

This PDF is a selection from a published volume from the National Bureau of Economic Research

Volume Title: Discoveries in the Economics of Aging

Volume Author/Editor: David A. Wise, editor

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-14609-X (cloth); 978-0-226-14609-6 (cloth);
978-0-226-14612-6 (EISBN)

Volume URL: <http://www.nber.org/books/wise13-1>

Conference Date: May 9-11, 2013


Publication Date: June 2014

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

Chapter Author(s): David A. Wise, Richard Woodbury

Chapter URL: <http://www.nber.org/chapters/c12961>

Chapter pages in book: (p. 1 - 18)



Introduction

David A. Wise and Richard Woodbury

The long-anticipated aging of the baby boom generation across the threshold of eligibility for Social Security and Medicare has arrived. The 76 million Americans making up the baby boom generation are currently between ages forty-eight and sixty-seven, and their initiation of retirement benefits is accelerating. The societal impact of aging baby boomers is compounded by longer life expectancies, which have risen continually over many decades. The implications of these demographic trends are extensive and significant, yet they are just one part of the rapidly changing landscape of aging in the United States and around the world.

The changing landscape includes a number of long-term trends, such as increased saving in 401(k)-type retirement plans, rising health care costs and, as noted, age demographics. It also includes unanticipated pressures, such as volatility in financial and housing markets, and strained macro-economic conditions. The impact of the financial crisis and its continuing ramifications have emerged as key concerns, adding to the fiscal challenges of government, and complicating people's financial planning for later life. Research in the economics of aging seeks to understand the health and financial well-being of people as they age, and how well-being is affected by this changing landscape.

This is the fifteenth in a series of National Bureau of Economic Research (NBER) volumes synthesizing analyses of economics of aging research.

David A. Wise is the John F. Stambaugh Professor of Political Economy at the Kennedy School of Government at Harvard University. He is the area director of Health and Retirement Programs and director of the Program on the Economics of Aging at the National Bureau of Economic Research. Richard Woodbury is a senior administrator with the Program on the Economics of Aging at the National Bureau of Economic Research.

For acknowledgments, sources of research support, and disclosure of the authors' material financial relationships, if any, please see <http://www.nber.org/chapters/c12961.ack>.

The previous volumes in this NBER series 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*, *Research Findings in the Economics of Aging*, *Explorations in the Economics of Aging* and *Investigations in the Economics of Aging*.

The goal is to bring together studies that are at the forefront of research in the field. The volumes are not intended to cover the entire area of economics of aging research, but rather to highlight cutting edge research projects that together contribute to a more comprehensive understanding of health and economic well-being as people age. Many of the studies are 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. Through fifteen 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.

A particular focus of the research reported in this volume deals with health, and its relationship to financial well-being. Why is health so important? First, health is perhaps the most essential aspect of what constitutes well-being as we age. As people live longer, it is important whether those increased years of life are characterized by poor health and functional disability, or by good health and functional independence. Second, health affects one's ability to work at older ages, and is strongly associated with financial well-being. And third, health has societal implications, such as for labor markets, government finances, and health care costs.

In past work, we developed a structural framework for studying health and disability, which we summarize in figure I.1. This framework includes key factors that influence health (the "determinants" in the left section of the figure), the multiple dimensions through which health is measured or characterized (the "characteristics" in the center section of the figure), and some important implications of health (the "consequences" in the right section of the figure).

While the arrows in figure I.1 suggest a unidirectional flow from determinants to characteristics to consequences, much of the research reported in this volume suggests more complicated interactions between variables. For example, health may affect work and retirement decisions at older ages. But whether someone retires may also affect their health. Also emphasized in the volume is the potential for interventions and policy changes to improve health and well-being, using approaches that may be implemented throughout this system of health-related interactions.

The first three chapters of the volume deal with health measurement and health trends. They fit largely within the center section of the figure I.1 framework. Chapter 1 looks at trends in morbidity. Chapter 2 looks at the

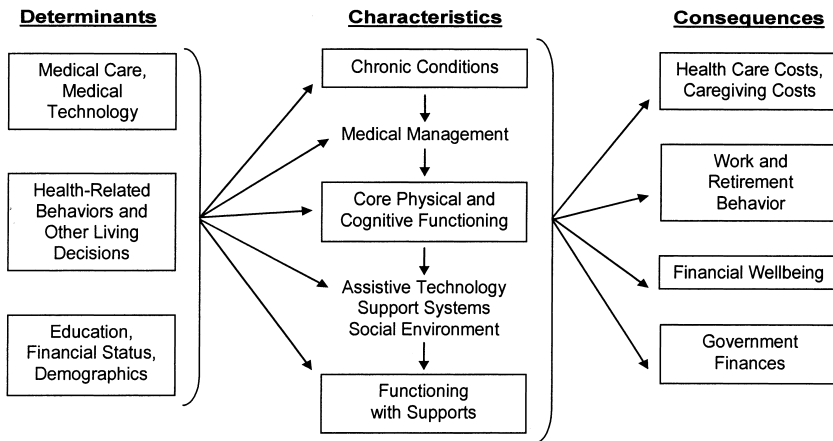


Fig. I.1 A framework for studying health and disability

lifetime risk of nursing home use. Chapter 3 analyzes the differences between various indices of health that have been used for research.

Chapters 4 and 5 look at the relationships and causal interactions between health and financial circumstances. As noted, these relationships are complex, as economic circumstances affect health, as health affects economic circumstances, and as both aspects of well-being interact continually over the life course. In the framework of figure I.1, financial circumstances are explicitly included as both determinants and consequences of health. Chapter 4 explores the extent to which better health and better financial circumstances are related to each other in the latter years of people's lives. Chapter 5 focuses on the causal relationships between health and economic well-being, and how they begin from early childhood.

The next four chapters in the volume consider how other aspects of people's lives affect their health. They fit best in the left sections of figure I.1, on the determinants of health. Chapter 6 looks at whether retirement improves or harms health. Chapter 7 looks at spousal effects on health. Chapter 8 looks at the effects of living with grandchildren. And chapter 9 looks at how aging affects optimism, uncertainty, and potential cognitive decline.

The last three chapters in the volume look at the potential for innovations, interventions, and public policies to improve health and financial well-being. Chapter 10 analyzes an experimental intervention to reduce anemia in a low-income region of the world. Chapter 11 looks at the uneven dissemination of medical advances. Chapter 12 looks at how the availability of a Roth 401(k) option has affected saving in employer-sponsored retirement plans.

The remainder of 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 work.

Part I: Health and Disability

Continuing increases in life expectancy are one factor in the changing landscape of aging in the United States. Using data from the National Center for Health Statistics, life expectancy at age sixty-two is currently about twenty years for men and twenty-three years for women. The number of years of life expectancy has increased by about a year every decade for at least the last four decades. Longer life is valuable to people, but it is even more valuable if the additional years lived are in good health. For the public sector as well, the consequences of longer lives depend on their quality. Medical spending for healthy seniors is modest, while spending for individuals with severely disabilities is much greater. Part I of the volume looks at trends in health impairments, as well as evaluating alternative measurements of health.

In chapter 1, David M. Cutler, Kaushik Ghosh, and Mary Beth Landrum present “Evidence for Significant Compression of Morbidity in the Elderly US Population.” The question of whether morbidity is being compressed into the period just before death has been at the center of health debates in the United States for some time. If morbidity is being compressed into the period just before death, the impacts of population aging are not as severe as if additional life involves many years of expensive care.

Empirical evidence on trends in morbidity is unclear. Some studies suggest that morbidity is being compressed into the period just before death, while others believe that the period of disabled life is expanding or that the evidence is more mixed. There are three reasons for this disagreement. First, there is not a single definition of morbidity. Some studies look at whether people report specific chronic conditions, which have increased over time, while other studies look at functioning. Second, it is often difficult to link health to the stage of life of the individual. If people are reporting more chronic disease, is that in the period just before the end of life, in which case the additional disease does not encompass many years? Or is the disease occurring in periods of time far from the end of life, in which case it represents many years of poor health? To answer this question one needs data on quality of life matched to time until death, and most cross-section data sources do not have such a link. Third, the data samples that tend to be used often focus on a particular subset of the population; for example, the noninstitutionalized. Since there are changes in the residential location of the elderly population over time, focusing on population subsets can give biased results.

Chapter 1 examines the issue of compression of morbidity, addressing these three concerns. The primary data source is the Medicare Current Beneficiary Survey (MCBS). The study analyzes health trends, linked to death records, for a representative sample of the entire elderly population between 1991 and 2009. The data are used in two ways. First, the authors examine

trends in various measures of morbidity by time until death. They consider a number of different metrics: the presence of disease; whether the person reports activities of daily living (ADL) or instrumental activities of daily living (IADL) disability; and various summary measure of functioning that draw together nineteen different dimensions of health. They show trends overall and by time until death.

As is well known, the MCBS data from the 1990s and 2000s show a reduction in the share of elderly people who report ADL or IADL limitations. A first result of this study is that this reduction in disability is most marked among those with many years until death. Health status in the year or two just prior to death has been relatively constant over time; in contrast, health measured three or more years before death has improved measurably. In chapter 1, these changes are translated into years of disability-free life expectancy and years of disabled life expectancy. The authors find that disability-free life expectancy is increasing over time, while disabled life expectancy is falling.

For a typical person age sixty-five, life expectancy increased by 0.7 years between 1992 and 2005. Disability-free life expectancy increased by 1.6 years; disabled life expectancy fell by 0.9 years. The reduction in disabled life expectancy and the increase in disability-free life expectancy were found for both genders and for nonwhites as well as whites. Hence, morbidity is being compressed into the period just before death.

A major question raised by these results is why this occurred. How much of this trend is a result of medical care versus other social and environmental factors? The results do not speak to this issue, but they give us a metric for analyzing the impact of changes that have occurred.

In their discussion of chapter 1, Daniel McFadden and Wei Xie highlight the significant differences between trends in disease prevalence and trends in functional morbidity. For many disease categories, for example, they show increases in disease prevalence, based on analysis of Medicare claims data from 1999 to 2010. When combined with the results from chapter 1, the implication is a sharp drop in the proportion of people with diagnosed diseases who also have functional morbidity. It is an open question, according to the discussants, whether this results from improved coping skills and functional aids, or improved and earlier diagnosis and treatment, or more aggressive disease coding that makes people with a diagnosed condition less sick on average. The discussants suggest further research on whether the measured increase in disease prevalence is a result of actual disease increases or more aggressive diagnosis, and further research on how people are managing their health conditions.

In chapter 2, Michael D. Hurd, Pierre-Carl Michaud, and Susann Rohwedder analyze “The Lifetime Risk of Nursing Home Use.” The risk of spending for long-term care is one of the most important risks faced by older households. However, finding data to estimate the risk has been difficult because of the necessity of following individuals over long periods

of time. The study in chapter 2 is based on ten waves of data from the Health and Retirement Study (HRS), following individual respondents for up to two decades. The data are used to assess the lifetime distribution of stays in nursing homes and, by consequence, the long-term care risk of nursing home use faced by households.

While the HRS only samples from the noninstitutionalized population at baseline, participants continue to be followed in subsequent survey waves, even if they move to a nursing home. As a result, after several waves of responses, the nursing home residence rates in the HRS sample closely reflect the residence rates in the population as a whole. In addition to the interviews with primary respondents in the sample, HRS data include proxy interviews, usually with a spouse or other close relative, for those unable to participate in a given interview wave. In addition, and particularly important for this study, the HRS data contain “exit interviews” with a proxy after the death of a primary respondent. The exit interviews allow investigators to estimate lifetime risk of a nursing home stay both nonparametrically and with a flexible transition model that simulates nursing home histories.

Similar results are found using both analytic approaches. Specifically, the authors find that a fifty-year-old has a 53 to 59 percent chance of ever staying in a nursing home in their lifetime. This likelihood is considerably higher than the risk reported in previous literature. Conditional on entering a nursing home, the average number of nights spent in a nursing home over the lifetime is just over a year (370 days). Of course, the 370-day average hides considerable variation in nursing home use across the population, including the extremely long stays of some individuals.

The study also looks at how sociodemographic factors affect the lifetime risk of using a nursing home. The results of this part of the analysis highlight two competing influences: first, the risk of entering a nursing home at any given age and second, the risk of dying younger as a result of poor health. Both relate to sociodemographic characteristics. For example, smokers have a higher risk of entering a nursing home at any age than nonsmokers. But since they also die younger than nonsmokers, on average, their lifetime exposure to nursing home risk is reduced. Combining both influences, the study finds that being female, white, and a nonsmoker are associated with higher lifetime risk, because average life expectancy is longer, and because nursing home use rises at older ages.

In his discussion of chapter 2, David Cutler proposes two extensions to the analysis for future study. First, he suggests differentiating between short- and long-term nursing home stays, which differ in both purpose and financing. Short stays are typically used to recover from acute events, and are generally covered by Medicare. Long stays are associated with frailty or severe and worsening impairment, such as from Alzheimer’s disease, Parkinson’s disease, and other degenerative impairments. Payment for these stays generally comes from individuals, their family, or Medicaid. Second,

Cutler suggests further exploration on the various substitutes for nursing home care. For example, inpatient rehabilitation services offer an alternative to skilled nursing facilities for shorter-term recovery, and assisted living is an alternative to nursing homes for long-term care.

Chapter 3 is “A Comparison of Different Measures of Health and Their Relation to Labor Force Transitions at Older Ages,” authored by Arie Kapteyn and Erik Meijer. Health can be characterized by a large number of indicators. For many analytic purposes, it is desirable to integrate multiple indicators into a single health index. A number of health indexes have been proposed in the literature, varying in statistical methodology and in the breadth of variables used to construct the index. Since different health indexes may be used for different purposes, there is no need to settle on any one preferred index. What matters is the statistical property of the index, what aspects of health are being described by the index, and how the index relates to economic behavior and outcomes.

Health indexes can be constructed using a number of different approaches. The simplest approach is to simply ask people to rate their own health on an ordinal scale. A second, more involved approach relates such self-reports to other explanatory variables, such as health conditions or difficulties with activities of daily living. Regressions can be used to weight the explanatory variables in the construction of the health index. A third approach considers health to be a latent construct for which a number of indicators exist; and the indicators can then be used to estimate the underlying latent variable. The study in chapter 3 compares these approaches.

The data are from the eleven countries that are in both waves 1 and 2 of the Survey of Health, Ageing and Retirement in Europe (SHARE). The traditional health measure is self-reported general health (SRH), which has five categories: excellent, very good, good, fair, and poor. SRH generally correlates strongly with objective measures of health. It is a short and easy question, and is widely available in many data sets. This makes it a useful measure for many purposes. However, it is also a crude measure, and it appears to be incomparable across countries without corrections. Hence, for comparing health across countries, it is not very suitable.

Three other health indices are considered in the chapter, all of which draw on a larger number of explanatory health variables. Their potential advantages over SRH are continuous values, greater reliability, and improved comparability across countries. The authors label these indices as MKA, PVW, and Jue, referring to the investigators who constructed them. A goal of the study is to describe the theoretical and empirical differences between these indices, so that researchers who want to include a measure of health in their analyses can make an informed choice as to which index is most appropriate, and so that readers can interpret differences between results from papers that use different indices.

The most important difference between the indices is in the choice of

variables that are included in their construction. Among the explanatory variables used to construct one or another of the indices are mobility limitations, ADLs, IADLs, self-reported health, physical attributes like grip strength and body mass index, specific health conditions, pain, and health care utilization. Indices may also draw on variables that may be correlated with health, such as gender, age, living with spouse or partner, household size, education, and net worth. The chapter helps to understand the uses of these various indices for different research applications.

In his discussion of chapter 3, Steven F. Venti elaborates on the issues that complicate the construction of health indices; and particularly those that can be applied in a cross-national research context. The discussion covers three considerations in developing a health index. The first is the choice of a statistical model that translates available health measures into a single index. The second is the choice of health measures to include in the construction. The third is how to account for country-specific reporting bias. Cross-national variation in respondent reported health measures may arise from genuine differences in health, or from the way that residents of each country answer questions. The challenge, Venti emphasizes, is distinguishing between genuine health effects and reporting bias.

Part II: Health and Financial Well-Being

The studies in part II of the volume analyze relationships and causal interactions between health and financial circumstances. In chapter 4, James M. Poterba, Steven F. Venti and I look at “The Nexus of Social Security Benefits, Health, and Wealth at Death.” Our study focuses on the drawdown of assets between the first year an individual is observed in the Asset and Health Dynamics Among the Oldest Old (AHEAD) data (1995) and the last year that individual is observed before death. We relate the drawdown of assets over this period to an individual’s health, Social Security benefits, and other annuity benefits. By considering income from Social Security and defined benefit (DB) pensions jointly with changes in asset stocks, we develop a more complete picture of the financial resources available to the elderly. We are also interested in the association between health status and these other variables.

We find that a significant fraction of people approach the end of life with few financial assets and no home equity, relying almost entirely on Social Security benefits for support. Whether people reach late life with positive nonannuity wealth depends importantly on health, which is quite persistent over the lifetime. People in poor health in old age have a higher-than-average probability of having experienced low earnings while in the labor force, which puts them at greater risk of having low Social Security benefits in retirement. While the progressivity of the Social Security benefit formula provides a safety net to support low-wage workers in retirement, a notice-

able fraction of people, especially those in single-person households, still have income below the poverty level in their last years of life. Many of these individuals have few assets to draw on to supplement their income, and are in poor health.

In addition to confirming the strong relationship between health and financial well-being in later life, our results also show that higher Social Security income and higher defined benefit pension benefits are strongly “protective” of nonannuity assets. Those with larger income flows from Social Security and defined benefit pensions are less likely to exhaust their nonannuitized assets.

While these are our general conclusions, it is difficult to summarize the drawdown of assets in any simple way; there is enormous variation across people. Because many individuals were observed in 1995 with relatively low levels of nonannuity assets, the median percent drawdown is sometimes quite large even though the dollar amount of drawdown is small. People who remained single and married persons predeceased by a spouse experienced median asset reductions of 30 to 50 percent between 1995 and the last year observed before their death. The reductions for persons whose spouse outlived them were much smaller.

In his discussion of chapter 4, Jonathan Skinner suggests that further work on consumption at older ages will be important to understanding more fully asset trends in later life. Poor health, Skinner agrees, is central to declines in wealth. He references previous work suggesting that mean levels of out-of-pocket expenditures in the last five years of life are remarkably large. His suggestion for future research is to focus in more detail on the components of consumption that are most likely to be variable near death.

In chapter 5, Till Stowasser, Florian Heiss, Daniel McFadden, and Joachim Winter report on “Understanding the SES Gradient in Health Among the Elderly: The Role of Childhood Circumstances.” They introduce their study as the classic “chicken and egg” problem. We know that people with high socioeconomic status (SES) tend to be in better health and live longer than their economically disadvantaged counterparts—but we are not sure which came first. Do economic resources determine health (hypothesis A)? Or does health influence economic success (hypothesis B)? Or are both health and wealth dependent on some third unaccounted factor (hypothesis C)?

The traditional view that causality flows from SES to health is especially common among epidemiologists. Often cited causal pathways are the affordability of health services, better health knowledge and lifestyles among the higher educated, environmental hazards associated with poorly paying occupations and low-income living conditions, or the mere psychological burden that comes with a life of constant economic struggle. Economists were among the first to argue that causality may also work its way from health to economic outcomes. For example, physical frailty is likely to have adverse effects on educational attainment, occupational productivity and,

consequently, the accumulation of wealth. In addition to these direct causal pathways, the observed correlation between health and financial well-being is, at least in part, likely caused by factors that jointly affect both. Family circumstances in childhood, for example, may have an influence on both health and financial well-being later in life.

While many past studies have explored these relationships, the research in chapter 5 draws on the increasing availability of retrospective life-history data within large panel studies. These data innovations are relevant, because of the potential long-term influences of early life circumstances on health and financial well-being at older ages. First, by incorporating longer health histories, one can construct a more realistic model of health dynamics. Second, to the extent that retrospective data also covers information on family backgrounds and parental SES, it will be possible to study factors that may jointly influence both health and wealth. Third, controlling for both historic and contemporary variables may elucidate when the association between SES and health is established.

The results confirm that childhood health has lasting predictive power for adult health. The study also uncovers strong gender differences in the intertemporal transmission of SES and health. While the link between SES and functional as well as mental health among men appears to be established later in life, the gradient among women seems to originate from childhood.

In his discussion of chapter 5, Robert J. Willis provides additional insights both on this study and on the studies that preceded it. An original 2003 study was controversial, because it suggested noncausation from SES to health, a finding that Willis emphasizes was narrowly applicable to a particular sample of quite elderly people who were largely retired and covered by Medicare. A 2012 follow-up study used a larger sample, a longer period of observation, and a wider age range, and the findings suggested that any of the three causal pathways were possible. Willis interprets the findings from this chapter as reinforcing that multidimensional conclusion, noting that any causal account of the determinants of the SES-health gradient is likely to be very complex, with room for feedback loops involving causation running in multiple directions.

Part III: Determinants of Health

The studies in part III of the volume consider other determinants of health, including retirement, marriage, living with grandchildren, and life expectations.

In chapter 6, Axel Börsch-Supan and Morten Schuth consider “Early Retirement, Mental Health, and Social Networks.” Early retirement is popular in Europe, as it is in other parts of the world. It is widely viewed as a social achievement that increases personal well-being, particularly among employees who suffer from work-related health problems. First introduced in

the 1970s and 1980s, generous early retirement provisions in most European countries were instituted with minimal actuarial adjustments. In response to financial pressures, the costs of early retirement have come under increased scrutiny, leading to reforms in many European countries since the 1990s.

The question addressed in this study is whether early retirement actually improves well-being. An immediate benefit from early retirement is the receipt of income support without the necessity to continue working, enabling individuals to enjoy more leisure. Moreover, early retirement relieves workers who feel constrained in their place of work, whether due to stressful job conditions or to work-impeding health problems. For such individuals, early retirement should manifest itself in an improvement of well-being and, potentially, also health. On the other hand, early retirement might also be harmful, because individuals who stop working may lose social connections, or a sense of purpose in life. This might, in turn, decrease subjective well-being and mental health.

Research on the causal impact of early retirement on health is complicated by the fact that survey measures of well-being, cognition, and health may suffer from justification bias. That is, early retirees may report worse health in order to justify their early exit from the workforce. Moreover, early retirement is not an exogenous outcome; it is related to health. The aim of the study in chapter 6 is to disentangle these relationships.

The analysis takes advantage of innovative social network data in wave 4 of the Survey of Health Ageing and Retirement in Europe (SHARE). SHARE wave 4 includes a name generator that identifies people with whom the respondent “discuss things that are important to them,” such as “good or bad things that happen to you, problems you are having, or important concerns you may have.”

The study finds a significant erosion of social networks after retirement. Retirement in general and early retirement in particular, reduces the size of the social network, and in particular the number of friends and other non-family interpersonal contacts. Put differently, social contacts are a side effect of employment that keeps workers mentally agile. The study presents evidence that early retirement has negative effects on people’s social networks which, in turn, accelerates cognitive aging.

In her discussion of chapter 6, Elaine Kelly highlights the challenge of evaluating causation in this type of investigation. Importantly, the timing of retirement may be determined by both current and expected future health and cognition. This makes it especially difficult to analyze how retirement causally affects future health and cognition. Similar difficulties in estimation arise from the interactions between social networks, the timing of retirement, and cognition. Kelly discusses potential identification strategies to address these analytic challenges. She also suggests further research on how these connections vary across the characteristics of individuals as a way to better understand the mechanisms through which retirement, cognition, and social networks interrelate.

In chapter 7, “Spousal Health Effects: The Role of Selection,” James Banks, Elaine Kelly, and James P. Smith look at the tendency for people to choose a spouse with similar characteristics as themselves. For example, if healthy people marry healthy people, unhealthy people marry unhealthy people, and the health of a spouse affects one’s own health, then partner selection will exacerbate health inequalities in a population.

Health histories of partners may matter for at least three reasons. First, individuals may select their partners based in part on their partner’s health history and current health status. Second, partner selection may depend on factors such as education and health behaviors (smoking, drinking, and exercise), which are correlated with current and future health. Third, couples typically share a common lifestyle and household environment, leading to more closely correlated health outcomes over time.

Chapter 7 explores these issues in the context of England and the United States. The investigators find a strong and positive association in family background variables including education of partners and their parents. Adult health behaviors such as smoking, drinking, and exercise are more positively associated in England compared to the United States. Childhood health indicators are also positively associated across partners. In general, these correlations are more positive for first than for subsequent partnerships. Especially for women, poor childhood health is associated with future marital disruptions in both countries.

The study explores in greater depth the pre- and postpartnership smoking behavior of couples. The results indicate that smokers are much more likely to partner with smokers and nonsmokers with nonsmokers; and this relationship is far stronger in England compared to the United States. In the United States, the influence of a partner’s smoking behavior on one’s own smoking behavior is asymmetric. Men’s premarriage smoking behavior influences his female partner’s postmarriage smoking behavior. But women’s premarriage smoking behavior does not appear to influence their male partner’s postmarital smoking. These influences are much more symmetric across genders in England.

In his discussion of chapter 7, Amitabh Chandra relates the study to some well-publicized prior work by Nicholas Christakis on “Mortality after the Hospitalization of a Spouse.” That research suggested that having a sick spouse was bad for a partner’s health, increasing their mortality risk. Drawing on the results from chapter 7, Amitabh notes that some of this relationship is likely explained by partner selection, rather than entirely by the causal effects of bad health across spousal partners.

In chapter 8, Angus Deaton and Arthur A. Stone present “Grandpa and the Snapper: The Well-Being of the Elderly who live with Children.” This study lies at the intersection of two literatures, one on whether children bring well-being to those who live with them, and one on the living arrangements of the elderly. Whether or not children make their parents’ life better is an

old question that remains unsettled. Some even suggest a more complicated relationship in which both are true: parents of children gaining more happiness and more enjoyment, as well as more stress and more worry.

The literature on the living arrangements of the elderly in the United States argues that the elderly value their ability to live independently. Those who are living with children under eighteen, therefore, are more likely to be doing so because of low income or poor health. On the other hand, outside of the United States and other rich countries, it is common for the elderly to live in multigenerational families. Where this is the case, there is less reason to believe that there is negative selection into living with children among the elderly. In such places, we should observe something closer to the direct effects of living with children.

This study analyzes two large data sets collected by Gallup, one for the United States, the Gallup-Healthways Well-Being Index, and one for 161 countries around the world, the Gallup World Poll. They include measures of life evaluation as well as a range of emotional well-being measures. They also have the advantage of using identical questions in all locations. These advantages are offset by incomplete information on living arrangements. In particular, we have information on one respondent from each household, and know only whether or not there is a child at home, not the relationship of the respondent to that child.

The study finds that elderly Americans who live with people under age eighteen have lower life evaluations than those who do not. They also experience worse emotional outcomes, including less happiness and enjoyment, and more stress, worry, and anger. In part, these negative outcomes come from selection into living with a child, especially selection on poor health, which is associated with worse outcomes irrespective of living conditions. Yet even with controls, the elderly who live with children do worse. This is in sharp contrast to younger adults who live with children, likely their own, whose life evaluation is no different in the presence of the child once background conditions are controlled for. Parents, like elders, have enhanced negative emotions in the presence of a child, but unlike elders, also have enhanced positive emotions.

In parts of the world where fertility rates are higher, the elderly do not appear to have lower life evaluations when they live with children; such living arrangements are more usual, and the selection into them is less negative. They also share with younger adults the enhanced positive and negative emotions that come with children.

In his discussion of chapter 8, David Laibson emphasizes that the relationships between living with children and life satisfaction need not be causal. He makes the case that selection probably lies behind the results of the study. Laibson discusses four kinds of selection that may be relevant. The first is adverse selection on the characteristics of older adults: “Grandpa is disabled so he’s going to move in with us so we can take better care of him.” The

second is adverse selection on the characteristics of the middle generation: “We need to move in with Grandpa, since we can no longer afford to live independently.” The third is advantageous selection on the characteristics of older adults: “Grandpa is rich and has invited us to move in with him.” The fourth is advantageous selection on the characteristics of the middle generation: “We have decided to ask Grandpa to move in with us since we are doing so well.” Laibson describes how the results of the study are consistent with these theories of selection, including the differences between developed and developing countries.

In chapter 9, Gábor Kézdi and Robert J. Willis explore “Expectations, Aging, and Cognitive Decline.” They use data from the Health and Retirement Study (HRS) to document general patterns in expectations in various domains with respect to aging and to investigate the potential role of cognitive decline in those patterns. They focus on two aspects of expectations: optimism and uncertainty. People who assign higher probabilities to events with positive consequences are considered more optimistic. People who respond to survey questions with “don’t know” or “50 percent” are considered more uncertain. The measures are based on subjective beliefs about stock market returns one year in the future, the chance of a future economic depression, whether tomorrow will be a sunny day, whether one’s income will keep up with inflation, job loss, and survival to a specific age.

Aging appears to decrease optimism and increase uncertainty. Optimism with respect to stock market expectations, expectations that income will keep up with inflation, and expectations of sunshine the next day all decline strongly with age. The increase in uncertainty is less robust and depends on the measure of uncertainty.

Aging could have these effects for several reasons. The authors speculate that cognitive decline associated with aging may affect an individual’s view of the world and their ability to process information about the world, causing a person to overstate the likelihood of negative events and to hold less precise probabilistic beliefs. Another possibility is that the increase in the awareness of mortality that accompanies aging leads to decreased attention to events that are farther in the future, or to the relevance of particular types of economic events.

In his discussion of chapter 9, John B. Shoven notes the emphasis of the chapter on how cognition changes as an individual person ages. He suggests that an interesting question for future research might disentangle the effect of individual aging on cognition (which this chapter addresses), the selection effect (the fact that healthier people live longer, causing cross-sectional measures of cognition by age to differ from longitudinal measures of cognition by age), and the cohort effect (the fact that the cognitive health of a person at any given age may be improving over time, as is the case with other aspects of health). Shoven also notes the complexity in interpreting responses from the HRS survey. For example, answering “I don’t know” to a question about

Dow Jones Industrial stocks might actually be more likely from someone with excellent cognition, whereas someone with reduced cognition might guess at an answer.

Part IV: Interventions to Improve Health and Well-Being

The studies in part IV of the volume explore the potential for innovations, interventions, and public policy to improve health and financial well-being.

Chapter 10, by Abhijit Banerjee, Sharon Barnhardt, and Esther Duflo, looks at “Nutrition, Iron Deficiency Anemia, and the Demand for Iron-Fortified Salt: Evidence from an Experiment in Rural Bihar.” Iron deficiency anemia (IDA) is frequent among the poor worldwide. For children, IDA is associated with slower physical and cognitive development with potentially long-lasting effects. For adults, IDA may lower energy, productivity, and physical performance, and accelerate cognitive declines at older ages. Severe anemia during pregnancy can lead to low birth weight and child mortality.

While IDA can be prevented with the appropriate supplement or food fortification, these programs often do not reach the poorest. Providing supplements to a large population, particularly pregnant women, is a standard policy in many countries. However, it faces two problems. The first is that it relies on public health infrastructure and local providers that are difficult to monitor. The second is that individuals often do not comply with the protocol. A second approach is to add iron to foods that are a regular part of the local diet. Fortification is a compelling solution in locations where households regularly purchase packaged foods that can be fortified centrally during mass production. These channels do not effectively reach low-income populations in remote locations, however, because such populations do not buy as much processed grain.

The experiment described in chapter 10 explores an alternative approach, which is to subsidize salt that is fortified with iron and iodine, known as double-fortified salt (DFS). The chapter describes first steps and preliminary analysis of baseline data, from a large scale randomized controlled trial in 400 villages in Bihar. The baseline survey strongly suggests the need for an intervention to fight anemia. The study finds that 53 percent of women age 15–49 have hemoglobin levels under 12 g/dL and 21 percent of men have a hemoglobin level under 13, the rough cutoffs for anemia. A large majority of households (94 percent) purchase iodized salt, which makes an intervention to provide DFS potentially promising.

The baseline survey indicated that anemia is prevalent, and may be both caused by and a cause of poverty: households with low expenditure per capita and with low diversity in their diet are more likely to have an anemic member. Anemic individuals are weaker, sicker, and perform worse on cognitive tests than nonanemic individuals. Finding a way to solve this issue on a large scale is important for policy, and would also give us an opportunity

for the first time to reliably measure the impact of a plausible instrument to fight IDA on health and economic outcomes.

The chapter also presents results from a small-scale experiment to assess willingness-to-pay for double-fortified salt using randomly assigned discount vouchers. The results show that the take-up of DFS falls quickly with price. At a price point of 45 percent of the retail price of DFS sold in major Indian metros, the take-up of DFS is 30 percent in private stores. The study also assesses the impact on purchase behavior of three separate information campaigns, though no differential impact was found among them, and the effects were small when used without price incentives.

In his discussion of chapter 10, Amitabh Chandra places this study in the broader context of human behavior, asking, “Why do people not always make decisions that are in their best economic interest?” After referencing a range of other research findings in which behavior appears to counter people’s best interests, Chandra introduces a theory about how differences between people may affect such behavior. He notes, in particular, differences in time and hassle costs among people with more and less stress and complexity in their lives; differences in the side effects of treatment across individuals; and differences in the benefits of certain behaviors, conditional on those side effects.

In chapter 11, Amitabh Chandra, David Malenka, and Jonathan Skinner analyze “The Diffusion of New Medical Technology: The Case of Drug-Eluting Stents.” Their focus is on the wide variation across hospitals and geographic regions in the diffusion, using drug-eluting stents as an illustrative case study. Drug-eluting stents are a commonly used approach to treating the narrowing of coronary arteries.

Before 2003, only bare metal stents were available to cardiologists seeking to perform revascularization for blockages in the heart. These cylindrical wire meshes were designed to keep arteries from narrowing. Yet bare metal stents were also subject to restenosis, or a renarrowing of the artery, leading to restricted blood flow. In April of 2003, the Food and Drug Administration (FDA) approved the use of coated antiproliferative drug-eluting stents, designed to reduce restenosis. In the same month, Medicare allowed for a higher reimbursement for drug-eluting stents to cover their higher cost. Adoption of the new technology was rapid; by December 2003 more than 65 percent of all stent placements in the Medicare population were drug eluting rather than bare metal stents. Yet different hospitals exhibited very different diffusion rates. In the bottom quintile of diffusion, drug-eluting stents comprised just 33 percent of total stents for the year following FDA approval, while in the top quintile the equivalent was 83 percent. The study in chapter 11 analyzes why some hospitals adopt drug-eluting stents earlier than others.

There are a variety of suggested explanations. One is profitability. Drug-eluting stents may not by themselves be more profitable than previous treat-

ments, but they could confer a competitive advantage to hospitals seeking to charge insurance companies and employers higher prices for high-quality care. A second explanation is based on provider expertise; the possibility that higher quality hospitals are the first to adopt drug-eluting stents, because they have better knowledge about technological advances. A third hypothesis stresses knowledge spillovers; diffusion based on area norms or copycat behavior. A final hypothesis is that producers allocate drug-eluting stents to those hospitals whose patients are most likely to benefit from them.

The hypothesis most consistent with the empirical findings is that better quality hospitals adopt technology quicker. There is also suggestive evidence that hospitals whose patients are most likely to benefit from technology are quicker to adopt it. There is no support for models of competition, knowledge spillovers, or profit motivations.

The authors note that rapid adoption of new technologies is not always welfare improving. For example, drug-eluting stents were subsequently found to have more risks than previously understood in the early months of their introduction. So in this case, and likely others, there do not appear to be large welfare costs associated with the uneven diffusion rates.

In his discussion of chapter 11, Jay Bhattacharya elaborates on the biological function of stents, and on the prevention and treatment of heart disease more generally. He also describes the historical dissemination of drug-eluting stents from their first patent in 1997 through their testing, FDA approval, and dissemination over the subsequent fifteen years. Bhattacharya also reiterates the dual implications of early adoption of new technologies: providing the benefits of the new treatment, on the one hand, but serving as test subjects for the development of new technology, on the other.

In chapter 12, John Beshears, James J. Choi, David Laibson, and Brigitte C. Madrian explore “Who Uses the Roth 401(k), and How Do They Use It?” Beginning in 2006, employers sponsoring a 401(k) plan were allowed to offer a Roth option in their plans. Like contributions to a Roth individual retirement account (IRA), employee contributions to a Roth 401(k) or 403(b) are not deductible from current taxable income, but withdrawals of principal, interest, and capital gains in retirement are tax free. Roth contributions are advantageous to households whose current marginal tax rate is lower than their marginal tax rate in retirement. If households understand this fact, then we would expect younger employees to be more likely to allocate contributions to the Roth. Employees with transitorily low income would also be expected to utilize the Roth 401(k). If households are uncertain about whether their marginal tax rate will be higher or lower in retirement, they may wish to hedge this risk by contributing to both Roth and before-tax accounts in their 401(k).

Chapter 12 describes the characteristics of employees who take advantage of the Roth 401(k) options. The study uses administrative 401(k) plan data from twelve companies that introduced a Roth 401(k) option between 2006

and 2010. The results suggest somewhat limited use of Roth 401(k) contributions to date. One year after the Roth was introduced at these companies, just 8.6 percent of 401(k) participants had a positive balance in their Roth account, only 5.4 percent of contributions were to Roth accounts, and Roth balances made up only 1.8 percent of total 401(k) balances at these companies. Roth contributions were more significant for those who chose to make them. Conditional on having a positive Roth contribution rate, 66 percent of employee contributions go to the Roth. Consistent with the existence of a tax diversification motive, 55 percent of employees who contribute to the Roth also contribute to another 401(k) account.

The low usage of the Roth 401(k) may reflect an active preference against the Roth, but it can also be partially explained if employees who enrolled in the 401(k) when the Roth was unavailable failed to update their 401(k) elections in response to the introduction of the Roth. Supporting the importance of the passivity channel, the study finds that 19 percent of 401(k) participants who were hired after the Roth's introduction had a positive balance in the Roth after one year, while only 8 percent of 401(k) participants hired before the Roth's introduction had a positive balance.

The young are more likely to use the Roth and to allocate a larger fraction of their contributions to it. This correlation could be consistent with a rational response to the Roth's tax incentives, since Roth contributions are advantageous to those whose current marginal tax rate is lower than the marginal tax rate at which those contributions will later be withdrawn. Roth usage declines with age, is less likely among women, and only weakly correlated with salary and tenure once one controls for other employee characteristics.

In his discussion of chapter 12, James M. Poterba focuses on the complexities of comparing tax rates while employed with anticipated tax rates at retirement, thereby complicating the choice between traditional and Roth contributions. He describes individual uncertainties, such as one's future earnings or the age a person will retire. He describes the complexities of tax laws, such as interactions with the Earned Income Tax Credit while working, or the tax treatment of Social Security benefits when retired. There is also uncertainty in what reforms may be enacted to future tax policies. In response to the long-term fiscal gap in the federal budget, for example, one solution might raise tax rates to close the gap, another might broaden the tax base while retaining or even lowering rates, and a third might implement a value added tax that could also lower income tax rates.

As a future research project, Poterba suggests analyzing the behavior of individuals who contribute at the legal contribution limit of \$17,500. These participants are an interesting subsample, because their only option for increasing the effective set-aside for retirement is by shifting to the Roth option. Thus they may have more incentive to use the Roth option than savers who are below the contribution limit.