive effect of social security systems: Is a significant part of the threat homemade because pension systems provide overly strong incentives to retire early? This "pull" view—that labor supply has declined because early retirement provisions pull old workers out of employment—is in contrast to the "push" view—that a secularly declining demand for labor has created unemployment, which pushes older workers into early retirement.

The pull view is prominently advanced in a recent volume edited by Gruber and Wise (1999). The authors from eleven countries argue that the declining old-age labor force is strongly correlated with the incentives created by generous early retirement provisions. Formal econometric analyses (Stock and Wise 1990a for the United States; Meghir and Whitehouse 1997 for the United Kingdom; and Börsch-Supan 1992, 2000 for Germany) find strong incentive effects of public and private pension rules.

Incentive effects of pension rules are usually estimated under the assumption that the institutional environment provides a single pathway for retirement age, or an opinion value of postponing retirement at any prospective retirement age. However, most countries provide competing pathways that include several early-retirement options in addition to normal retirement, typically at age sixty-five. In particular, this is true in Germany, where early retirement due to a "disability" before age sixty has been the most common pathway to retirement.

Social Security Incentives for Retirement in the United States

In their paper on "Social Security Incentives for Retirement," Courtney Coile and Jonathan Gruber direct attention to the incentives for retirement by considering the distribution of Social Security wealth accrual at different ages for participants in the HRS. Although Social Security benefits are intended to be actuarially fair between ages sixty-two and sixty-five, Coile and Gruber show that there is enormous variation in accrual depending on individual circumstances. For example, for about one-third of survey participants, accrual is negative between ages sixty-two and sixty-three, so that for these people there is a penalty on work at age sixty-two.

One of the most striking labor force phenomena of the second half of the twentieth century has been the rapid decline in the labor force participation rate of older men. In 1950, for example, 81 percent of sixty-two-year-old men were in the labor force; by 1995, this figure had fallen to 51 percent.

Much has been written about the proximate causes of this trend, in particular about the role of the Social Security program. A large number of articles have documented pronounced spikes in retirement at ages sixty-two and sixty-five, which correspond to the early and normal retirement ages for Social Security, respectively. Although there are some other expla-

nations for a spike at age sixty-five, such as entitlement to health insurance under the Medicare program or rounding error in surveys, there is little reason for a spike at sixty-two other than the Social Security program. Indeed, this spike at age sixty-two emerged only after the early-retirement eligibility age for men was introduced in 1961.

The presence of these strong patterns in retirement data suggest that Social Security is playing a critical role in determining retirement decisions. In order to model the impact of Social Security reform on retirement behavior, however, it is critical to understand what this role is. The evidence of a spike at age sixty-two, for example, is consistent with at least three alternative hypotheses. The first is that there is an actuarial unfairness built into the system penalizing work past age sixty-two, so that there is a "tax" effect that leads workers to leave the labor force at that age. The second is that workers are liquidity constrained: They would like to retire before age sixty-two, but cannot because they are unable to borrow against their Social Security benefits and have no other sources of retirement support. In this case, there will be a large exit at age sixty-two as benefits first become available. The third is that workers are myopic or information constrained: They either do not understand or do not appreciate the actuarial incentives for additional work past age sixty-two, so they retire as soon as benefits become available.

The Coile-Gruber paper is an investigation of the tax effect along four dimensions. First, the authors assess whether the tax rate Diamond and Gruber (1999) compute (using a synthetic individual with annual earnings at the median of his cohort) is similar to the tax rate of the real median person. We might expect a difference, as the shape of the earnings history is a significant determinant of Social Security incentives—through the dropout years provision—and this is not appropriately reflected with a synthetic earnings history. Second, the authors assess the distribution of retirement incentives across the population. Even if there is no significant disincentive for the typical worker, disincentives for a large subset of workers could still be associated with a spike in the aggregate retirement data. Third, they assess the importance of considering incentives for retirement in the next year versus incentives for retirement over all future years, drawing on the insights of the option value model of Stock and Wise (1990a,b). Finally, they incorporate the role of private pensions, an important determinant of retirement for a large share of workers.

The strategy is to consider a set of real individuals, the older persons surveyed in the HRS. These data allow the authors to compute carefully the retirement incentives from Social Security and pensions, both for the median individual and across all survey participants.

The authors confirm that there is, in fact, a small subsidy to work at ages sixty-two to sixty-four at the median. However, they also show that there is substantial heterogeneity across persons. There is a net tax on

work at age sixty-two for about one-third of their sample, which is consistent with a spike in the hazard rate at age sixty-two.

Bequests and Dissaving

Wealth accumulation before retirement and decumulation after retirement have been the subjects of longstanding and intensive study by economists. In their study of "Anticipated and Actual Bequests," Michael D. Hurd and James P. Smith use new data from the HRS to deepen our understanding of the role of bequests in household decisions about these processes.

Important advances have recently been made in documenting the process of wealth accumulation by households. New data have increased our knowledge about the facts surrounding the distribution of household wealth and, to a lesser extent, household saving behavior. However, this improved factual base has not yet been translated into a deeper understanding about the theoretical reasons people save. The candidates remain much the same: life cycle timing, risk aversion, and bequests. Understanding bequest motives has been particularly difficult, in part due to the inherent difficulties in measuring anticipated and actual bequests.

In their paper, Hurd and Smith study the role of inheritances and bequests in shaping household decisions about wealth accumulation. They study bequests by using new methods of measuring anticipated and actual bequests. They examine actual bequests made by deceased individuals and compare them with their previously stated bequest intentions. Using panel data with two measurements of subjective bequest probabilities, they explore the reasons an individual might revise his or her bequest expectations. Among other things, these reasons may include new information on health or economic conditions of household members. Their results are based on wealth, anticipated bequests, and actual bequests from two waves each of the HRS and AHEAD. Because the paper uses two new types of data, considerable attention is directed to validating these new data.

Actual bequests are measured in exit interviews given by proxy respondents for 774 AHEAD respondents who died between waves 1 and 2 of the AHEAD survey. Among other things, these exit interviews provide data about the medical and nonmedical costs associated with the illnesses of the deceased respondents and the value and distribution of the estates. Even though the deceased were quite ill before they died, medical expenses did not cause substantial reduction in their estates. Because the exit interview obtained estate information that is representative of the population, the distribution of these estate values is quite different than one would suppose from estate records, which are obtained for only a wealthy subset of the population.

Anticipated bequests were measured in two waves of HRS and AHEAD by the subjective probability of leaving bequests. The authors study the reasons for between-wave revisions of the subjective bequest probabilities. They find that increases in the subjective probability of surviving, increments in household wealth, and widowing were all associated with increases in bequest probabilities, whereas out-of-pocket medical expenses reduced the likelihood of a bequest. By comparing bequest probabilities with baseline wealth, Hurd and Smith are able to test a main prediction of the life cycle model: that individuals dissave at advanced old age. The authors conclude that respondents anticipate substantial dissaving before they die.

References

Advisory Council on Social Security. 1997. *Report of the 1994–96 Advisory Council on Social Security.* Vol. 1, *Findings and recommendations.* Washington, D.C.: GPO.

Börsch-Supan, A. 1992. Population aging, social security design, and early retirement. *Journal of Institutional and Theoretical Economics* 148:533–67.

———. 2000. Incentive effects of social security on labor force participation: Evidence in Germany and across Europe. *Journal of Public Economics* 78:25–49.

Diamond, P., and J. Gruber. 1999. Social Security and retirement. NBER Working Paper no. 7830. Cambridge, Mass.: National Bureau of Economic Research.

Gruber, J., and D. A. Wise. 1999. *Social Security and Retirement around the World.* Chicago: University of Chicago Press.

Hurd, M. D., D. McFadden, and L. Gan. 1998. Subjective survival curves and life cycle behavior. In *Inquiries in the economics of aging,* ed. D. A. Wise. Chicago: University of Chicago Press.

Life Insurance Market Research Association (LIMRA) International. 1993. Profiles in coverage: The widening gap in U.S. life insurance ownership. Windsor, Conn.: LIMRA.

Meghir, C., and E. Whitehouse. 1997. Labour market transitions and retirement of men in the UK. *Journal of Econometrics* 79:327–54.

Poterba, J. M., S. F. Venti, and D. A. Wise. 1998a. Implications of rising personal retirement saving. In *Frontiers in the economics of aging,* ed. D. A. Wise, 125–67. Chicago: University of Chicago Press.

———. 1998b. Lump sum distributions from retirement saving plans: Receipt and utilization. In *Inquiries in the economics of aging,* ed. D. A. Wise, 85–105. Chicago: University of Chicago Press.

Schieber, J., and J. B. Shoven. 1999. *The real deal: The history and future of Social Security.* New Haven, Conn.: Yale University Press.

Stock, J. H., and D. A. Wise. 1990a. The pension inducement to retire: An option value analysis. In *Issues in the economics of aging,* ed. D. A. Wise, 205–30. Chicago: University of Chicago Press.

———. 1990b. Pensions, the option value of work, and retirement. *Econometrica* 58 (5): 1151–80.