Title: Introduction to "Analyses in the Economics of Aging"

Author: David A. Wise

URL: http://www.nber.org/chapters/c10355

Introduction

David A. Wise

This is the tenth in a series of volumes on the economics of aging. The previous volumes were *The Economics of Aging, Issues in the Economics of Aging, Topics in the Economics of Aging, Studies in the Economics of Aging, Advances in the Economics of Aging, Inquiries in the Economics of Aging, Frontiers in the Economics of Aging, Themes in the Economics of Aging,* and *Perspectives on the Economics of Aging.*

This introduction provides a summary of the papers and draws heavily on the authors' own language. Analyses in prior volumes have included discussion of the saving effect and the spread of personal retirement accounts, especially 401(k) plans, and of the effect of defaults on the 401(k) saving decisions of employees. This volume includes introductory work on evaluations of risk in 401(k) plans and continues work on defaults with theoretical discussion of optimal defaults. There are papers on the growth of Medicare cost and on the efficiency of Medicare. Three papers consider different aspects of disability and the consequences of health shocks.

There is a paper on the evolution of, and the relationship between, health, wealth, and living arrangements as people age. The volume also includes papers on saving in other countries compared to the United States. Prior volumes have included many papers on aging issues in other countries. This tradition is continued in this volume with a paper on population aging and the plight of widows in India.

Finally, modern surveys like the Health and Retirement Study and the

David A. Wise is the John F. Stambaugh Professor of Political Economy at the John F. Kennedy School of Government, Harvard University, and the director for Health and Retirement Programs at the National Bureau of Economic Research.

Asset and Health Dynamics among the Oldest Old survey ask respondents to provide estimates of subjective probabilities of survival and other uncertain outcomes. Specific features of the responses limit the usefulness of the data. A paper in this volume provides a method to obtain useable distributions of survival probabilities based on the survey responses.

401(k) Plans: Risk and Optimal Defaults

The shift from employer-provided defined benefit pension plans to employee-controlled personal retirement accounts in the United States has drawn increased attention to the effect of participants' asset allocation decisions on their retirement resources. In "Utility Evaluation of Risk in Retirement Saving Accounts" James M. Poterba, Joshua Rauh, Steven F. Venti, and I begin research to evaluate the implications of asset choices in personal accounts. In this paper we develop a stochastic simulation algorithm to evaluate the effect of holding a broadly diversified portfolio of common stocks compared to a portfolio of index bonds on the distribution of 401(k) account balances at retirement. We compare the alternative distributions of retirement wealth both by showing the empirical distribution of potential wealth values and by computing the expected utility of these outcomes under standard assumptions about the structure of household preferences. Our analysis highlights the critical role of other sources of wealth, such as Social Security, defined benefit pension annuities, and saving outside retirement plans, in determining the expected utility value (based on a constant relative risk aversion utility function) of holding equities in the retirement account.

Given the historical pattern of returns on stocks and bonds, a household that is not highly risk averse would achieve a higher expected utility by holding a portfolio of stocks rather than bonds. We have explored the robustness of this finding to reducing the expected return on corporate stocks by 300 basis points per year. While this shifts the distribution of retirement balances to lower values and reduces the expected utility of holding stocks, we still find that only highly risk-averse investors would choose not to hold corporate stocks.

First, we present "pictures" of the distribution of wealth outcomes for different investment allocation rules. This approach is closely related to the techniques used by many financial planners, who show clients the set of outcomes that they might achieve under a given set of assumptions about future returns and investment strategy. It is also the approach that we, and others, have used in past studies that considered the returns to different investment strategies. Second, we tried to synthesize the information in the distribution of wealth outcomes by computing an expected utility measure corresponding to each distribution. This approach allows for the possibility that the marginal utility of wealth declines with wealth, so that a given

increment to wealth is more valuable when wealth is at a low level than when it is high. We conclude that both approaches to the evaluation of risk can be important and that the appropriateness of each is likely to depend on the specific goal of the evaluation.

Given the apparent inconsistencies between standard behavior and the predictions of standard utility functions, we hope to gain a better understanding of individual preferences over uncertain levels of future retirement assets by developing a set of survey questions designed to elicit respondent preferences over alternative wealth outcomes. We hope to include these questions on household surveys like the Health and Retirement Study and other surveys.

In the previous volume in this series, James J. Choi, David Laibson, Brigitte C. Madrian, and Andrew Metrick reported evidence on the effect of plan defaults on 401(k) participation and contribution rates. Based on this and other research, they conclude that defaults have an enormous effect on employee choices. In "Passive Decisions and Potent Defaults" in this volume, they develop a theory of optimal defaults based on these considerations: Defaults matter because opting out of a default is costly and these costs change over time, generating an option value of waiting. In addition, people have a tendency to procrastinate.

In a world of saving preference, the authors conclude that it may sometimes be optimal to set *extreme* defaults that are far away from the mean optimal savings rate. This effect arises for two reasons. First, a default that is far away from an employee's optimal savings rate may make that employee *better off* than a default that is closer to the employee's optimal savings rate. Intuitively, if an employee suffers from a procrastination problem, then a "bad" default—that is, one that is far from the consumer's optimal savings rate—will be more motivating than a better default. Hence, sometimes bad defaults make people better off than better but imperfect defaults. Second, their theory implies that optimal defaults are highly sensitive to the actual distribution of optimal savings rates. In particular, optimal defaults are often associated with the modal optimal savings rate and not the mean optimal savings rate. Since these modes are sometimes extreme (e.g., minimum or maximum contribution rates), optimal defaults will sometimes be extreme as well.

At the end of their paper Choi, Laibson, Madrian, and Metrick calibrate their model and use it to calculate optimal defaults for employees at four different companies. For two of these companies, the optimal default is close to the mean optimal savings rate, whereas for the others the optimal defaults are extreme: 0 percent and 15 percent, respectively. Their analysis suggests that optimal defaults are likely to be at one of three savings rates: the minimum savings rate (0 percent), the employer match threshold (typically 5 percent or 6 percent), or the maximal savings rate (around 15 percent in their sample of companies from the late 1990s).

Medicare Cost Growth and the Efficiency of Medicare

The increasing cost of Medicare is an ongoing concern. In "Characterizing the Experience of High-Cost Users in Medicare" Thomas MaCurdy and Jeff Geppert provide new insights into the source of cost growth over the past decade. Their study reveals several valuable insights into the growth of Medicare expenditures in recent years, both in projecting aggregate trends and in discovering the extent to which the concentration of spending on high-cost users contributes to overall expenditures. Considering the growth in Medicare program payments from 1989 to 1999, they attribute 20–30 percent of total growth to an increase in the participation rate, 50–60 percent to an increase in average program payments per service recipient, and the remainder to higher enrollment.

In sharp contrast to the first half of the 1990s, total Medicare costs fell in the late 1990s. Whereas the lower percentiles of the expenditure distribution continued to increase throughout the period, expenditures fell for the highest-cost users for Medicare services. Published statistics extending beyond their sample period suggest that this was only temporary; starting in 2000, the overall growth in Medicare expenditures reverted to its previous rate and may have accelerated. Annual Medicare spending is highly concentrated among a small segment of the beneficiary population, and the share of spending attributable to high-cost users has remained remarkably stable over the 1989–99 decade even though growth rates have varied considerably during the period. Those beneficiaries classified in the top 2 percent of the annual expenditure distribution account for about one quarter of total expenditures, and those in the top 5 percent cover almost half of annual expenditures. Considering spending by months only reinforces this picture of concentration. The top 2 percent of months with spending during a year account for around two-fifths of total annual expenditures, and the top 5 percent of months cover nearly two-thirds of yearly Medicare expenditures. While high-cost episodes account for a large proportion of total spending during any year, the majority of the elderly experience such episodes at some point over a decade, implying far less concentration in expenditures when viewed over lifetimes. Three-fifths of beneficiaries experience at least one 95th-percentile month in the decade, and two-fifths realize one or more 98th-percentile months. Knowledge of the incidence of 95th-percentile months alone explains nearly two-thirds of Medicare spending over the decade, and spending accumulated for those in the 98th-percentile months comprises almost four-fifths of total decade expenditures.

MaCurdy and Geppert emphasize that a major challenge faced by researchers involves identifying intense users of health care services along with the factors leading to the incidence and duration of their utilization. Such a task requires a detailed understanding of the patterns of medical care for periods much shorter than a year to adequately capture the onset

of health events and relationships linking the persistence of costs in both the short and long runs. Many surmise that restricting attention to those in their last year of life identifies most high-cost users, but existing work shows that the majority of the most intense users measured by accumulating expenses annually live beyond a year after incurring their high expenses. Further, previous work has met with limited success in associating large fractions of the highest-cost users with particular diagnoses or chronic conditions. The availability of the monthly longitudinal Medicare data greatly enhances options for improving our understanding of the sources of high-cost utilization.

The United States spends more on health care in per capita terms and as a percentage of gross domestic product (GDP) than any other developed country. This can be interpreted in two ways. One is that the elevated spending is symptomatic of failure in the health care system. Money is wasted through administrative overhead, the overuse of fully insured health care, or the provision of expensive tertiary but only marginally useful technology. A different view is that U.S. citizens demand, and get, a higher quality level of health care than anywhere else in the world. It may be expensive, but the technological advances provided in the United States have led to dramatic improvements in functioning and life expectancy. In "The Efficiency of Medicare," Jonathan Skinner, Elliott S. Fisher, and John E. Wennberg use the idea of "natural randomization" to evaluate the efficiency of the Medicare program.

The authors attempt to test whether the Medicare program is broadly consistent with the efficiency criterion commonly used in public economics, where the marginal social value of the last dollar spent on specific types of health care (in each region) is equal to the marginal social benefits of the dollar that could have been spent for other worthy causes. They used data on survival rates, Medicare expenditures, and health status measures across 306 hospital referral regions in the United States to test these hypotheses. Their best estimate of the incremental value of Medicare spending with regard to effective care suggests that spending for these types of services is too low, especially considering how this type of care is associated with overall Medicare expenditures. On the other hand, the supply-sensitive dimension of care is a major factor in explaining overall Medicare expenditures—roughly 20 percent annually—but does not show any impact in terms of improving survival rates across regions. These results, the authors conclude, suggest that the inefficiency inherent in the Medicare program is as much as 20 percent of total Medicare expenditures.

A Focus on Disability and the Implications of Health Shocks

Substantial recent evidence shows a reduction in disability among the elderly in the United States, on the order of 25 percent in the past two

decades. The major issue raised by these findings is why disability has declined. In "Intensive Medical Technology and the Reduction in Disability" David M. Cutler investigates the role of intensive medical technologies in the decline in disability. Using data from the National Long-Term Care Survey, he documents that increased use of intensive procedures might be associated with some reduction in disability but probably does not account for the majority of the decline.

Cutler presents preliminary evidence on the role of increased medical procedure use in response to acute episodes of disease. These procedures consist largely of major surgical operations—open-heart surgery for people with cardiovascular disease and hip replacements for people with fractures or severe arthritis, for example. Use of these procedures has diffused widely, in many cases doubling or tripling over a decade, suggesting they could play a large role in improved health.

Cutler documents two facts about procedure receipt and disability change. He shows that most of the reduction in disability is not from people having fewer disabling conditions. The share of people with a stroke, fracture, or other serious condition has increased over time. Rather, reduced disability is because fewer people who have these conditions become disabled. Since intensive medical care is most important after a person has an acute event, this suggests the potential role of increased utilization of medical services in reducing disability.

Examining specific technologies shows that receipt of intensive procedures is associated with some, but not a large, reduction in disability. People with musculoskeletal problems and circulatory diseases are much more likely to get surgery now than in the past, and disability for people with these conditions has fallen. But this does not account for a large share of the total decline. Other medical and nonmedical interventions are more important in aggregate. Cutler speculates about what other factors might be relevant, but he does not provide conclusive evidence.

Self-reported health status (SRHS) is an imperfect measure of nonfatal health, but it allows examination of how health status varies over the life course. Although women have lower mortality than men, they report worse health status up to age 65. The SRHS of both men and women deteriorates with age. There are strong gradients, so that at age 20, men in the bottom quartile already report worse health than do men in the top quartile at age 50. In the bottom quartile of income, SRHS declines more rapidly with age, but only until retirement age. These facts motivate Anne Case and Angus Deaton to study the role of work, particularly manual work, in health decline with age, as reported in their paper, "Broken Down by Work and Sex: How Our Health Declines."

Case and Deaton start from the observation that SRHS worsens with age and that it does so much more rapidly among those at the bottom of the income distribution, who also start their working lives with lower

health. They originally expected that, because manual work involves more wear and tear on the body, the health of manual workers would decline more rapidly than that of nonmanual workers, thus offering an explanation for their starting facts. The data from the National Health Interview Survey (NHIS) show that the health of manual workers does in fact decline more rapidly during the working years than does the health of nonmanual workers, in spite of the existence of health-based selection out of manual work, which artificially inflates the health of those who remain. The authors do not find this result at all implausible. Instead, the implausibility lies in the health repair technology that is routinely assumed in the health economics literature.

The authors find that manual workers have worse health than do nonmanual workers, and although their health declines more rapidly, the major factor accounting for the differences in health and health decline in different parts of the income distribution is whether or not people are in the labor force, a mechanism where causality runs from health to income, not the reverse. Even so, both income and education have independent protective effects on health for those who are in work, and these effects are reduced but not eliminated by controlling for occupation. With only a few exceptions, Case and Deaton find a marked similarity in all of these results between men and women.

Considerable analysis has been directed to understanding and disentangling the multiple ways in which socioeconomic status (SES) may influence a variety of health outcomes. Much less attention has been directed to the impact health may have on SES. In "Consequences and Predictors of New Health Events," James P. Smith aims to estimate the effect of new health events on a series of subsequent outcomes that are both directly and indirectly related to SES. These outcomes include out-of-pocket medical expenses, labor supply, health insurance, and household income.

The analysis is based on the HRS and AHEAD surveys and thus considers persons first observed when they were between ages fifty-one and sixty-one. The author concludes that among people in their preretirement years, feedbacks from health to labor supply, household income, and wealth are realities that should be neither ignored nor dismissed as of secondary importance. Working is the critical link in this chain, with out-of-pocket medical expenses of lesser importance. The negative income and wealth consequences of new health shocks appear to decay with age and are much smaller in an already retired population. What these consequences would be ten or twenty years earlier in age is an important and yet unanswered question. The evidence in this paper, along with that available in other studies, shows that we can say with confidence that health has quantitatively strong consequences for several dimensions of SES, particularly financial, in certain age groups.

Smith observes that only tentative conclusions are warranted with re-

spect to the effect of SES on the future onset of disease. Household income never appears to predict any future onset over the horizon of about a decade. However, even after controlling for an extensive list of baseline health conditions, education is a strong predictor if the future onset of disease.

Health, Wealth, and Living Arrangements

Health, wealth, and where one lives are important indicators of living conditions. In "Healthy, Wealthy, and Knowing Where to Live: Trajectories of Health, Wealth, and Living Arrangements among the Oldest Old," Axel Börsch-Supan, Florian Heiss, and Michael D. Hurd investigate the joint evolution of these three conditions as persons age. The elderly reach their early postretirement years in an initial status that is characterized by housing wealth, nonhousing bequeathable wealth, annuity income, health status, and family connections. The broad goal of the paper is to describe the trajectories of health, wealth, and living arrangements as people age; to understand how the trajectories of health status, wealth position, and living arrangements are interrelated; and to be able to predict how health and living arrangements would evolve when economic and other conditions change.

The authors reproduce the finding that wealth and health are strongly related to each other. Wealthier persons live longer and are healthy longer. This interaction is moderated by where elderly persons live. Remaining in the lowest wealth quartile is most likely when an elderly person lives in a nursing home, and it is least likely when this person lives with others. The reverse pattern is true for the probability of remaining in the highest wealth quartile. For the initially wealthy, living with others decreases the expected future wealth relative to living alone. This is consistent with the notion that cohabitation implies intrafamily transfers to the needy. For those who remain living independently, home ownership declines, but the speed of reduction is slower than it is for financial wealth. The authors emphasize that the results in the paper are descriptive and imply no causality. Further research will apply more sophisticated econometric methods to attempt to identify patterns of causality.

Saving: The United States Compared to Other Countries

In addition to relying on public retirement provisions, households prepare for retirement through tax-sheltered and after-tax savings. They may invest these funds in a wide variety of assets, including housing, stocks, bonds, savings accounts, and so on. These asset types differ in their risk, return, and liquidity characteristics as well as in their fiscal treatment. Economic theory postulates that households allocate their portfolios accord-

ing to their risk aversion, time horizon, uncertain out-of-pocket medical expenditures, income risk, informal (family) risk sharing arrangements, and the like. Both testing and quantification of the theory are hampered by the fact that important variables do not exhibit sufficient variation within a country to establish their relative importance for portfolio choice or, more generally, for retirement saving and investment. This gap is partially filled by Arie Kapteyn and Constantijn Panis in "Institutions and Saving for Retirement: Comparing the United States, Italy, and the Netherlands."

In this paper, they consider retirement saving and portfolio choice in three countries with widely varying institutional arrangements for retirement income: the United States, Italy, and the Netherlands. The authors emphasize that the number of countries considered and the data are limited. Nonetheless they propose these implications:

- Americans should save more for retirement than the Dutch or the Italians.
- Americans should save more due to more exposure to uninsurable income and consumption risk.
- Italians should save more due to severe borrowing constraints in their country.
- The Dutch should have relatively low stockholdings due to the low level of private wealth.
- Stock ownership in the United States should be higher than in Italy because of the more developed capital markets in the United States.

The authors emphasize that while each of these implications is borne out by the data, it is in general not possible to establish the relative magnitude of factors influencing wealth accumulation or portfolio choice. For instance, both low replacement rates at retirement and higher consumption and income risk in the United States imply that Americans should save more than Europeans.

In "Household Saving in Germany—Results of the First SAVE Survey," Axel Börsch-Supan and Lothar Essig take a fresh look at the saving behavior of German households. They use the first wave of the SAVE panel. It is a preliminary look, since many aspects of saving can only be understood using longitudinal data. The authors take this paper on German saving behavior as both a comparison with and a contrast to the large literature on the saving behavior of U.S. households.

Germany is an interesting country in which to study household saving behavior since it appears to contradict the familiar textbook version of the life-cycle theory of consumption and saving. First, there is little borrowing by young households. More striking is that nearly everyone—whether in the middle income bracket or richer—saves substantial amounts at older ages. Only in households that earn less than 25 percent of average income between the ages of 60 and 75 is there dissaving.

Overall, the findings show a savings pattern that is extraordinarily stable. Germans save regularly, in a manner that is planned and often with a clearly defined purpose in mind. German households appear not to save in order to balance out transitory income fluctuations. Rather, they appear to save out of income components that are stable in the long run. German labor income has less individual variation than labor income in the United States. This should reduce the precautionary savings motive, all else being equal, relative to the United States. In addition, German public pension replacement rates are much higher than those of the U.S. replacement rates. This should reduce the savings motive for old-age provision relative to the United States. The findings on German savings motives, however, contradict these predictions: The authors find that precaution and old-age provision are the two most important savings motives in Germany. Less developed credit markets may explain the high saving rate relative to the United States.

Aging Issues in India

The elderly are a very large and rapidly growing population in India. The well-being of widows in particular has attracted considerable attention. In his paper, "Caste, Culture, and the Status and Well-Being of Widows in India," Robert Jensen pursues two issues. The first is how norms, attitudes, and practices vary across groups in India and how they affect the well-being of the elderly. He first examines the influence of caste on well-being. Caste is a social organization prevalent in India that creates a well-defined social ordering. The objective is to examine the consequences for widows, using indicators for individual nutritional status and health, rather than relying on household income or expenditure per capita. The second issue is whether social and cultural institutions and the status of widows have any underpinnings in economic factors. In particular, Jensen explores whether the production of crops for which women and the elderly are able to make larger economic contributions leads to improved status of widows.

Jensen draws two primary lessons from his paper. The first lesson is that issues of intrahousehold allocation are essential for assessing individual well-being, especially for the elderly. While this has been widely appreciated in economics for some time, very little is done about it in practice, and most studies focus only on household per capita measures in assessing living standards. This observation also has implications for studies on the relationship between SES and health. Further, the results also show that the relationship between SES and health or nutrition in India is more complex than simply the purchasing power potentially implied by income or expenditure. In particular, other factors, such as the treatment of individuals within the household, mediate this relationship. Jensen finds that widows are much better off in forward-caste households when measured in terms

of per capita expenditure, but in using body mass index (BMI) as a crude proxy for consumption of the elderly, forward-caste persons are no better off than lower-caste households, suggesting the share of household resources is not measured well by expenditure per person.

The second lesson is that the status, treatment, and well-being of widows have a foundation in potential economic value, either through bargaining power within households or through a cultural underpinning to the evolution of cultural norms. The claim is not that economic factors are the only, or even the largest, determinant of "culture" or to argue that factors such as history and lineage are not important. However, the evidence indicates that economic factors appear to play at least some role in the well-being and status of widows. The implication is that programs (such as micro-enterprise ventures) that expand economic opportunities for women or the elderly attempt to minimize age or gender discrimination in private-sector employment, or gender- and age-sensitive hiring schemes for public projects, especially in places where the state employs a significant number of people, may improve women's and widows' status.

Using Survey Responses on Subjective Probabilities

Many economic models are based on the forward-looking behavior of individuals. Although it is often said that expectations about future events are important in these models, it is the probability distributions of future events that are contemplated in the models. For example, an individual's consumption and saving decisions are assumed to depend on future interest rates, the likelihood of dying, and the risk of substantial future medical expenditures. According to typical theories, individuals have subjective probability distributions about these and other events and use these subjective distributions to make decisions about their saving practices. In "Individual Subjective Survival Curves," Li Gan, Michael D. Hurd, and Daniel McFadden use responses from HRS and AHEAD survey data to estimate individual subjective survival probabilities.

The HRS and AHEAD surveys ask individuals for their expectations on the probability of given future events. On average, the subjective probability of a future event is consistent with the observed probability that the event does occur. For example, in general, the averages of individual survival probabilities reported by survey respondents are consistent with those from life tables.

However, at the individual level, the subjective probability responses in HRS and AHEAD suffer serious problems of focal responses at 0.0 and 1.0. Consequently, applications of subjective probabilities will be extremely limited if "true" subjective survival probabilities are not recovered.

In this paper, the authors suggest a Bayesian update model to account for problems caused by focal responses of 0.0 and 1.0. As a result, individ-

ual survival curves derived from the model do not suffer the problems of focal responses. The authors also propose two approaches to model individual heterogeneity in subjective survival curves. One approach modifies the life table hazard rates while another approach models the subjective aging process, which is different from the life table aging process. The model is estimated from the observed survival information in the sample the authors use. From the estimated model, they construct several optimistic indexes for each individual and conduct a test that is based on out-of-sample prediction. These optimistic indexes are used to create individual subjective survival curves that have considerable variation and are readily applicable to economic models that require individual subjective survival curves.