Introduction
Discoveries in the Economics of Aging

David A. Wise and Richard Woodbury
June 2013

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 48 and 67, and their initiation of retirement benefits is accelerating. The societal impact of aging baby boomers is compounded by rising life expectancies 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 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 macroeconomic 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 wellbeing of people as they age, and how wellbeing is affected by this changing landscape.

This is the fifteenth in a series of NBER volumes synthesizing analyses of economics of aging research. 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 wellbeing 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 wellbeing. Why is health so important? First, health is perhaps the most essential aspect of what constitutes wellbeing 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 wellbeing. 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 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).

Figure 1
A Framework for Studying Health and Disability

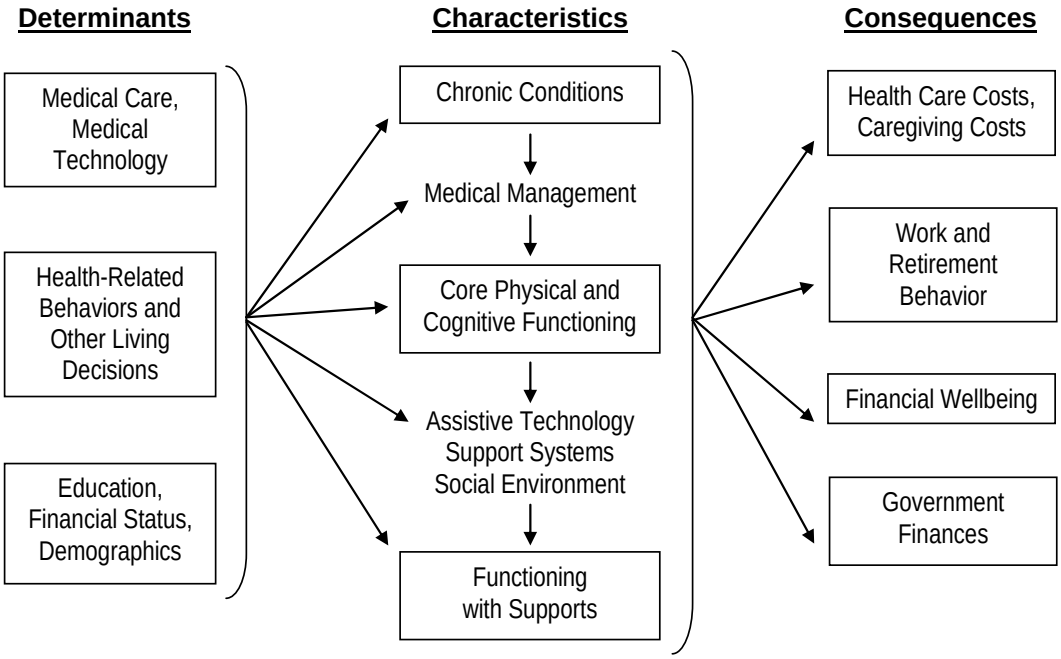

While the arrows in figure 1 suggest a unidirectional flow from determinants to characteristics to consequences, much of the research reported in this volume suggests more complicated interactions between variables. For example, health may affect work and retirement decisions at older ages. But whether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

The first three chapters deal with health measurement and health trends. They fit largely within the center section of the figure 1 framework. Chapter 1 looks at trends in morbidity. Chapter 2 looks at the 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 wellbeing interact continually over the life course. In the framework of figure 1, financial circumstances are explicitly included as both "determinants" and a "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 wellbeing, 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 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 wellbeing. 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 62 is currently about 20 years for men and 23 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 Cutler, Kaushik Ghosh and Mary Beth Landrum present "evidence for significant compression of morbidity in the elderly U.S. 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 non-

institutionalized. Since there are changes in the residential location of the elderly population over time, focusing on population subsets can give biased results.

Chapter one 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 ADL or IADL disability; and various summary measure of functioning that draw together 19 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 aged 65, 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 was found for both genders and for non-whites 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 chapter 2, Michael 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 relies on data from the Health and Retirement Study (HRS) 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 non-institutionalized 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 HRS closely reflects the nursing-home population because of turnover in nursing homes. Nearly all those in nursing homes at baseline will have died and been replaced by persons initially residing in the community and represented by HRS respondents.

The study relies on 10 waves of HRS data, including cohorts added after the original HRS cohort. 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 importantly 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 non-parametrically 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 50 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 results also highlight two competing forces that affect lifetime nursing home use: first, the risk of entering a nursing home at any given age, and second, the risk of dying younger. Both relate to socio-demographic characteristics. For example, smokers have a higher risk of entering a nursing home conditional on being alive. But since they also face higher mortality risks, this reduces lifetime exposure to nursing home risk. The study finds that being female, white and a non-smoker are associated with higher lifetime risk, because average life expectancy is longer, and because nursing home use rises at older ages.

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 purposes it is desirable to integrate multiple indicators into a heath index. Several health indexes have been proposed in the literature, varying in statistical methodology and the breadth of variables included in the index. Since health indexes may be constructed 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 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 11 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 paper, all of which draw on a larger number of explanatory health variables. Their potential advantages over SRH are continuous values, greater reliable, and 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 a better informed choice as to which index is most appropriate for their analyses, and readers can interpret differences between results from different 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 on 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, 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 purposes.

Part II – Health and Financial Wellbeing

The studies in part II of the volume analyze relationships and causal interactions between health and financial circumstances. In chapter 4, James Poterba, Steven Venti and I look at "the nexus of Social Security benefits, health, and wealth at death." Social Security benefits are the most important component of the income of a large fraction of older Americans. 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 non-annuity wealth depends importantly on health, which is quite persistent over the life-time. 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 noticeable 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 general, low assets and low income in old age are strongly related to poor health.

Our study focuses on the drawdown of assets between the first year an individual is observed in the 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 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 conclude that Social Security and defined benefit pension benefits are strongly "protective" of non-annuity assets, with a negative relationship between these income flows and the likelihood of exhausting non-annuity assets. We note that this result may in part reflect population heterogeneity in saving propensities. We also find that poor health is an important determinant of the drawdown of non-annuity wealth.

While these are our general conclusions, it is difficult to summarize the draw-down of assets in any simple way; there is enormous variation across people. Because many individuals were observed in 1995 with relatively low levels of non-annuity assets, the median percent drawdown is sometimes quite large even though the dollar amount of draw-down is small. Persons 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 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 socio-economic 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, the most important channel being the development of human capital: Physical frailty is likely to have adverse effects on educational attainment, occupational productivity and, consequently, the accumulation of wealth. Finally, the observed correlation between morbidity and SES is, at least in part, likely caused by unobserved individual heterogeneity which affects both health and wealth, due for instance to genetic disposition and other family effects which have an impact on preferences and health-relevant behaviors.

Discriminating among these rivaling hypotheses is important since policy recommendations will critically depend on the nature and the sources of the gradient. While many past studies have explored these issues, the research strategy used for the study reported here rests 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 wellbeing in later life.

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 proxy-control for some of the individual heterogeneity that is suspect of exerting a common influence on health and wealth. Third, controlling for both historic and contemporary variables may elucidate when the association between SES and health is established, which has important policy implications: If future outcomes are predetermined during childhood, resources spent on policies that aim at improving access to health care for adults and retirees may in fact be more wisely invested into educative and financial measures for young families.

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.

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,” based on European experience. 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 wellbeing, 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 in fact improves wellbeing. 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 wellbeing 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 wellbeing 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 ageing.

In chapter 7, "spousal health effects: the role of selection," James Banks, Elaine Kelly and James 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 post partnership smoking behavior of couples. The results indicate that smokers are much more likely to partner with smokers and non-smokers with non-smokers; 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 pre-marriage smoking behavior influences his female partner's post-marriage smoking behavior. But women's pre-marriage smoking behavior does not appear to influence their male partner's post-marital smoking. These influences are much more symmetric across genders in England.

In chapter 8, Angus Deaton and Arthur Stone present "grandpa and the snapper: the wellbeing of the elderly who live with children." This study lies at the intersection of two literatures, one on whether children bring wellbeing 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. Others suggest a more complicated relationship with 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 US argues that the elderly value their ability to live independently. Those who are living with children under 18, 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 multi-generational 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 Wellbeing Index, and one for 161 countries around the world, the Gallup

World Poll, with sample sizes of 1.8 million and 1.1 million individuals respectively. These data sets are rich in well---being questions, and include measures of life evaluation as well as a range of emotional wellbeing measures or hedonics. 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 18 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. The misery of the elderly living with children is one of the prices of the demographic transition.

In chapter 9, Gábor Kézdia and Robert 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.

Part IV – Interventions to Improve Health and Wellbeing

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

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 x 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% of women age 15-49 have hemoglobin levels under 12 g/dL and 21% of men have a hemoglobin level under 13, the rough cutoffs for anemia. A large majority of households (94%) 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 non-anemic 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% of the retail price of DFS sold in major Indian metros, the take-up of DFS is 30% 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 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 re-narrowing of the artery, leading to restricted blood flow. In April of 2003, the FDA approved the use of coated anti-proliferative 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% 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 x 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 treatments, 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 chapter 12, John Beshears, James Choi, David Laibson, and Brigitte 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 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% of 401(k) participants had a positive balance in their Roth account, only 5.4% of contributions were to Roth accounts, and Roth balances made up only 1.8% of total 401(k) balances at these companies. Roth contributions were more significant for those who choose to make them. Conditional on having a positive Roth contribution rate, 66% of employee contributions go to the Roth. Consistent with the existence of a tax diversification motive, 55% 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 fail 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% 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% 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.