Volume Title: Explorations in the Economics of Aging

Volume Author/Editor: David A. Wise, editor

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-90337-0
ISBN13: 978-0-226-90337-8

Volume URL: http://www.nber.org/books/wise09-2

Conference Date: May 2009

Publication Date: March 2011

Chapter Title: Introduction

Chapter Authors: David A. Wise and Richard Woodbury

Chapter URL: http://www.nber.org/chapters/c11930

Chapter pages in book: (1 - 19)

Introduction

David A. Wise and Richard Woodbury

The next twenty years mark a new phase in the demographic transition of the United States, as the baby boom generation becomes eligible for Social Security and Medicare. This large population mass, which has been in a prime working and earning phase of their careers, and supporting a comparatively smaller population of older retirees, reaches an age when they too may retire. The social and economic transition inherent in this demographic shift will be shaped to a significant extent by the labor market decisions of this population, how much they will have saved for their retirement, the cost of their health care, and the future evolution of retirement and health policy.

Research in the economics of aging is not limited to the United States, or to changing age demographics, or to studies of the elderly. The field covers aging issues around the world, in both developed and developing countries. It involves research on both health and economic circumstances, for both individuals and populations, and with particular focus on how health and economic circumstances evolve interactively over the life course. It encompasses advances in research methodology, data resources, experimental interventions, and the evolution of public policy in health, work, disability, and retirement.

This is the thirteenth in a series of NBER volumes on the economics of aging. The goal of this series is to present studies that are at the forefront of economics of aging research along multiple dimensions. The topics addressed in each volume have evolved with the science, though some core

David A. Wise is the John F. Stambaugh Professor of Political Economy at the Kennedy School of Government, Harvard University, and director of the program on aging at the National Bureau of Economic Research. Richard Woodbury is a senior administrator with the program on aging at the National Bureau of Economic Research.

themes have persisted. The volumes have generally included work on the evolving financial circumstances of individuals as they age, and the factors that influence financial well-being. They have included work on health and disability, medical practice patterns, health care costs, health care financing, and health policy. They have included work on aging around the world, including an increased focus on aging in developing countries. They have included work on the relationships between health and economic circumstances. And they have included work on research methodology, database development and, most recently, on what can be learned from experimental interventions. Through thirteen volumes, the large majority of this research has been funded by the National Institute on Aging, which has made a long-term commitment to advancing the economics of aging field.

The previous volumes are *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, Frontiers in the Economics of Aging, Themes in the Economics of Aging, Perspectives on the Economics of Aging, Analyses in the Economics of Aging, Developments in the Economics of Aging,* and *Research Findings in the Economics of Aging.* This introduction provides an overview of the studies contained in the volume, relying to a significant extent on the authors' own language to summarize their findings.

The volume is organized into four substantive topic areas, not dissimilar from past volumes, though with a continuing evolution of emphasis. The first section contains studies on retirement saving in the private sector. The second deals with trends in health care costs and coverage. The third focuses on relationships between socioeconomic circumstances and health. The fourth looks at issues in aging in less developed countries. There are many issues that cut across this organization. For example, in addition to the health issues addressed in the second and third section, chapter 1 in the first section draws out the important implications of poor health for the post-retirement evolution of assets, and chapter 10 in the fourth section considers a way to fight anemia in India. As noted, the studies are all components of longer-term research themes of the NBER program on aging, and an attempt is made to place these new studies in the context of our larger agenda.

Each chapter also includes comments provided by a discussant. These comments add a depth of perspective on each research topic. In some cases, the discussant comments put the primary study into a larger context. In some cases, they are critical commentary. And in other cases, they are expansions of either the theoretical underpinnings or empirical findings that are reported in the primary studies. The result is a richer treatment of each topic addressed. Because the volume focuses on studies that are at the forefront of economics of aging research, they are by their nature more exploratory or innovative. The discussant comments provide a certain grounding or

breadth of perspective that is particularly valuable in assessing these more exploratory and innovative research directions.

Retirement Saving

Each of the thirteen volumes in the NBER's Economics of Aging series has looked at a new dimension in the evolving financial circumstances of individuals as they age. This theme has persisted, first because of the major changes over twenty-five years in how individuals finance their retirement. Studies have documented the changing composition of financial resources available at older ages, projected future changes, and followed cohorts of individuals affected by these changes. The "financial circumstances" theme has also persisted because of a realization that people's behavior is so dramatically influenced by the design of public and private programs and policies. This has opened up an entirely new area of research in understanding how the determinants of financial decision making, as well as behavior, change research—how to induce financial decision making that leads to better outcomes.

Financial support at older ages comes from a combination of Social Security, traditional defined-benefit pension benefits offered by employers, retirement saving, and income from work. The landscape and relative weight of these components has changed dramatically over the past two decades, and will likely change again in the decades ahead. The major change of the last two decades has been the growth in retirement saving programs, and particularly 401(k) plans, and a parallel decline in traditional defined-benefit pension programs. The trends to date suggest that private saving will be far more important and widespread in its implications than employer-provided pensions achieved at their peak. The implications of this saving will play out over time, as new retirees will have spent an increasing portion of their careers contributing and accumulating assets in private retirement accounts. The first two chapters in this volume consider issues in retirement saving.

Chapter 1, on "Family Status Transitions, Latent Health, and the Post-Retirement Evolution of Assets," is the most recent in a series of NBER studies on retirement saving by James M. Poterba, Steven F. Venti, and David A. Wise. In most of our past work, we have focused on the accumulation of assets in retirement saving accounts, and the implications of retirement saving for the financial circumstances of those that will be retiring in the future. As noted already, the implications are profound. In this study, we redirect our attention from the accumulation phase to the use of accumulated assets in later life. This redirection is a natural evolution of our research agenda, as 401(k)-type plans have matured. What is happening as more new retirees have substantial 401(k) balances that few previous retirees had available? How will this growing financial resource affect older households, both now and in the future?

A key question is the extent to which people's accumulated assets are (a) annuitized, thereby assuring that one does not outlive their financial resources, no matter how long they live; (b) drawn down gradually over time, thus providing an ongoing and steady stream of financial support for post-retirement consumption, but without a long-life guarantee; versus (c) "saved for a rainy day," thereby conserving resources for irregular or unanticipated financial needs.

Our initial work on asset drawdown suggests that accumulated assets, including both housing equity and retirement accounts, tend to be conserved for transitional life events, such as adverse health changes, widowhood, or nursing home entry. The aim of chapter 1 is to analyze more carefully these relationships. We focus our analysis on two questions: first, whether the drawdown of assets is triggered by shocks to family status, and second, how the evolution of assets is related to health status. We use longitudinal Health and Retirement Study (HRS) data from 1992 to 2006, following households across two-year intervals between waves, and evaluating how their asset balances change as they age from one survey wave to the next. There are several key results.

First, households without major family status transitions tend to conserve their assets, even allowing them to grow further into older age. For example, among two-person households who were age fifty-six to sixty-one in 1992 (aging to seventy to seventy-five years old in 2006), and who remained two-person households across the survey period, the average wave-to-wave increase in total assets was 6.3 percent. Even for older cohorts, those who were seventy to seventy-five years old in 1993 (aging to eighty-three to eighty-eight years old in 2006), the average wave-to-wave increase in total assets was 4.6 percent. For continuing one-person households, the median estimates are smaller, but assets still tend to increase, except at much older ages, when they are closer to constant across waves.

Second, people in households that experience a family status transition during an interval—by becoming widowed or divorced—often experience a large decline or no increase in total assets. The wave-to-wave change in assets following a family status transition is very different from the continuing asset growth of households without this transition. Substantial declines in asset balances are associated with divorce, and the declines are statistically different from zero. The total assets of persons entering widowhood increase on average but the increase is not significantly different from zero. Households that experience a family status transition also have substantially lower levels of assets than continuing two-person households.

Third, both the level and evolution of household assets is strongly related to health. People in better health have more assets to begin with, and more continuing growth of those assets into older ages. For example, among continuing two-person HRS households age fifty-six to sixty-one in 1992, the ratio of assets of people in the top health quintile to the assets of people in

the bottom quintile was 1.7 in 1992, and grew to over 2.2 by 2006. Among continuing one-person HRS households age fifty-six to sixty-one in 1992, the ratio of assets of persons in the top health quintile to the assets of persons in the bottom grew from 2.8 in 1992 to 4.1 in 2006. Similar differences are found for older Asset and Health Dynamics Among the Oldest Old (AHEAD) households.

Our research in chapter 1 is empirical; we analyze data on how asset balances change after certain life events. In his discussant comments, David Laibson introduces a theoretical framework for modeling how individuals use their accumulated assets in later life. It provides a thoughtful theoretical perspective on the empirical findings that we report, and will inform our future work on the topic.

Chapter 2 is also part of a continuing series of studies on retirement saving, in this case by a research team made up of John Beshears, James J. Choi, David Laibson, and Brigitte C. Madrian. Most of their previous work has looked at how the characteristics of 401(k) plans affect the savings decisions of plan participants. The key finding from this work is that plan features like automatic enrollment and the default provisions of the plans have tremendous influence on the likelihood of participation, the amount saved, and the allocation of assets across the investment alternatives in the plan. More recent work has focused more extensively on portfolio allocation and its implications, such as the extent to which fund management fees are considered in investment decisions in 401(k) plans, and how the salience of management fees may impact investment decision making. Chapter 2 is another contribution in this ongoing series, asking in particular, "How Does Simplified Disclosure Affect Individuals' Mutual Fund Choices?" It adds one more piece to our understanding of how people make financial decisions, and how financial decision making might be improved through simple policy changes or interventions.

Motivated by concerns about the complexity of mutual fund prospectus statements, and by evidence that many investors may ignore them, the Securities and Exchange Commission (SEC) recently proposed and subsequently adopted a new simplified disclosure document. Mutual funds now have the option of sending investors this two- to four-page document, dubbed the "Summary Prospectus," instead of the more lengthy statutory prospectus. The Summary Prospectus contains key information about the mutual fund's investment objectives, strategies, risks, costs, and performance. This information can also be found in previously extant fund literature (the statutory prospectus, the Statement of Additional Information [SAI] and the shareholder report). The question explored in the study is whether people's investment behavior changes with the simplified presentation.

The methodology used in the study involves a behavioral experiment. The authors recruited 186 nonfaculty, white-collar Harvard employees to participate. The subjects were asked to allocate two hypothetical savings

account portfolios: one among four actively managed equity mutual funds, and one among four actively managed bond mutual funds. The amount of each subject's payment for participation depended on how their chosen portfolios actually performed subsequent to the experimental session.

The subjects were assigned randomly into one of three groups. In the first group, subjects received only the funds' more lengthy statutory prospectuses. In the second group, subjects received only the funds' simplified Summary Prospectuses. In the third group, subjects received the Summary Prospectuses but could additionally request the statutory prospectuses. Subjects were randomly assigned to be paid based on either their subsequent *one-month* portfolio return or their subsequent *one-year* portfolio return. This distinction was important because the funds with initial investment fees, or "loads," should have been viewed as less attractive when the time horizon of the investment was shorter.

The study finds that providing the Summary Prospectus does not alter subjects' investment choices. The dollar-weighted average fees and past returns of mutual fund choices are statistically indistinguishable across the three experimental groups. However, subjects receiving the Summary Prospectus spent less time on their investment decision. Thus the principal welfare gain from the Summary Prospectus appears to come from allowing investors to spend less time and effort to arrive at the same portfolio decision they would have come to after reading only the statutory prospectuses.

It is not clear that the portfolio allocation decisions that people make are good ones, however. Lack of understanding of sales loads, in particular, appears to be widespread. For example, mutual fund purchase fees, or loads, as well as redemption fees at liquidation, should be avoided to a greater degree when the investment horizon is short. Yet even when the subjects in the study had a one-month investment horizon—where minimizing loads is the only sensible strategy—they did not avoid loads. Indeed, they chose portfolios with loads plus redemption fees that were on average 200 basis points higher than the load-minimizing options. This implies that investors are either confused about loads, overlook them, or believe (implausibly) that their chosen portfolio will achieve an annual rate of return that is as much as 24 percentage points higher than the load-minimizing portfolio. The simplified Summary Prospectus developed by the SEC appears not to alleviate these kinds of errors.

As emphasized in the discussant comments by Steven F. Venti, the lack of responsiveness to fees, loads, and expenses in investor decisions could be far more important than other components of the prospectus presentation. Based on the findings from chapter 2, Venti suggests that a great deal of care should be taken in the design of materials provided to investors, such as reframing the prospectus to focus more on fees and loads and less on past returns.

Health Care Costs

Most of the our past NBER volumes on the economics of aging have included new investigations of health, health care, or health policy, and this thirteenth volume is no exception. With continuing advances in medicine, continuing increases in health care costs, and major health policy changes at the national level, health care will almost certainly remain at the forefront of economics of aging research going forward as well. Highlighted in this volume are three topics that have grown in importance to the well-being of the elderly, and that have brought them to the forefront of economics of aging research: growing out-of-pocket medical costs as an increasing risk to financial security, continuing growth of Medicare costs, and implementation of Medicare Part D.

In chapter 3, "The Risk of Out-of-Pocket Health Care Expenditure at the End of Life," Samuel Marshall, Kathleen McGarry, and Jonathan S. Skinner analyze the magnitude, variation, persistence, and composition of out-of-pocket medical expenditures among older households. They are interested in the extent to which high out-of-pocket burdens represent temporary financial shocks, such as from a more sudden acute illness, or persistent financial burdens, such as from a long-term chronic illness. They focus in particular on patterns of out-of-pocket medical spending in the period leading up to death. The study follows up on earlier work on out-of-pocket spending by these investigators, but with a new emphasis on spending variation among households across longer time periods, and in the period before death.

Based on HRS data from 1998 to 2006, the results suggest that out-of-pocket expenditures near death are considerable, averaging over $10,000 per person in the last 1.2 years of life. Not surprisingly, the distribution of out-of-pocket spending is highly skewed, with a median level of spending under $4,800, but a ninety-fifth percentile of nearly $38,000. There is also considerable persistence in out-of-pocket costs over longer time horizons. Mean out-of-pocket expenditures for the last five years (including both single and married households) is about $33,000, with a median of $21,416 and a ninety-fifth percentile of over $100,000. For married couples only, the study finds mean and median estimates of combined spending of $55,672 and $41,693 over the six-year period surrounding the death of a spouse.

The largest components of out-of-pocket spending include nursing home care, hospital care, insurance premiums, and helpers. Combined spending for nursing home and hospital care averages $2,658 per person. (Nursing home care makes up two-thirds of this amount.) The ninety-fifth percentile level of spending for this category is much higher, at $17,499. Insurance premiums are also important components of spending, averaging $1,975. The third major component is helpers, averaging $1,911 in out-of-pocket spending. Spending on helpers is highly concentrated among relatively fewer

households; the seventy-fifth percentile of spending is $0, the ninety-fifth percentile is over $13,000.

Out-of-pocket expenditures increase with both income and wealth. For example, the average out-of-pocket spending for the highest-income quartile is $11,469, compared with $7,443 for the lowest income quartile. The bottom wealth quartile spent an average of $5,731, compared with $14,324 in the top wealth quartile. These differences appear to be driven mostly by greater spending in the home on helpers, home health care, and other services to improve the independence of people living at home. Put differently, a large part of the incremental spending by households with greater financial resources is more discretionary spending in the home, buying independence and avoiding institutional care.

In his discussion of the chapter, David R. Weir adds perspective on the broader changes in finances that are likely to affect individuals in the period leading up to death, and that may explain changes in asset holdings that are larger than out-of-pocket spending for health care. He notes transfers to children, charitable donations, perhaps travel, and lost earnings (by self or spouse) as illustrations of how asset profiles might be altered by terminal illness.

Chapter 4 moves from out-of-pocket spending to Medicare spending. In "Cost Growth in Medicare: 1992 to 2006," Amitabh Chandra, Lindsay Sabik, and Jonathan S. Skinner examine trends in Medicare spending from 1992 to 2006 across all hospital referral regions (HRRs) in the United States. They decompose cost growth into changes in the distribution of the population among high- and low-spending HRRs, spending per enrollee, number, or medical encounters, and spending per encounter. They also look at the differences between Part A and Part B spending growth, and between different service categories within Part A and Part B. The study is an updated version of NBER work on Medicare cost trends that was conducted by other investigators in past years of the program, and it identifies some differences in cost trends in more recent years.

The study looks first at aggregate trends in Medicare spending and enrollment. Overall per capita Medicare fee-for-service spending grew at a real rate of 3.8 percent per year between 1992 and 1999, and at a rate of 2.7 percent per year from 1999 to 2006. Throughout both of these periods growth in Part B spending was higher than growth in Part A spending. Among subcategories of Part A and Part B spending, reimbursements for inpatient short stays and skilled nursing facilities had the highest absolute growth during the earlier years, while outpatient hospital services and medical care services had the highest absolute growth during the latter years. Relative growth in reimbursements for hospice services was high across both periods, though baseline spending for this category was low. Per enrollee spending in the highest cost area, hospital short stays, was relatively flat, while per enrollee spending on hospital outpatient services and physician medical care

services more than doubled over the fourteen-year period. Reimbursements for home health care services grew quickly from 1992 to 1996, but dropped off after the Balanced Budget Act (BBA) of 1997 changed the reimbursement rules for home care. Patterns of growth in total Medicare spending are similar to per capita patterns, though total spending increases more rapidly in recent years with the acceleration of Medicare enrollment.

Both the number of medical encounters and the amount spent per encounter increase over the study period; though in latter years, the growth comes almost entirely from increased spending per encounter. Of course, within these aggregates, different types of services have different usage trends. For example, there are increases in physician medical care services and diagnostic, lab, and x-ray services. These result from more Medicare enrollees receiving these services, and from increasing frequency of use among those who receive them. For hospital outpatient services, the study finds similar increases in both the percent of enrollees using services and the number of encounters per enrollee. There is also wide variation in growth rates by procedure, with procedures like coronary artery bypass graft (CABG) and hip fracture remaining relatively flat or declining over the study period, while back surgery, hip replacement, knee replacement, and percutaneous coronary intervention (PCI) increase substantially. In general the procedures that experienced the greatest increases in use are those that may be more discretionary, rather than clearly indicated clinically.

Analyzing the data geographically, the large majority of Medicare spending growth is found to result from increased spending within HRRs, rather than from a redistribution of the population into higher spending HRRs. The rate of spending growth, however, varies considerably across HRRs. For Part A spending, the average annual growth rate of the slowest growing quintile of HRRs was 1.9 percent, while the fastest growing quintile had an average growth rate of 2.7 percent. Growth in Part B averaged 4.4 percent in the slowest growing quintile and 6.1 percent in the fastest growing quintile. High growth HRRs can be found with both high and low baseline levels of spending, and in larger and smaller referral regions.

A final section of the study looks at how "financial entrepreneurship" may differ across hospital referral regions. The authors seek to identify financial entrepreneurship by looking at the impact on home health care spending of the regulatory tightening in the Balanced Budget Act of 1997. Those affected more significantly by the regulations may have been more proactive in taking advantage of the looser reimbursement rules when they were in effect before the BBA. Thus the change in home care reimbursement following the BBA may be an indicator of more active financial entrepreneurship. Interestingly, both Part B reimbursements and purchases of durable medical equipment were found to be significantly higher in these more "entrepreneurial" HRR markets.

While chapter 4 analyzes historical spending trends in Medicare Parts A

and B, chapter 5 focuses on the new Medicare Part D prescription drug program. Participation in Medicare Part D requires individuals to make active enrollment and plan choice decisions. Active decisions are made initially when first enrolling in the program, and can be revisited periodically during the program's annual open enrollment periods. Whether consumers make wise choices at these decision points is of interest not only for the evaluation of this particular market, but for consumer-directed health care more generally. In chapter 5, "The Demand for Medicare Part D Prescription Drug Coverage: Evidence from Four Waves of the Retirement Perspectives Survey," Florian Heiss, Daniel McFadden, and Joachim Winter explore how people decided among the prescription drug plan options made available through Medicare Part D, based on their prior drug use, self-rated health, and other factors.

To obtain data on health, prescription drug use, and Part D selection, the authors conducted four Internet-based interviews with samples of older Americans. They call their survey the Retirement Perspectives Survey, or RPS. The first wave of the RPS survey was conducted in November 2005, just before enrollment into Part D began. After the initial enrollment period closed on May 15, 2006, the authors reinterviewed the same respondents to elicit their actual Medicare Part D decisions for 2006. Third and fourth waves of the RPS survey were conducted in March/April 2007 and March/April 2009, enabling an expansion of the database to the decision making that occurs during the annual open enrollment periods. The latter waves also contained added measures of decision-making competence, planning horizon, and attitudes toward risk.

The analysis in chapter 5 concentrates on "active" decision making, either by enrolling in a Medicare Part D plan for the first time, or actively switching plans during the eligible open enrollment windows. The aim is to understand whether choices were related to the salient features of the program and the economic incentives they generated. Given the structure of the program, an individual's expected drug costs for the first year should be the most important determinant of plan selection. The results suggest that this is indeed the case. Enrollment in Medicare Part D was found to be driven strongly by the number of drugs used on a regular basis in the previous year. In all RPS waves, respondents who take three or more prescription drugs on a regular basis are much more likely to have stand-alone Part D coverage than those who take fewer drugs. Similarly, respondents whose self-rated health is "excellent" are less likely to have stand-alone Part D coverage.

All other variables, including those measuring decision-making competence, planning horizon, and attitudes toward risk, had little predictive power in explaining plan enrollment. The overall conclusion from the analysis is that consumers respond to the immediate incentives as they relate to their current health status and drug expenditures.

In her discussion of the chapter, Amy Finkelstein highlights the success

of the study in identifying adverse selection, as distinct from moral hazard, in Medicare Part D coverage. She notes that it is rare to find situations in panel data where one can use past behavior (in this case prescription drug use) to demonstrate adverse selection in subsequent decision making about insurance. She also provides perspective on whether choice among Medicare Part D coverage options leads to better or worse social outcomes.

Socioeconomic Circumstances and Health

The third section of the volume deals with the relationship between socioeconomic characteristics and health. This too has been a continuing theme through most of the NBER's Economics of Aging books. Measures of well-being in one domain (income, wealth, or education, for example) are highly correlated and causally related to measures of well-being in the other domain (self-reported health, chronic illness rates, disability, or mortality, for example). The range of studies advancing our understanding of this relationship has encompassed a breadth of likely causal pathways, characteristics of the relationship, and implications; and has drawn on both single-country and comparative cross-national investigations.

The next four chapters in this volume address topics that continue to advance our understanding of how socioeconomic circumstances and health interrelate. The issues covered include the implications of differential mortality by income on the progressivity of the Social Security system, cognition and economic outcomes, religion and health, and scale variations in self-reported work disability. Each moves the science forward, either through an incremental substantive question, new data resource, or methodological development. The first two analyze issues in the United States; the latter two are multinational studies.

In chapter 6, "Differential Mortality by Income and Social Security Progressivity," Gopi Shah Goda, John B. Shoven, and Sita Nataraj Slavov assess the implications of differential mortality by income for the lifetime progressivity of the "old-age" or retirement portion of Social Security. It is well known, for example, that Social Security has a highly progressive benefit formula that applies in the determination of the monthly benefit amounts from the program. Workers with low lifetime earnings get a monthly payment stream with a much higher replacement rate than workers with high lifetime earnings. However, because of differential mortality by income, those with low lifetime earnings will on average receive their Social Security benefits for a shorter period of years. Thus, some of the progressivity in the benefit amount is counterbalanced by the longer average lifetimes experienced by higher lifetime income recipients of Social Security. The goal of this study is to quantify these offsetting effects.

The methodology used to study overall lifetime progressivity of age-based Social Security benefits is to calculate internal rates of return (IRRs) and

net present values (NPVs) for the program under assumptions of differential mortality, as compared with population-average mortality rates. The study also considers how progressivity has changed from older to younger cohorts of program participants, as differential mortality by income has increased over time. The analyses are conducted using both hypothetical (or illustrative) earnings profiles, and the actual earnings histories of a sample of Social Security participants.

The key finding from the study is that differential mortality can change significantly the distributional characteristics of the program, though its impact varies by gender and cohort. Among older cohorts, differential mortality makes virtually no difference for women and a relatively small difference for men. For these older cohorts, the differences in mortality by income are not strong enough to offset the basic progressivity of the benefit formula. For more recent cohorts, differential mortality makes a substantial difference for men in particular. For example, in the 1938 birth cohort, differential mortality reduces the IRR from 1.51 percent to 1.07 percent for low income men, and raises it from 0.75 percent to 1.28 percent for high-income men. That means that for this cohort, once differential mortality is taken into account, men in the seventy-fifth percentile of income receive a higher rate of return on their payments into the system than men in the twenty-fifth percentile of income. The result is driven by the fact that mortality inequality is much larger for the younger cohorts. At least in terms of rates of return, an apparently progressive system becomes regressive.

Other finds from the study are that women experience higher IRRs and NPVs than men, because women have longer life expectancies than men. Also, for both men and women, later cohorts experience higher IRRs and NPVs than earlier ones, due to increases in overall life expectancy for these later cohorts. Of course this cross-cohort trend could reverse again in the future, depending on how Social Security taxes, benefits, and ages of eligibility for benefits continue to evolve going forward.

The discussion by Michael Hurd compares the results of this chapter with previous findings about Social Security progressivity, and explains some of the differences between them. In doing so, Hurd notes certain flaws in the current analysis, and a more comprehensive perspective from which to evaluate Social Security progressivity. The discussion also offers perspective on the term progressivity, and its relationship to public policy objectives.

Chapter 7 looks at another aspect of the socioeconomic status (SES)-health dynamic, focusing on cognitive skills as potentially influencing economic well-being over the life course. Indeed, the growing participation and saving in 401(k) plans may have created an increased impact of cognition, as individuals must decide for themselves how much to contribute to the plans, and how to invest their accumulated savings among the options available in the plan. In "Cognition and Economic Outcomes in the Health and Retirement Survey," John J. McArdle, James P. Smith, and Robert Willis examine

the association of cognitive skills with wealth, wealth growth, and wealth composition for people in their pre- and post-retirement years.

The analysis relies on selected waves of HRS data, supplemented by a cognitive economics survey (CogEcon) that measures cognitive capabilities in more depth. The HRS is well-known for its high-quality measurement of many key SES variables, including income and wealth. In some waves of the HRS, cognitive measures are also assessed, including immediate and delayed memory recall, episodic memory and intact mental status, numeracy, numerical reasoning, and retrieval fluency. The study explores the relationship between these various cognitive measures and wealth outcomes. The study also compares the cognition-wealth relationship for the person in a household who is most responsible for finances with the cognition-wealth relationship for the person who is less financially responsible.

The study finds that "numeracy," as measured by answers to three relatively simple mathematical questions, is by far the most predictive of wealth among all cognitive variables in the HRS sample. The association between numeracy and wealth holds for both the financial and nonfinancial respondents, but the magnitude of the estimated impact is much higher for the financial respondent. In quantitative terms, the estimated effect of the financial respondent answering one of the "numeracy" questions correctly is a roughly $30,000 increase in household wealth; while the estimated effect of the nonfinancial respondent answering a question correctly is a roughly $10,000 increase in household wealth.

Other cognitive measures were also predictive of household wealth, but less powerfully than numeracy. For example, the limited HRS data on number series (as a broader measure of numerical reasoning) showed some impact on wealth. The more complicated and time intensive measurement of number series in the face-to-face component of the CogEcon sample performed much better in predicting wealth. Episodic memory (or word recall) also appears to be related to the total and financial wealth holdings of the family, and applies to both the financial and nonfinancial respondent. The variables on mental status and retrieval fluency had very weak and erratic relationships with the financial outcomes measured.

The authors emphasize that their research is not about the causal pathways underlying the cognition-wealth association, but is more purely exploratory about the existence of the relationship. In his discussion, Finis Welch also highlights the exploratory character of the chapter, and introduces some of the questions one would want to address in follow-up research.

In chapter 8, Angus Deaton explores aging, religion, and health (also the chapter title). He considers both the determinants and consequences of religion. For example, determinants of religiosity and religious practice may include demographic characteristics such as age and gender, economic factors such as income and education, and attitudes of government toward religion. Among the consequences of religion may be health-related

behavior and health status, as well as broader measures of well-being. To explore these issues, Deaton relies on data from the Gallup World Poll, which randomly samples individuals from 146 countries around the world.

The chapter works with a model in which religiosity and religious practice are caused by income, education, age, and sex according to a stable set of patterns. Religion in turn is considered as a causal factor influencing health. While Deaton acknowledges substantial uncertainty, multidirectionality, and theoretical ambiguity in the causal relationships between religion and other variables, the study is presented as a largely exploratory first step in learning about these relationships. The study estimates these patterns for each country separately, and then examines similarities and differences across countries. It is a largely descriptive analysis of patterns of aging, religiosity, and health throughout the world.

In the vast majority of countries, women are found to be more religious than men, and the elderly more religious than the young. These two phenomena are related in that the difference in religiosity between men and women is positively and strongly correlated with the difference in religiosity between the young and the old. It is difficult to separate out age from cohort effects, but at least some of the evidence is consistent with pure age effects that are roughly consistent with rational choice theory—that religion should be postponed until late in life, that lower wages promote religiosity, and that the acquisition of religion can be postponed when life is longer. There is no obvious link between long-term income growth and the gap in religiosity between young and old, which is contrary to income-driven secularization.

The study also finds that, at least on average, over all countries, and over countries sorted into income groups, religious people do better on a number of health and health-related indicators. These protective effects appear to be strongest in poorer countries. Deaton emphasizes that none of these results show that the health benefits of religion can be obtained by joining a church, or even by undertaking a serious conversion. People who are religious are almost certainly different from nonreligious people in ways that go beyond their religiosity and beyond the basic educational and demographic controls that are used in the chapter. Even so, some of the correlations identified in the chapter appear remarkably universal across the religions and countries of the world.

Like the previous chapter, this study is highly exploratory; it is a very new area of economics of aging research, using a newly available data resource. Given the novelty of the research topic, and the novelty of the data, a real value in James P. Smith's discussant comments is the companion perspective offered. He interprets the empirical relationships presented in the chapter, and provides a second view, or a second set of observations on the apparent relationships between economic circumstances, religion, and health, and what causes those relationships.

In chapter 9, "Work Disability, Work, and Justification Bias in Europe and the United States," Arie Kapteyn, James P. Smith, and Arthur van Soest look at response scale differences in self-reported health measurement across countries and socioeconomic groups. The study aims to apply an innovative data calibration methodology known as vignettes to interpret self-reported measures.

As motivation for the study, the authors note the significant variation across Western European countries in both the fraction of workers receiving disability insurance benefits and, separately, the fraction that report having a work-limiting disability. But to what extent are differences in disability enrollment a result of differences in functional disability? Or, alternatively, to what extent do differences in program enrollment lead people to self-report differences in health? The complication in trying to make meaningful comparisons across countries is that self-reported disability status (or self-reported health measurement of any kind) may not mean the same thing to different people in different circumstances in different places.

Put another way, the response scales on which people describe their health may be different from one person to the next, and may be systematically related to culture, demographic group (such as age, gender, race, or ethnicity), socioeconomic circumstance (such as income, education, or wealth), public policy, work status, or other characteristic. Self-reported responses to health questions may also be subject to "justification bias"—the greater likelihood of reporting a disabling health condition as justification for not working, for example.

One question explored in the study is the magnitude of the variation in response scale and justification bias across countries. A second question is whether the response scale variations and biases can be corrected in a way that allows for meaningful cross-country comparisons of health. The study seeks to identify justification bias and other systematic response variations across countries and across socioeconomic groups in the United States and Europe, focusing specifically on self-reported work-limiting disability.

The primary methodology used in the chapter is to apply anchoring vignettes. The basic idea behind vignettes is as a calibration tool. A vignette describes the health-related circumstances of a hypothetical individual, and asks the respondent to make an assessment of that person's work-limiting health status on a scale matching the scale used to self-report their own work-limiting health status. For example, is the hypothetical person "not at all limited" in the work they do, "mildly limited," "moderately limited," "severely limited," or "cannot do any work?" The responses to the vignettes help to interpret what the respondent means in describing their own health status; they help calibrate the self-reported disability status measurement to a kind of benchmark response.

The study finds that people's self-reported definition of what constitutes a health-related work limitation is related to the generosity of earnings

replacement schemes and employment protection in different countries. In countries with more generous income replacement programs, there is an increased likelihood of reporting a health-related work limitation that is independent of actual health status. The study finds further that the variation in self-reported disabilities across European countries declines substantially when vignettes are used to calibrate responses, and to make the response scales more comparable across countries.

There are also noticeable differences in response scales between the United States and the European countries studied. Americans appear to be less likely to self-report themselves as work disabled than Europeans. The results suggest as well that justification bias plays a larger role in response patterns in the United States; Americans appear to use health-related work disability as a justification for not working, whereas Europeans do not feel the need to do so. The effect of health limitations on employment at older ages is estimated to be about twice as large in the United States as it is in Europe, reflecting at least in part the earlier departure from the labor force that occurs in many European countries regardless of health.

The discussant comments by Angus Deaton provide a critical perspective on the use of vignettes as a data calibration tool. Deaton argues that the use of vignettes rejects the assumption that people's self-reports of disabilities are internationally comparable, and replaces it with an assumption that their capacity for empathy is internationally comparable. Since the two assumptions are very similar, and similarly plausible or implausible, he sees little gained from their use. The ability of vignettes to correct response scales across different languages and cultures is also questioned. Together, the study and the critique provide informative lessons, an illustrative application and analytic perspective on this innovative methodological approach.

Aging in Less Developed Countries

In the last several NBER volumes on the economics of aging, we have added work on aging, health, and living circumstances in less developed countries. These studies have informed our knowledge of the particular challenges of aging in very poor regions of the world. The research has also explored the potential for improving health and economic circumstances, based on public policy changes that have been implemented, and experimental interventions that have been tested. The collection of high-quality data on these poor regions of the world has been a particularly important goal of past work, and several papers have been produced describing household circumstances, based on these new data.

The last two studies in this volume also deal with aging in less developed countries. One involves an experimental health intervention in a very poor region of rural India. The other is a study of how high spending on funerals affects the household finances and life circumstances of families in a

very poor region of South Africa. Both studies are extensions of past work on health and economic circumstances in these regions by the respective investigator teams.

In chapter 10, "Is Decentralized Iron Fortification a Feasible Option to Fight Anemia Among the Poorest?," Abhijit Banerjee, Esther Duflo, and Rachel Glennerster describe the impact of a village-level health intervention to fortify locally milled flour. Iron deficiency is believed to be the most common nutrient deficiency in the world today, and is thought to cause reduced productivity, increased susceptibility to illness, and cognitive difficulties in childhood. Iron supplementation of foods is considered an attractive means to reducing anemia, because it requires no additional effort on the part of the consumer, and can be done cheaply in centralized locations. However, for very poor and isolated populations, such as the population in the tribal district of Udaipur, where this study was conducted, centralized food fortification is not a practical solution: most households consume their own grain, and do not purchase any goods that could easily be fortified. The only way to reach these households is to fortify flour at the village level. So for this experimental study, local millers were trained and supplied with simple equipment to fortify flour in a safe and easily implemented way.

The intervention was implemented in 68 villages, randomly chosen out of 134. A first objective of the evaluation was to assess the logistical feasibility of the intervention: is it possible to recruit, train, and monitor millers and to keep them regularly supplied? Will the population accept the program? A second objective was to determine the impact of the program on anemia. To this end, the researchers collected data on hemoglobin levels at baseline, midline, and end line. The third objective was to determine whether the program had any health effect, beyond a possible reduction in anemia. To achieve this, the investigators collected rich data on health at baseline and end line and administered a unique monthly health survey. The final objective was to assess whether there would be any economic impacts of the program, such as through increased work capacity and higher school attendance.

The results were mixed. Program take up increased steeply over the first six months, but subsequently declined. Ultimately take up was quite low (around 30 percent of flour was fortified). The program was effective in reducing anemia as long as the take up was high enough, but ineffective when and where take up was low. It also reduced symptoms of fatigue when take up was sufficiently high. It did not lead to other improvement in health, or to increases in labor supply.

Because the tangible health benefits of the program may have seemed modest to the people receiving the iron supplements, their willingness to pay (or to add even a minor additional inconvenience) appears to be low. For example, the drop in program participation was faster among people whose nearest miller was not fortifying flour as well as for those who had

to walk more than 1.5 kilometers to find a fortifying miller. Ultimately, low demand from the households seems to be at the root of the decay of the program. This is despite some positive impact on symptoms of weakness in the program's initial phase.

The discussion of the chapter by Amitabh Chandra and Heidi Williams explores explanations for the mixed success of the experimental intervention, and potential changes or additions to the intervention that could improve its effectiveness.

In chapter 11, "Requiescat in Pace? The Consequences of High-Priced Funerals in South Africa," Anne Case and Alicia Menendez highlight the substantial fraction of a household's economic resources that are expended on funerals in South Africa, and the implications of this spending for economic well-being. According to the authors, funerals in South Africa are generally considered an individual's most important rite of passage. As a result, they tend to be more elaborate and expensive than weddings, graduations, or naming ceremonies for children. Households may spend the equivalent of a year's income for an adult's funeral, borrowing from money lenders if need be to have a funeral that befits the status of the household and of the person who died. The study is an attempt to quantify these costs, and to analyze their impact on the circumstances and functioning of the household.

The study analyzes data that the authors helped to collect in the Agincourt Demographic Surveillance Site in South Africa in 2004. The data describe the economic and health circumstances of about 3,000 individuals in close to 500 households. The authors find that the average amount spent on an adult funeral is 3,400 rand—equivalent to 40 percent of average annual total household expenditures. The implications of this spending for the lives of surviving families appear to be substantial. For example, households that experienced a death in the past five years have significantly lower expenditures per person than do other households. Adults in households that experienced a death report significantly more symptoms of depression and anxiety, and significantly more problems in their households. Children in households that experienced a death in the past five years are significantly less likely to be enrolled in school than are other children their age.

Many of these effects on households appear to result from the amount of money that the household spent on the funeral. In other words, the differences in household circumstances are found to be associated not just from the death itself, but from the amount of financial resources that a household contributes to the funeral. For example, the larger the household's financial contribution to the funeral, the less likely it is that children are enrolled in school, and the more likely it is that adults are depressed and anxious.

The authors find little evidence that households that experienced deaths were different in observable ways from other households. They find no association between the death of family members aged six or older and

assets holdings, maximum education of a member, or the type of household dwelling. Drawing on multiple aspects of the database, households that experienced a death appear much like other households in the demographic surveillance area. Thus the occurrence of the death, and the amount spent on the funeral, appear to result in poorer household circumstances along a number of dimensions. The implication is that reigning in the size of funerals may serve to improve post-funeral household circumstances.

The discussion of the study by Esther Duflo notes other countries where funeral spending is high, and where public policy has been used to try to contain it. The discussion also considers the reasons for high funeral spending, such as social and cultural norms, as well as the role of life or burial insurance.