Introduction
Investigations in the Economics of Aging

David A. Wise and Richard Woodbury
Draft: June 14, 2011

One of the most persistent findings from research in the economics of aging is the depth of the relationship between health and financial circumstances. These dimensions of wellbeing are not only highly correlated and causally connected. They are also changing rapidly over time; and in conjunction with an aging population, and with a broader evolution in how people live in later life. For example, one major change of the last two decades has been the growth in retirement saving programs, and particularly 401(k) plans, and a parallel decline in traditional defined-benefit pension programs. Reforms to Social Security, Medicare, and Medicaid are also being implemented, as necessitated by changing population demographics and mounting fiscal pressures. Health and health care are evolving as well, with continuing advances in medicine, better disease management, improving health and functional ability, and increasing longevity. The dynamic interrelationship between evolving health and evolving financial wellbeing is a core substantive focus of current research in the economics of aging field.

This is the fourteenth in a series of NBER volumes compiling economics of aging research. The goal is to present studies that are at the forefront of research in the field. The volumes are not so much integrative summaries of the field, but rather collections of cutting edge research on topics that are most current, encompassing advances in research methodology, data resources, current trends, and changing policies in health, work, financial wellbeing, and retirement. 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 fourteen volumes, the large majority of this research has been funded by the

National Institute on Aging, which has made a long-term commitment to advancing the economics of aging field.

The previous volumes 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* and *Explorations in the Economics of Aging*.

This volume continues the series with a collection of *Investigations in the Economics of Aging*. The volume is divided into two sections: the first weighted more heavily toward people's changing financial circumstances as they age, and the second weighted more heavily toward people's changing health and health care. Even in the first section, however, the emerging importance of out-of-pocket health care costs, or the risk of such costs, is emphasized as a growing need for financial resources at older ages. Indeed the interactions between health and financial resources, and between health care costs and resource adequacy, are the general topic of the first four chapters in the volume. The studies analyze people's finances in the years leading up to their death, the extent to which large out-of-pocket costs increase the risk of financial hardship at older ages, how health care cost growth has limited the inflation-protection in Social Security, and how loan provisions affect asset accumulations in 401(k) plans.

The second half of the volume begins with a study that more comprehensively decomposes the dimensions of health and functional ability, and how health is changing over time along these various dimensions. Following this broader perspective on health is a more narrowly-focused study on progress in cancer treatment and outcomes specifically. Another

more methodological study improves our ability to interpret self-reported disability measures through the use of reference groups. The last three chapters consider in greater depth the relationship between health and economic circumstances, the causal pathways from socioeconomic circumstances to health, the lifelong impact of childhood health on later-life health outcomes, and how subjective wellbeing is affected by financial and macroeconomic conditions.

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

Part I: Financing Retirement

The composition of financial support at older ages and the demands on financial resources later in life are two key aspects of the changing retirement landscape in the United States. As noted, the trends to date suggest that private saving, particularly through 401(k) plans, will be far more important and widespread in its implications than employer-provided pensions

achieved at their peak. The implications of this saving will play out over time, as new retirees will have spent an increasing portion of their careers contributing and accumulating assets in private retirement accounts. While financial market declines have had an impact, the most recent generation of new retirees is the first to have significant numbers of people with more substantial 401(k) accumulations. This has inspired a redirection of research in the field from the accumulation phase in these retirement savings accounts, while people are working, to how the accumulated assets are used in later life.

On the expense side, out-of-pocket health care costs are a significant and growing demand on financial resources in later life. This burden may be particularly acute for individuals in poorer health, who not only face higher costs, but who (on average) have less financial resources to support them. Thus the first several chapters deal with interactions between health and financial resources, and between health care costs and resource adequacy.

In Chapter 1, James Poterba, Steven Venti and I analyze people's finances in the years leading up to their death. We ask: "were they prepared for retirement? financial status at advanced ages in the HRS and AHEAD cohorts." The study's methodology relies on panel data to analyze the evolution of each household's financial assets over this later life period. The sample consists of households who were age 70 and older in 1993. Each household is followed until they die, or until the most recent survey period in 2008. Separate analyses are conducted for households in each of three family pathway groups: (1) one-person households, (2) households that begin in 1993 as two person households, but with a deceased spouse over the study period, and (3) two-person households in which both spouses are living when last observed. Much of the analysis is restricted to people who are known to have died.

Several results stand out. First, median total wealth was relatively high in the year last observed for each of the three family pathway groups. Second, wealth in the last year before death is greatest for persons who were in two-person households the longest period of time. For example, the average assets in the last year observed were $141,606 for persons in one-person households in 1993 whose last year observed before death was 2006, $252,849 for persons in two-person households in 1993 whose spouse was deceased when last observed in 2006, and $691,588 for person in two-person households in 1993 whose spouse was alive when last observed in 2006. Third, for total wealth and for each of the asset sub-categories there is a strong correspondence between the level of assets in 1993 and the number of years a person survives after 1993. Persons who lived longer had higher initial assets. Fourth, for each family pathway group, there is a very strong relationship between health status and wealth in the last year observed. Thus there is a strong association between health and wealth even among persons who would die within the next two years. And fifth, despite the appearance of substantial assets at the median, a substantial fraction of people die with income less than $10,000 and with no financial assets and with zero housing wealth.

A rather large fraction of the original single-person households have low income judged by the income poverty thresholds. We find that 12.1 percent are below the poverty threshold and have no financial assets, and that 23 percent are below twice the income poverty line and have no financial assets. To put the results in context we first compare the total income in the last year observed to total income in 1993, the first year the AHEAD data were collected. Total income in the last year observed was about 4 percent higher, on average, than total income in 1993. We also compare total median income in the last year observed to median earnings (in 2008 dollars) of the same persons when they were between ages 57 and 62. While the difference is hard to

evaluate because the two measures are not clearly comparable, overall median income in the last year observed was approximately 50 percent lower than median earnings of the same persons at ages 57 to 62.

There are also important differences across the pathways. Consider for example the proportion of persons with annuity income less than $20,000 (approximately twice the poverty level for single persons over age 65) and financial assets less than $10,000: 52 percent of persons in the single-household pathway fall below these thresholds, but only 36.4 percent of those in the two- to one-person pathway, and only 26.3 percent of those in the two-person pathway. Similarly, consider the proportion of persons with annuity income in the $10,000 to $20,000 interval and financial assets less than $10,000 who have zero housing wealth: of those in the single family pathway 63.3 percent have zero housing wealth, 53.0 percent of those in the two-to one-person pathway, and only 25.9 percent of those in the two-person pathway. A perhaps striking similarity across the pathways is that given income and housing wealth the health status of the persons in the three pathways is very close. The median health percentile of persons with annuity income in the $10,000 to $20,000 interval and financial assets less than $10,000 is 28.2 for persons in the single-household pathway, 23.9 percent for those in the two-to one-person pathway, and 28.2 for persons in the two-person pathway.

The results raise several issues. First, a noticeable fraction of persons die with virtually no financial assets—46.1 percent with less than $10,000. Based on a replacement rate comparison, many of these may be deemed to have been well-prepared for retirement, in the sense that their income in their final years was not substantially lower than their income in their late 50s or early 60s. Yet with such low asset levels, they would have little capacity to pay for unanticipated needs such as health or other shocks or to pay for entertainment, travel, or other

activities. This raises a question of whether the replacement ratio is a sufficient statistic for the "adequacy" of retirement preparation. In addition, this group relies almost entirely on Social Security benefits for support in retirement. These people balance on only one leg of the oft touted three-legged stool that is said to provide retirement support—Social Security, pension benefits, and personal saving. If the one leg is judged inadequate it raises the question of how to strengthen the other legs which in turn may, for example, increase interest in the spread of 401(k)-like plans to low-wage workers in firms with high turnover.

In discussing the results from chapter 1, David Laibson compares actual savings behavior and accumulations in 401(k) plans with the asset accumulations that might be necessary to assure resource adequacy in retirement. He does this by applying a rough conversion formula between asset amounts, on the one hand, and their annuitized income equivalent, on the other. Compared with his ideal vision for savings-based retirement planning, he raises concerns about resource adequacy resulting from many factors, including pre-retirement withdrawals, individuals without access to plans, low employer match rates, high management fees and risks of low returns, non-participation among those who are eligible, starting to save too late, the continuing decline of defined-benefit pension arrangements, reductions in housing equity at older ages, increased credit card debt, and rising out-of-pocket costs for medical care.

Chapters 2 and 3 focus on health care costs as a growing demand on financial resources at older ages. A key issue is the risk of being inadequately prepared. For example, a person may seem to have been well-prepared for retirement if they turn out to have lived an average lifespan, earned an average return on their investments, and had average out-of-pocket medical spending needs. But what if their experience turns out to be different than the average? Specifically, what if they live longer than average, or earn less than average on their investments, or have larger

than average medical care needs? In that case, the same household in retrospect may have been unprepared. Risk and uncertainty raise diverse questions about the magnitude of retirement assets needed, the annuitization of assets, insurance and the distribution of financial preparedness across the population of retirees.

In Chapter 2, Michael Hurd and Susann Rohwedder, look at "the effect of the risk of out-of-pocket spending for health care on economic preparation for retirement." It follows up on their earlier work that focuses on longevity risk. In their prior research, which excludes medical spending risk from their models, Hurd and Rohwedder find that a substantial majority of those just past the usual retirement age are adequately prepared for retirement, in that they will be able to finance a path of consumption that begins at their current level of consumption and then follows an age-pattern similar to that of current retirees. However, almost half of singles without a high school education will be likely to be forced to reduce consumption, and a substantial number of married college graduates will also have to reduce consumption. The model used for this work did not account for health-care spending risk.

The goal of the study in chapter 2 is to explicitly account for the risk of large out-of-pocket spending on health care. Variation in out-of-pocket spending for health care changes quite sharply the percentage of single persons adequately prepared for retirement. When there are no shocks, 57% are adequately prepared. With serially correlated health shocks, just 44% are adequately prepared—even though average out-of-pocket spending on health care does not change (by design). The difference for single women specifically is even larger: 58% with no shocks versus 39% with. For single women without a high school degree, the percent adequately prepared drops from 33 percent to just 15%.

The effects on married people are not nearly as large because many couples have wealth substantially in excess of what is required to deal with most health-care spending shocks. In this sample, the percentage of those adequately prepared declined from 80% with no health care spending shocks, to 73% with shocks

Averaging both single and married households, the percentage adequately prepared is 72% when out-of-pocket spending does not have a stochastic component. But when out-of-pocket spending is stochastic (with the distribution recorded in the 2008 wave of the HRS survey) and when spending is serially correlated (as estimated from 2004 and 2005 Medicare Current Beneficiary Surveys), the percentage of those who are financially prepared for retirement declines by 9 percentage points to about 63%.

In his discussion of chapter 2, Robert Willis notes the sequencing of asset use, based on post-retirement needs. Annuitized resources, such as Social Security and traditional pension benefits, are used as the initial or baseline resource, followed by financial asset savings, and housing equity as the longest protected asset. Willis also raises questions about the availability of annuity and other insurance products that may better protect against later life risks, as well as the roles of bequest motives and "excess wealth" that is not consumed before death.

In chapter 3, Gopi Shah Goda, John Shoven and Sita Slavov ask "how well are Social Security recipients protected from inflation," after accounting for increased Medicare premium and other out-of-pocket health care costs?

Social Security is widely believed to protect its recipients from a number of risks, including uncertainty regarding length of life and inflation, due to the inflation-indexed life annuity form of the benefit. The inflation protection comes from the fact that Social Security benefits are indexed to the Consumer Price Index for Urban Wage Earners and Clerical Workers

(CPI-W). The CPI-W is based on the spending patterns of a broad group of workers, representing approximately 32 percent of the U.S. population. However, the CPI-W may not accurately reflect the experience of retirees for two reasons. First, retirees generally have higher medical expenses than workers, and medical costs, in recent years, have tended to rise faster than the prices of other goods. Second, even if medical costs did not rise faster than the prices of other goods, individual retirees would still, on average, need to devote a larger share of income to medical spending as they age. This means that individual retirees would still see a decline in the real income they have available for non-medical spending. Chapter 3 explores both of these factors, the extent to which they undermine the inflation protection in Social Security, and what alternative methods of indexing benefits might be considered.

The study accounts for two major components of medical costs. First, most Social Security recipients are also participants in Medicare Part B. These premiums are automatically deducted from Social Security retirement benefits and have increased approximately 1,600 percent between 1975 and 2011, while the Social Security cost-of-living adjustments have accumulated to just over 300 percent over this period. Second, retirees often have substantial out-of-pocket medical expenses, including Medicare deductibles and co-payments, and payments for services with limited Medicare coverage, such as nursing home care.

The authors find that, after subtracting both of these components of health spending from Social Security benefits, available income net of medical expenses for a Social Security participant with average out-of-pocket medical spending has, in fact, been increasing more slowly than a price index of non-medical goods and services. For example, the average man born in 1918 has seen his monthly Social Security benefit, net of medical expenses, rise from $528 at the end of 1983 (when he was 65) to $867 at the end of 2007 (when he was 89). However, if his

net-of-medical-expenses benefit had kept pace with inflation in the prices of non-medical goods over that time period, he would have had $1,086 per month in 2007 after medical expenses. That is, his net-of-medical-spending benefit has declined by around 20 percent, relative to the non-medical goods price index. Similarly, the average woman born in 1918 has seen her net-of-medical-expenses benefit decline by around 27 percent relative to the non-medical goods price index.

An alternative inflation index is CPI-E, an experimental measure which targets the cost of living for people aged 62 and older. If Social Security benefits had been indexed to the CPI-E instead of the CPI-W, men born in 1918 would have $961 net of medical expenses, falling only 11.5 percent short of the $1,086 needed to hold non-medical expenditures constant in real terms; similarly, women born in 1918 would fall only 18.1 percent short. The reason indexing to the CPI-E does not fully compensate retirees for inflation is that, even if medical costs remained constant over time for the elderly, they tend to spend more on out-of-pocket medical expenses as they age, crowding out non-medical spending. Thus, each cohort's Social Security benefit net of average out-of-pocket medical spending would tend to decline in real terms even if the price of medical care rose at the same rate as the prices of other goods, or alternatively, even if the average retiree's real net Social Security benefit remained constant.

Both the authors and Michael Hurd's discussion of the study highlight limitations in both the CPI-W and CPI-E measures of inflation. For example, the failure to account properly for technological progress can be quite serious when it comes to health care. Higher medical costs likely reflect the consumption of better quality medical care, and retirees may be better off even if they are left with less to spend on other non-medical goods. The study's authors, therefore, emphasize that they cannot draw any conclusions about changes in the utility of Social Security

recipients from this analysis. All they show is that Social Security benefits may not be fully inflation-indexed in the sense that recipients with average out-of-pocket medical spending cannot, from one year to the next, purchase the same bundle of non-medical goods with their Social Security benefits.

Chapter 4 looks narrowly at one aspect of 401(k) plans which has not been studies extensively in past research: "the availability and utilization of 401(k) loans." These loans may decrease accumulation in 401(k) plans if their repayment displaces new contributions to the plans. They may increase asset accumulations, however, if they enable people to make larger contributions to the plans, knowing that they can still access their savings if necessary before retirement. In chapter 4, John Beshears, James Choi, David Laibson and Brigitte Madrian provide an introductory overview of how 401(k) loans work, and how often they are used. This paper is the first step in a research agenda on how the availability of 401(k) loans affects retirement wealth accumulation.

About 90% of 401(k) participants are in plans that offer a loan option. Within those plans, about one in five eligible participants has a loan outstanding at a given point in time. Loan utilization rates follow hump-shaped patterns with respect to age, tenure, compensation, and plan balances, reaching peaks for participants in their 40s, those with 10 to 20 years of tenure, those earning $40,000 to $60,000 per year, and those with $20,000 to $30,000 in plan balances. Conditional on having a loan, the loan balance to 401(k) balance ratio is declining in age, tenure, compensation, and 401(k) plan balance.

Despite the prevalence of 401(k) loans, they constitute only 2.5% of total plan assets among plans with a loan option. For some individuals, however, 401(k) loans can be an important source of credit. Our empirical analysis finds that 401(k) loan utilization is correlated

with the types of loan rules adopted by firms. Loans are more likely to be used in plans that charge low interest rates, and conditional on taking a loan, loan sizes are larger when multiple loans are allowed to be outstanding simultaneously, the maximum loan duration allowed is long, and the loan interest rate is high.

In her discussion of the study, Gopi Shah Goda suggests numerous directions for further research, most notably on the initial question of whether the loan provisions increase or decrease asset accumulation in the plans. She also asked what factors determine whether a plan adopts a loan policy, what factors determine whether an individual takes out a loan and, importantly, whether 401(k) plan participants have other borrowing (such as credit card debt) that could be attained less expensively through 401(k) loans.

Part II – Health and Health Care

The relationships between financial circumstances and health are emphasized throughout this volume. Part II is more heavily weighted toward studies of health. Understanding changes in the health of the elderly is a central policy issue. A healthier elderly population is able to work to later ages, spends less on medical care each year, and requires less informal care from family and friends.

By many metrics, the health of the elderly has improved over time. For example, the share of elderly people with basic physical impairments has declined markedly over the past two decades. By other metrics, however, the health of the elderly is worsening. Problems with more advanced functional measures such as stooping and walking moderate distances have increased over time, and obesity among the elderly has soared along with weight in the non-elderly population.

Researchers have attempted to combine these different measures of disability into one summary measure, but these summaries are generally ad hoc and difficult to interpret. The most common single measures of disability ask whether the person has impairments in Activities of Daily Living (ADLs, such as bathing or dressing) or Instrumental Activities of Daily Living (IADLs, such as doing light housework or managing money). In the Medicare Current Beneficiary Survey, which is analyzed in chapter 5, the share of the elderly population that is disabled by this definition has declined from 49 percent in 1992 to 43 percent in 2007. However, this summary measure exhibits somewhat different trends in different surveys and for different measures of health, and ignores some aspects of functional impairment (e.g., can the person walk a reasonable distance), cognitive problems such as memory loss, and sensory impairments such as difficulty seeing and hearing.

In chapter 5, David Cutler and Mary Beth Landrum aim to more carefully decompose the "dimensions of health in the elderly population," and how health is changing over time along these various dimensions. In the first part of the paper, they consider how to optimally combine different measures of health into a smaller number of summary measures. Of course, the best way to summarize multiple measures of health depends on the question being asked. The optimal measure to predict medical spending may be somewhat different than the optimal measure to predict health transitions, for example. Cutler and Landrum estimate factor models for 19 indicators of health. These measures include specific ADL impairments, IADL impairments, functional impairments, cognitive limitations, and sensory impairments.

The study shows that these 19 dimensions can be compressed into three broad summary measures. The dominant factor is impairment in very basic physical and social tasks such as dressing, eating, transferring in and out of bed, preparing meals, doing light housework, and

managing money. This encompasses many of the ADLs and IADLs, but not all. The second factor loads heavily on functional limitations and includes measures such as walking moderate distances, stooping, and reaching. The third dimension is sensory impairments – trouble seeing and hearing.

After determining these factors, the authors analyze the evolution of these health dimensions over time. They show that the set of physical and social limitations and sensory impairments have declined rapidly over time. Functional ability was flat or increasing, after declining early in the time period.

These results suggest many possible patterns. One possibility is that the community-dwelling population is increasingly concentrated among the less severely ill, with more severely ill individuals in nursing homes or having died. The study shows that composition changes – both people leaving the sample and new people entering the sample – cannot explain a change in the health of the community-dwelling population. In a second scenario, it may be that people are recovering from severe disability more frequently in later years in the sample, thanks to better medical care or other environmental changes.

The third part of chapter 5 investigates the evolution of health states, comparing the early years in the sample (1991-96), middle years (1997-2001) and later years (2002-2007). The study shows that health deteriorates less rapidly in later years of the sample than in earlier years. We do not know exactly why this occurs, but we show that the average trend masks significant variation across the population. Even as health deteriorates overall as people age, health is improving for a significant minority of people. The next step in this research agenda is to develop a richer model of the change in health over time. To what extent is the improvement in

health a result of fewer new conditions developing, existing problems being cared for better, or changes in the social and environmental circumstances that the elderly face?

In his discussion of chapter 5, David Weir emphasizes the importance of age standardization to the effective interpretation of the results. He notes that the study is more about the disability aspects of health than about chronic illness and disease. Because issues of disability are dramatically more prevalent in the older old population, such as people in their 80s and older, one needs to differentiate between the 65-plus population generally, and this older old group. This is particularly important in interpreting long-term population trends, as the baby boom generation moves into their 60s now, their 70s soon, and their 80s further in the future.

Chapter 6, by Alan Garber, Jay Bhattacharya, Matt Miller and Daniella Perlroth, focuses more narrowly on "the value of progress against cancer in the elderly." Cancer remains one of the most common causes of death in the elderly, but the number of cancer deaths in the United States began falling in the early 2000s. This trend of improved overall survival, including for the most common types of cancers – breast, prostate, lung and colorectal cancers - continued at least through the end of the last decade, when the latest national statistics were available.

Improvements in health outcomes from cancer and other diagnoses have come at a time of unsustainable growth in health expenditures. This has led to research that more carefully evaluates the relationship between rising costs, on the one hand, and health outcomes, on the other. The goal of the research in chapter 6 is to match changes in survival for cancer matched with changes in spending for those conditions. The study begins by evaluating the relative contribution of changes in diagnosis and treatment to changes in survival after a diagnosis of breast, prostate, lung or colorectal cancer, separately for men and women. Understanding the relative effectiveness of early diagnosis and treatment advances for these cancers is important in

deciding where increasingly limited anti-cancer resources should be allocated. The analysis is applied over two study periods: 1988-1994 and 2000-2004.

The study finds that changes in treatment during these periods may have improved outcomes for some but not all cancers evaluated, and even under favorable assumptions would only be considered cost-effective for a subset of cancers. The years of analysis most corresponding to the recent treatment era did show generally cost-effective medical advances for cancer treatment, with the possible exception of men with lung cancer. The findings suggest that advances in treatment did not lead to improved survival for women during the former period, but did for men in the former, and for both men and women in the latter period. Costs increased for all periods of analysis. The incremental cost-effectiveness of medical progress for cancer varied by treatment era and cohort. In the recent treatment era, we find a consistent positive effect on survival as compared with previously, and lower incremental cost-effective ratios for recent cancer advances.

Chapter 7 focuses on the measurement and interpretation of self-reported disability, as a more specific measure of health. The problem in interpreting self-reported health is that it may be influenced, or biased by people's frame of reference, or by social norms. Most people do not live in social isolation. Instead, they interact repeatedly with family, friends, and neighbors. As a consequence of those pervasive interactions, they allow themselves to be transformed in many ways, a transformation of which they may often be unaware.

One type of transformation involves the formation of social norms about what normal or acceptable behavior might be. These social norms then fix the scales that they may be using in responding to questions about their own behaviors and current situations. If they had different neighbors and friends, their self-descriptions about their lives may well be quite different. While

this may be true within a country where there exists a shared history and culture, it is especially likely to be the case when cross-national comparisons are made, such as between the Netherlands and the United States.

In chapter 7, “Self Reported Disability and Reference Groups,” Arthur van Soest, Tatiana Andreyeva, Arie Kapteyn, and James Smith test the importance of social interactions in people’s self-reported work disability. The research uses data from a household survey representative of the Dutch population. As part of the survey, respondents are asked to evaluate work disability of hypothetical people with some work related health problem, described in vignettes. They are also asked how many people among their friends and acquaintances receive DI benefits. Combining the vignette responses with self-reports on the number of people receiving disability insurance benefits (DI) among one’s friends and acquaintances, the authors estimate a model describing the influence of DI prevalence in one’s reference group on the subjective scale used to report work disability.

The main feature of the model is the notion that response scales for reporting no, mild, or severe work disability, can be affected by a "peer group effect," i.e., by the number of people in the reference group receiving disability benefits. The study finds that DI benefit receipt in one’s reference group has a significant effect on response scales in the expected direction. The effects are sufficiently strong to explain a good deal of the higher rates of self-reported work disability in the Netherlands compared to the United States, as the Dutch population appears to have much more lenient thresholds about what constitutes a work disability.

These findings are suggestive of how policy programs affect social norms. If a policy makes receipt of DI benefits more attractive or easier (e.g., by loosening eligibility requirements) thus increasing the number of DI recipients, this changes social norms. Individuals are now more

likely to label a given health condition as work limiting and the prevalence of self-reported work will rise.

David Cutler's discussion of chapter 7 focuses more intently on the theoretical foundations of why reference groups matter. He emphasizes three hypotheses for why people who have more disabled peers are more likely to describe themselves as disabled. The first is that it lowers the threshold for reporting someone as disabled. The second is that people sort themselves into groups with similar characteristics as themselves. The third is that having more disabled peers makes you feel worse yourself, so that your perception of your own disability changes. While the study emphasizes the importance of the first explanation, Cutler makes the case for follow-up work that more comprehensively differentiates between these hypotheses.

The last three chapters consider in greater depth the relationship between health and economic circumstances, the causal pathways from socioeconomic circumstances to health, the lifelong impact of childhood health on later-life health outcomes, and how wellbeing is affected by financial and macroeconomic conditions.

Throughout the expansive literature in the economics of aging, there is little dispute that the socio-economic status (SES) of individuals is positively correlated with their health status. The size of the body of literature documenting that wealthy and well-educated people generally enjoy better health and longer life is impressive. The robustness of this association is underscored by the fact that the so-called health-wealth gradient has been detected in different times, countries, populations, age-structures, and for both men and women. Moreover, the results are largely insensitive to the choice of SES measures (such as wealth, income, education, occupation, or social class) and health outcomes.

While the existence of the gradient may be uncontroversial, the same cannot be said about its explanation. Medical researchers, economists and other social scientists have developed a large number of competing theories that can broadly be categorized as follows: there may be causal effects from SES to health, causal effects that work in the opposite direction, and unobserved common factors that influence both variables in the same direction without a causal link between the two. A full explanation of the gradient likely encompasses aspects of all three categories of causation, though the relative weighting of explanations remains controversial.

Chapter 8 revisits a 2004 study by Peter Adams, Michael Hurd, Daniel McFadden, Angela Merrill, and Tiago Ribeiro, originally entitled "Healthy, Wealthy, and Wise? Tests for Direct Causal Paths between Health and Socioeconomic Status." Using an innovative research methodology and drawing on the first three waves of the AHEAD survey, that study found that that in an elderly US population, causal channels that operate from wealth to health are an exception rather than the rule. Considering these strong results, as well as the methodological novelty of the approach used, it is not surprising that their work has subsequently been the subject of vivid debate within the literature, focused in particular on the validity of the study's identification strategy in general.

The new research: "Healthy, Wealthy and Wise Revisited: An Analysis of the Causal Pathways from Socioeconomic Status to Health," by Till Stowasser, Florian Heiss, Daniel McFadden, and Joachim Winter, reviews the methodological challenges involved in testing the causal relationships between socioeconomic status and health. It also repeats the earlier analysis using the full range of HRS and AHEAD data that have become available since the original study was completed. This enriches the original study along multiple dimensions. First, the same individuals can be tracked for a longer period of time. Second, the analysis can be extended to

new cohorts of respondents, and the working sample can be widened by including younger individuals aged 50 and older. Third, because of the wider age range, there is variation in health insurance status that is not available in a Medicare-eligible population. To understand which of these data changes contribute to any deviating conclusions, the study does not apply the whole bundle of modifications at once. Instead, the model is estimated multiple times, by applying it to several different data samples, which are gradually augmented along these dimensions.

The new analysis shows that causal inference critically depends on which time periods are used for estimation. Taken together, the estimations with more extended HRS data alter quite significantly the previous study's conclusions about SES causation. Using the information of many (ideally all) waves at once has the greatest effect on results, with many health conditions moving to the column of illnesses for which SES causality may well play a role. Adding younger individuals to the sample has a very similar effect, reducing the number of medical conditions for which the existence of causal links can be statistically rejected even further. This represents a stark contrast to the original findings, where the rejection of structural causality was the most frequent outcome. Given that the greatest changes are triggered by the addition of panel waves, the driving force behind this reversal in results is most likely an increase in test power as sample sizes soar.

Robert Willis's discussion of chapter 8 emphasizes the important bridge the paper provides between the original analysis, the large amount of commentary and critique that resulted, the updated application of the methods, and future work on causal pathways. He emphasizes the complexity of identifying causation, and distinguishing between correlation and causation, particularly when the pathways between health and financial wellbeing have such multi-faceted theoretical underpinnings.

Chapter 9 looks at the sources of later-life health differences between England and the United States. Based on self-reported prevalence of seven important illnesses (diabetes, heart attack, hypertension, heart disease, cancer, diseases of the lung, and stroke), Americans were much less healthy than their English counterparts. These differences were large at all socioeconomic status levels, for both self-reported measures and biological markets, among both men and women, and largely independent of conventional risk factors such as smoking, obesity, and drinking. Much of the US-English difference in later life adult health remains unexplained.

Considerable evidence has emerged that variation in health outcomes at middle and older ages may be traced in part to health and other conditions during childhood. In chapter 9, "Childhood Health and Differences in Late-life Health Outcomes between England and the United States," James Banks, Zoe Oldfield and James Smith analyzes how much of the large differences in illness at middle and older ages in America compared to England can be explained by differences that originated when these people were children and adolescents.

The results suggest that differences in prevalence of childhood diseases between England and the United States and a higher rate of transmission into poorer adult health in the United States do appear to contribute to higher rates of adult illness in the United States compared to England. Based on comparable retrospective questionnaires placed in the HRS survey in the U.S. and the ELSA survey in England, the origins of poorer adult health among older Americans compared to the English traces back right into the childhood years. The transmission rates of childhood illness into poor health in later life also appears to be higher in America compared to England.

A larger question raised by the results is why conditions in America appear to make people of all ages sicker than the English. In his discussion of chapter 9, Amitabh Chandra focuses on this question. He raises a number of possible explanations, ranging from differences in health behaviors, neighborhood or environment effects, differences in health insurance or health care, genetics, stress, socioeconomic circumstances, or socioeconomic disparities. He speculates that the U.S. may do comparatively worse at preventive health care, but comparatively better at treatment, once a health problem arises.

Chapter 10 is the most exploratory in the volume, analyzing daily measures of self-reported wellbeing, and how they respond to financial and macroeconomic circumstances. Since January 2008, the Gallup Organization has been collecting daily data on 1000 Americans each day, with a range of self-reported well-being (SWB) questions. These data provide an opportunity to examine how large changes in the macroeconomic environment affected the emotional and evaluative lives of the population, as well as of subgroups within it. The period covered by these data is characterized by very substantial volatility in financial markets, unemployment, income and macroeconomic performance.

The financial crisis that began in the summer of 2008 saw a rise in the unemployment rate from 4.8 percent in April 2008 to 10.6 percent at its peak in January 2010, a 4.4 percent drop in employee compensation over five months in 2009–10, large stimulus-associated tax credits and rebates, 4.7 percent of personal disposable income in May 2008 and 1.7 percent in May 2009, as well as a collapse and subsequent recovery of the stock market—the S & P 500 Index on March 6th, 2009 had fallen to 40 percent of its all time high of October 2007, and then more than doubled again by end 2010. Through the fall in the market and the fall in the prices of housing and other assets, sixty percent of households saw their wealth decline between 2007 and

2009, and 25 percent lost more than half of their wealth. These declines were widespread, affecting large shares of households across all age, income, and education groups.

These large fluctuations provide an unparalleled opportunity to examine how these events affected the standards of living, the emotional experiences, and life evaluations of those who lived through it. In chapter 10, "The Financial Crisis and the Well-being of America," Angus Deaton analyzes the impact of these economic events using data from the Gallup Healthways Well-being Index (GHWBI), which contains about a million observations on self-reported well-being, as well as on demographics, income, occupation, employment status and numerous health measures. These data allow daily tracking, not only of national averages, but of the outcomes of different groups.

Deaton finds that the wellbeing measures in the survey tracked the stock market to a significant degree. For example, in the fall of 2008, around the time of the collapse of Lehman Brothers, and lasting into the spring of 2009, at the bottom of the stock market, Americans reported sharp declines in their life evaluation, sharp increases in worry and stress, and declines in positive affect. By the end of 2010, in spite of continuing high unemployment, these measures had largely recovered, though worry remained higher and life evaluation lower than in January 2008.

The subjective wellbeing (SWB) measures do a much better job of monitoring short-run levels of anxiety as the crisis unfolded than they do of reflecting the evolution of the economy over a year or two. Even large macroeconomic shocks to income and unemployment can be expected to produce only small and difficult to detect effects on SWB measures.

SWB, particularly evaluation of life as a whole, is very sensitive to question order effects. For example, asking political questions before the life evaluation question reduces reported life

evaluation by an amount that dwarfs the effects of even the worst of the crisis; these order effects persist deep into the interview, and condition the reporting of hedonic experience and of satisfaction with standard of living. Because these order effects are so large and persistent, methods for controlling these effects need to be developed and tested if national measures are to be comparable over space and time.

In his discussion of chapter 10, Daniel McFadden raises the concept of "cognitive bubbles" that may be particular relevant to measures of subjective wellbeing. He suggests that news reports affect the salience of economic conditions (and other aspects of wellbeing), making people feel them more emotionally; and that social networks amplify and compound their impact on subjective wellbeing in the population more generally. A substantial aim of our continuing research effort at the NBER involves further analysis of subjective wellbeing, as another dimension of wellbeing that relates and to some degree incorporates both health and financial circumstances.