Volume Title: Research Findings in the Economics of Aging

Volume Author/Editor: David A. Wise, editor

Volume Publisher: The University of Chicago Press

Volume ISBN: 0-226-90306-0

Volume URL: http://www.nber.org/books/wise08-1

Conference Dates: May 10-13, 2007

Publication Date: February 2010

Chapter Title: New Age Thinking: Alternative Ways of Measuring Age, Their Relationship to Labor Force Participation, Government Policies, and GDP

Chapter Author: John B. Shoven

Chapter URL: http://www.nber.org/chapters/c8188

Chapter pages in book: (17 - 31)

New Age Thinking
Alternative Ways of Measuring Age, Their Relationship to Labor Force Participation, Government Policies, and GDP

John B. Shoven

This chapter is not about what you think it is. It is about how to measure age. I argue that there are better alternatives to the standard measure of years-since-birth. In fact, I claim that public policy would be better if age were more appropriately specified in the law. A particularly simple alternative to years-since-birth would be a measure of age based on mortality risk. Groups whose mortality risk is high would be considered old, those with low mortality risk would be classified as young, and those with the same mortality risk would be considered to be the same age. Another closely related approach would be to measure age from the other end of life, at least in expected terms. That is, remaining life expectancy (RLE) would be the measure of age—those with a short RLE would be considered elderly and those with a long RLE would be considered young. One advantage of the RLE approach is that it is measured in years, units that are widely understood, unlike mortality risk, which is measured in the percentage chance of dying within a year.

Even at a point in time, there are differences between the various ways of measuring age. For instance, RLE and mortality risk would reflect that a man at the conventionally defined age of sixty-five is roughly the same age as a woman of age seventy. The real differences between the proposed mortality-based measures and the conventional years-since-birth measure

comes when comparing populations at different points in time, such as comparing the 1965 population with the 2007 population or the projected 2050 population. The different measures will, for instance, give a very different answer to how many elderly people there will be in 2050. Later in the chapter we will look at how these various ways of measuring age would apply to labor force participation and also how different old-age dependency ratios might look under the alternative approaches.

To the best of my knowledge, there is not a large existing literature on alternative ways of measuring age. The paper that contains ideas most similar to mine was written by my Stanford colleague Victor Fuchs (1984). In his paper, Victor discusses using remaining life expectancy as a better measure of age and noted that when Social Security was designed in 1935, the gender-blended remaining life expectancy at age sixty-five was 12.5 years. By 1984, those who had 12.5 years of remaining life expectancy were seventy-two years of age. Victor went on to say that if sixty-five was the appropriate entry age for being categorized as elderly in 1935, then the entry age for that status should have been seventy-two in 1984. Thus, Victor already had the idea of an alternative measure of age and suggested that "nominal ages" could or perhaps should be adjusted to "real ages" based on mortality or remaining life expectancy.

Another paper that is closely related is Cutler and Sheiner (2001). The authors are concerned with the impact of demographic changes on medical spending both in the past and in future projections. They note that for acute care and nursing home care, demand is more a function of remaining life expectancy than it is of age. They also note the high medical costs associated with the last year of life and that, on average, the last year of life has been occurring at older and older ages. They do not quite reach the conclusion that I have—namely, that age itself could be defined as something other than years-since-birth—but their analysis suggests the need for a new measure of age.

In order to get started, figures 1.1 and 1.2 introduce the concept of mortality milestones—the first age at which men and women reach 1, 2, and 4 percent mortality risk. Figure 1.1 shows that in the year 2000, men first reached a mortality risk of 1 percent at age fifty-eight, they first reached a 2 percent mortality risk at age sixty-five, and they reached the 4 percent milestone at age seventy-three. The corresponding ages in 1970 were fifty-one, fifty-nine, and sixty-eight. The figure says that fifty-one-year-olds in 1970 and fifty-eight-year-olds in 2000 had the same mortality risk (1 percent), fifty-nine-year-olds in 1970 and sixty-five-year-olds in 2000 similarly had the same mortality (2 percent), and sixty-eight-year-olds in 1970 had the same mortality risk as seventy-three-year-olds in 2000. In just the thirty years between 1970 and 2000, the age at which 1 percent mortality is reached advanced seven years, the age at which 2 percent mortality risk is reached advanced six years, and the age at which 4 percent is reached advanced five

Fig. 1.1   Age of mortality milestones for men, 1940–2000
Note: Sixty-five-year-olds in 2000 had the same mortality risks as fifty-nine-year-olds in 1970.

Fig. 1.2   Age of mortality milestones for women, 1940–2000
Note: Sixty-three-year-olds in 2000 had the same mortality risks as fifty-nine-year-olds in 1970.

years. Any way you look at it, there was remarkable progress in age-specific mortality. A mortality-based age system would suggest that fifty-nine-year-old men in 1970 and sixty-five-year-old men in 2000 were the same age.

Figure 1.2 has the corresponding information for women. Two things are immediately apparent. First, women at the same number of years since birth are effectively younger than men of the same conventionally defined

age. In 2000, women first reach a 1 percent mortality risk at age sixty-three (compared with fifty-eight for men), a 2 percent mortality risk at seventy (compared with sixty-five for men) and a 4 percent mortality risk at age seventy-eight (compared with seventy-three for men). The mortality risk approach to measuring age would have seventy-year-old women in 2000, sixty-five-year-old men in 2000, and fifty-nine-year-old men in 1970 as all being the same age.

The measurement of age with different measures is not like choosing between measuring temperature on a Fahrenheit or Centigrade scale. The connection between the two temperature measures is linear and constant through time. In a very real sense, it does not matter which scale you use. However, the relationships between the different ways of measuring human age change over time and some apparently important phenomenon are primarily due to a particular method of age measurement. For instance, it is reasonably well-known that Medicare spends more on men than it does on women of the same age. The difference (being of the order of 30 percent) is not small. But, this result is a function of how age is measured. Medicare spends roughly the same amount on men and women with the same mortality risk or with the same remaining life expectancy (Shoven 2004). Of course, the fact that reconciles these observations is that seventy-five-year-old women are younger than seventy-five-year-old men, at least according to an age system based on mortality risk or remaining life expectancy.

Figures 1.1 and 1.2 also show that the rate of mortality progress was somewhat slower for women than for men, at least for the last thirty years of the twentieth century. The age at which women first reach a 1 percent mortality risk went up four years between 1970 and 2000 (versus seven years for men), the age at which mortality risk reaches 2 percent advanced two years for women (versus six for men) and the age where 4 percent mortality is "achieved" advanced three years for women (versus five years for men).

Figures 1.3 and 1.4 illustrate the same phenomenon slightly differently. They show mortality risk by years-since-birth for men and women in 1965 and 2005. Once again, we see that there was more mortality progress for men than women over this period. One way to look at it is the amount you would have to shift the 1965 curve to the right in order for it to overlap the 2005 curve. If you do it so that the curves match at roughly 3.5 percent for men (about at age sixty-five in 1965), then the required shift is about seven years. This is similar to the fact we saw in figure 1.1, but here we learn that seventy-two-year-old men in 2005 had about the same mortality risk as sixty-five-year-olds in 1965. The mortality curve for women would need to shift to the right far less to coincide between 1965 and 2005. The mortality risk of sixty-five-year-old women in 1965 was about 1.75 percent, roughly the same as sixty-nine-year-old women in 2005. So, at these ages we see that men in 2005 are effectively seven years younger than someone of the same age in 1965, whereas women are effectively about four years younger.

Fig. 1.3   Male mortality risk by age in 1965 and 2005, age fifty-five through seventy-nine

Fig. 1.4   Female mortality by age in 1965 and 2005, age fifty-five through seventy-nine

There are plenty of other demographic statistics we could look at regarding the measurement of age, but I will simply present one more relationship, the relationship between remaining life expectancy and mortality risk for men and women. The data are plotted in figures 1.5 and 1.6. You might have thought that the relationship between mortality risk and remaining life expectancy would be pretty stable across time (I did), since they are

Fig. 1.5   Remaining life expectancy by mortality risk for males
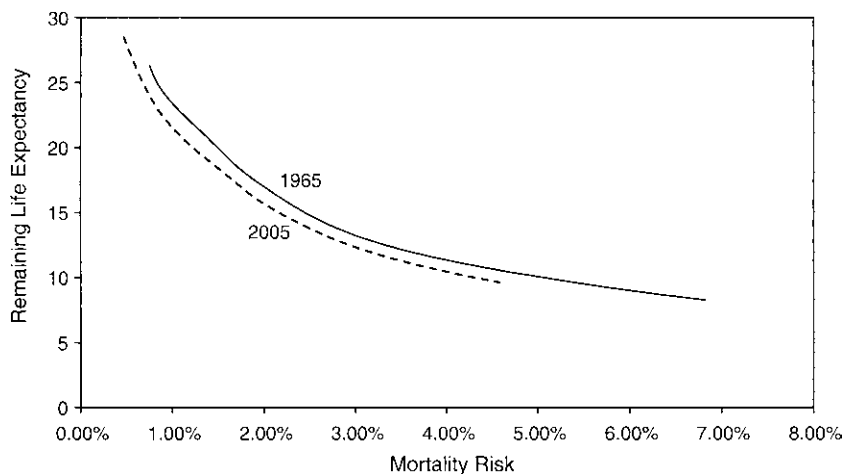
Fig. 1.6   Remaining life expectancy by mortality risk for females

alternative mortality-related measures of age. However, the figures show that the relationship has changed considerably in the forty years between 1965 and 2005. The basic pattern has a positive interpretation. The charts are drawn for people between age fifty-five and seventy-nine, just like for figures 1.3 and 1.4. At least for this age range, for any given remaining life expectancy, people had a lower mortality risk in 2005 than they did in 1965. For instance, men with a fifteen-year remaining life expectancy in 1965 had about a 3.00 mortality risk, whereas such a man in 2005 had about a 2.45

percent mortality risk. For women with fifteen years of expected remaining life, the corresponding mortality risks were 2.45 percent in 1965 and 2.15 percent in 2005. At least to me, this suggests that even with the same remaining life expectancy, people are healthier in 2005 than they were in 1965. This is consistent with a mild squaring of the survival curves and a concentration of high mortality in the last years of life.

1.1   How Much Aging Will Occur in the U.S. Population between Now and 2050?

There probably is as much attention paid to the anticipated aging of the U.S. population and how the economy will adjust to it as any demographic fact. Some of the predictions of the aging of the population are simply due to the use of the conventional years-since-birth measure of age. Consider two alternative definitions of who is elderly in the population, those who are currently sixty-five or older, and those who have a mortality risk of 1.5 percent or worse. Today, at least on a gender blended basis, the two definitions of elderly are equivalent, since the average mortality risk faced by sixty-five-year-olds is 1.5 percent. However, going forward being sixty-five and over and having a mortality risk of 1.5 percent and over will not be equivalent. Figure 1.7 tells the story. The Census Department predicts that the sixty-five and over population will increase from about 12.5 percent of the population today to about 20.5 percent between 2035 to 2050. In 2050, the Census predicts that the percent of the population that is elderly will continue to gradually increase. On the other hand, the percent of the population with mortality risks higher than 1.5 percent (currently also 12.5


Fig. 1.7   Elderly as a percent of the U.S. population, 2000 to 2050

percent of the population) never gets above 16.5 percent and is projected to be just slightly below 15 percent and declining by 2050. With the sixty-five and over criterion, the fraction of the population that is classified as elderly is projected to grow by approximately 66 percent by 2050; whereas with the 1.5 percent and above mortality criterion, the fraction of the population classified as elderly is projected to grow by only 20 percent. The point is the great aging of our society is partly a straightforward consequence of how we measure age. Another interpretation of figure 1.7 is that by 2050 there will be approximately 6 percent of the population that are over sixty-five years of age but who are young enough to have a mortality risk of less than 1.5 percent. By the standard criterion used today they would be classified as elderly but by any mortality-based definition of who is elderly they would not. This naturally leads to the topic of labor force participation by age.

1.2   Labor Force Participation

There have been significant changes in labor force participation by age over the past forty years. I am going to concentrate on male labor force participation because the dramatic increase of women in the workforce masks to some degree what is going on in the retirement behavior of women. The conventional graph of male labor force participation by age is shown in figure 1.8. There was a dramatic decrease in labor force participation by age between 1965 and 1985 with the labor force participation at age sixty-two falling from 79 percent to 51 percent and at age sixty-five falling from 56 percent to 30 percent. The change between 1985 and 2005 was less dramatic, with labor force falling slightly from ages fifty-five to sixty-one, but rising


Fig. 1.8   Male labor force participation by age, fifty-five to seventy

somewhat from sixty-two through seventy. On average, men were retiring almost three years earlier in 2005 than they were in 1965.

If we look at male labor force participation by age with one of the two mortality-based definitions of age, we get a somewhat different picture. Figures 1.9 and 1.10 plot male labor force participation by mortality risk and by remaining life expectancy, respectively. First, the figures show that with either mortality-based age measure, the fact that men are working more at older ages essentially disappears. In figure 1.9, we see that men of a given

Fig. 1.9   Male labor force participation by mortality risk

Fig. 1.10   Male labor force participation by remaining life expectancy

mortality risk have dramatically lower labor force participation in 2005 than in 1965, and that the shift from 1985 to 2005 was generally toward lower labor force participation. Figure 1.10 perhaps is the most revealing. While the conventional chart (figure 1.8) shows that men are retiring two-and-a-half to three years earlier than they did in 1965, figure 1.10 displays the more interesting fact that men are living roughly six years longer in retirement than they did in 1965. The six-year figure comes from noting that labor force participation was 50 percent in 1965 for men with a RLE of thirteen years, but in 2005 it was 50 percent for men with a RLE of nineteen years. Those that exit the labor force relatively early are leaving with even more than six extra years of remaining life relative to their counterparts in 1965, whereas those that exit relatively late are leaving with an extra five years of life expectancy. The overall average is that the RLE of male retirees increased six years over the forty years between 1965 and 2005 for an overall increase in the expected length of male retirement of nearly 50 percent. This percentage increase in the length of male retirement is right in the data, but it is not a well-known fact. Such a dramatic increase in the length of the average retirement has quite a bit to do with the financial strains faced by Social Security and defined benefit pension plans. Providing for a nineteen- or twenty-year retirement with a thirty-five- or forty-year career is much more difficult than providing for a thirteen-year retirement. Unless retirement ages begin to adjust with RLE, today's young people could spend 40 percent of their adult life out of the workforce.

All of the increase in life expectancy of adult men in the twentieth century was taken as retirement and not work. The expected length of retirement of men increased from approximately two years in 1900 to about nineteen years in 2000. It appears to be financially impossible that the same allocation of increased life expectancy to continue in the twenty-first century. However, pension laws and programs feature lots of conventionally defined ages that have not been adjusted for improvements in mortality and life expectancy. For instance, the 59.5 age after which money can be withdrawn from tax deferred retirement accounts has not changed since it was introduced decades ago. Similarly, the age of early eligibility for Social Security (sixty-two), the age of Medicare entitlement (sixty-five), and the age that one must begin withdrawing from tax-deferred saving accounts (70.5) have not changed in at least the past forty years, if ever. These critical ages will likely need to be adjusted if we expect much of the increase in life expectancy in the twenty-first century to be devoted to work instead of retirement.

In a book I wrote with George Shultz (Shultz et al. 2008), we calculate the difference in the total labor supply in the United States in 2050 between two scenarios: (a) people retire in the same pattern as they do today by conventionally defined ages, and (b) people retire with the same lengths of retirement as they do today; that is, with the same remaining life expectancy at the time of retirement. Obviously, the difference between retiring at the same

ages and retiring with the same retirement lengths depends on the amount of mortality progress between now and 2050. If we use the official Census forecast for mortality improvement, then the total labor supply (in aggregate hours of work) would be about 9.6 percent higher if labor force participation stayed constant relative to remaining life expectancy (that is, the 2050 graph in figure 1.10 looks like the 2005 one) than it would be if labor force participation stays constant in terms of age (if the 2050 curve in figure 1.8 were to look like the 2005 one). We do not think of this 9.6 percent number as precisely estimated by any means—it might be 8 percent or it might be 10 percent. On the other hand, 9.6 percent is our best estimate and we think that an increase in the size of the labor force by such a magnitude is rather enormous. The estimate takes account not only the change in the labor force participation by age under the two scenarios, but also takes account of the number of hours worked per week at different stages of the life cycle. One way to think about it is that in one scenario, all labor force behavior (both work week and retirement) remains constant as a function of years since birth between now and 2050 and in the other scenario, all labor force behavior remains constant as a function of age. But in the second scenario, age is defined as remaining life expectancy. A simple Cobb-Douglas aggregate production function would suggest that 9.6 percent more labor would result in about 7 percent more gross domestic product (GDP), even if the extra labor was not accompanied with a larger capital stock. If investment were correspondingly higher so that the capital stock was also 9.6 percent higher in 2050 under the second scenario, then GDP could also be 9.6 percent higher. These figures of an extra 7 to 10 percent of GDP are worth pursuing, particularly given the forecasts of how much more we will be spending as a society on health care by 2050. The question is how to encourage people to balance work and retirement relative to their age and guide them on how to think of age.

There are lots of policies within Social Security, Medicare, and the tax law that actually discourage long careers. In a paper coauthored with Gopi Shah Goda and Sita Nataraj (2007), I analyzed three changes in Social Security that would level the playing field with respect to career length. Currently, Social Security counts only the highest thirty-five years of indexed earnings in computing the initial monthly benefit for someone commencing benefits. With each year of work for the first thirty-five years, the year's earnings replace a zero in the benefit calculation. Once an individual has worked for thirty-five years, additional years either replace earlier nonzero indexed earnings or they do not count at all, because they are lower than the previous best thirty-five. In all cases, the thirty-sixth, thirty-seventh, and thirty-eight years (etc.) count far less than the thirty-third, thirty-fourth, and thirty-fifth years of work. The modal age at which men reach thirty-five years of covered service in Social Security is fifty-two. That is an incredibly young age to reduce or eliminate the connection between additional work and

contributions and higher benefits. So, one possible reform would be to raise the thirty-five years that go into the benefit formula to forty years and then index the number of years for improvements in life expectancy. A second reform would be to create a new category of workers, those that are "paid up" in terms of Social Security and Medicare contributions. After completing the years that count (forty under reform one plus whatever increase comes from further increases in life expectancy), workers would achieve this new paid-up status. They would be exempt from all payroll taxes if they choose to work further. Today, most of these workers face a pure tax with no increase in benefits to offset the additional payroll taxes that they face for working. Under this reform, they would neither pay taxes nor improve their benefits from further work. The third reform we examined was to have all of the years that count (forty), count the same. Currently, short careers are favored relative to long careers. For instance, someone who works for seventeen-and-a-half years instead of thirty-five years at the same real indexed wage rate will get significantly more than half the benefits of the full career. Social Security uses the same formula to achieve progressivity that it uses to treat people with different career lengths. Effectively, those with less than full careers are treated as if they are lower income and benefit from the progressivity of the system. There is a relatively simple fix for confounding these two effects. Progressivity can be set in terms of the average salary earned over the years worked. If one works a full career (thirty-five years under current law, forty under the first reform we examine), then the progressive formula would be used to calculate monthly benefits. However, if one works less than a full career, benefits would be reduced proportionately. A twenty-year career would generate half of the benefits of a forty-year career under the proposed reform. All of these proposals could be implemented in such a way as to preserve average benefits at today's levels.

The impact of the three proposals would be a rather dramatic change in the net payroll tax from continued work. The current law has long-career people facing a 10.4 percent payroll tax, whereas short career people actually face a Social Security wage subsidy. What is going on is that Social Security is a net subsidy (the extra benefits are worth more than the extra payroll taxes) as long as you remain on the first segment (the 90 percent segment) of the primary insurance amount-average indexed monthly earnings (PIA-AIME) formula. When you work enough to "graduate" onto the second segment of the PIA formula (the 32 percent section), you face an immediate 10 percentage point jump in marginal tax rates. Those relatively high income people who work long careers and "graduate" to the third segment of the PIA formula (the 15 percent segment) face yet another 3 percent jump in marginal tax rates. Finally, once the additional work stops qualifying in the "high 35" aspect of the PIA formula, the payroll tax becomes simply a tax with no offsetting benefit increases. All of these facts are shown on the left-hand graph in figure 1.11. Our concern is that those with long careers face

high marginal tax rates while those with short careers are subsidized. The three relatively simple reforms that we examine change everything in terms of work incentives, as shown on the right-hand side of the graph. All of the jumps are eliminated as is the predominate fact of increasing marginal tax rates for those with longer careers. To us, it makes sense to try to level the playing field for those with long careers.

Figure 1.12 is another way to look at the three reforms. It shows the benefit
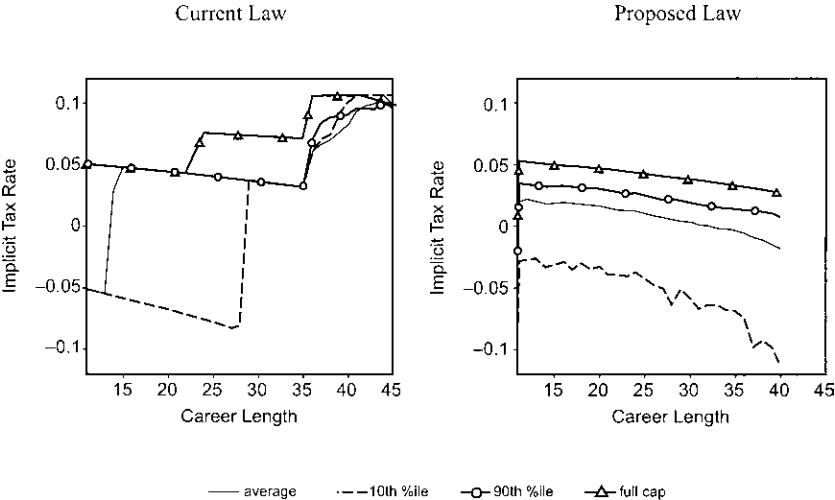

Fig. 1.11   Current law; proposed law


Fig. 1.12   Monthly primary insurance amount under current and proposed law, average income earner

levels for a person with average earnings for different career lengths under current law and under the three alternative benefit rules. What should be clear is that under the revised rules, each of the first forty years has the same marginal improvement of benefits. The current system has lower and declining marginal benefits (the flat region corresponds to the 32 percent bracket in the PIA formula) from lengthening a career. The net impact of the three rule changes would be to offer lower benefits for those with careers shorter than thirty years and higher benefits for those who work longer than thirty years. Given the improvements in life expectancy, these changes in incentives seem appropriate.

1.3   Conclusion

The current practice of measuring age as years-since-birth, both in common practice and in the law, rather than alternative measures reflecting a person's stage in the life cycle, distorts important behavior such as retirement, saving, and the discussion of dependency ratios. Two alternative measures of age have been explored, mortality risk and remaining life expectancy. With these alternative measures, the huge wave of elderly forecast for the first half of this century does not look like a huge wave at all. By conventional sixty-five and over standards, the fraction of the population that is elderly will grow by about 66 percent. However, the fraction of the population that is above a mortality rate that corresponds to sixty-five and over today will grow by only 20 percent. Needless to say, the aging of the society is a lot less dramatic with the alternative mortality-based age measures.

In a separate application of age measurement, I examined the consequences of stabilizing labor force participation by age with alternative age definitions. If labor force participation were to remain as it is today with respect to remaining life expectancy (i.e., if the length of retirement stayed where it is today) rather than labor force participation remaining fixed by conventionally-defined age, then there would be 9.6 percent more total labor supply by 2050 in the United States. This additional labor supply would be very helpful in terms of meeting the challenges of financing entitlement programs, among other things. The GDP might be almost 10 percent higher by 2050 if retirement lengths stabilize. Several policies were examined that would encourage longer work careers.

It is my opinion that the allocation of the extra lifetime in the twenty-first century cannot and will not continue the pattern of the twentieth century—namely, all extra adult lifetime is taken as retirement. Even average retirement ages today look like early retirement when age is measured by remaining life expectancy or mortality risk. In order to allow people to choose when to retire without encouraging an early departure from the workforce, many ages in the laws should be indexed for demographic changes. It is time to consider a new way to measure age.

References

Cutler, D. M., and L. Sheiner. 2001. Demographics and medical care spending: Standard and non-standard effects. In *Demographic change and fiscal policy,* ed. A. Auerbach and R. Lee, 253–91. Cambridge: Cambridge University Press.

Fuchs, V. R. 1984. "Though much is taken": Reflections on aging, health, and medical care. *The Milbank Memorial Fund Quarterly: Health and Society, Special Issue, Financing Medicare: Explorations in Controlling Costs and Raising Revenues* 62 (2): 142–66.

Goda, G. S., J. B. Shoven, and S. N. Slavov. 2007. Removing the disincentives for long careers in Social Security. NBER Working Paper no. 13110. Cambridge, MA: National Bureau of Economic Research, May.

Shoven, J. B. 2004. The impact of major improvement in life expectancy on the financing of Social Security, Medicare, and Medicaid. In *Coping with Methuselah: The impact of molecular biology on medicine and society,* ed. H. J. Aaron and W. B. Schwartz, 166–97. Washington, DC: The Brookings Institution.

Shultz, G. P., and J. B. Shoven, with M. Gunn and G. Shah Goda. 2008. *A citizen's guide to Social Security and health care reform.* New York: W.W. Norton & Company.

Comment   Erzo F. P. Luttmer

The key contribution of this chapter is a straightforward but profound insight. For policies that are age dependent because functioning depends on age, years-since-birth is generally not the best measure of age. For example, we might want to set the retirement age at an age where workers' health on average becomes too poor to reasonably expect work or where a worker has a certain number of remaining years to live. The chronological age at which this occurs is generally not constant over time or across groups. Thus, rather than defining retirement age as years-since-birth, this chapter argues it should depend on a "new age" measure such as mortality risk or remaining life expectancy. More generally, this chapter argues that in many policy settings the relevant measure of age is not chronological age but a measure of age that captures an individual's functioning. In short, it is a thought-provoking chapter that challenges the reader to think at a more conceptual level about what aging really means, rather than narrowly counting the years since birth.

The chapter also discusses a proposal to remove the incentives for early retirement in the Social Security system. This is a sensible proposal, but it is somewhat disconnected from the rest of the chapter because most of the substantive elements of this proposal have little to do with the switch

Erzo F. P. Luttmer is an associate professor of economics at the Kennedy School of Government, Harvard University, and a faculty research fellow at the National Bureau of Economic Research.