

This PDF is a selection from an out-of-print volume from the National Bureau of Economic Research

Volume Title: Pensions, Labor, and Individual Choice

Volume Author/Editor: David A. Wise, ed.

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-90293-5

Volume URL: <http://www.nber.org/books/wise85-1>

Publication Date: 1985

Chapter Title: Overview

Chapter Author: David A. Wise

Chapter URL: <http://www.nber.org/chapters/c7128>

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

1 Overview

David A. Wise

During the past two decades the labor market participation of older workers has fallen dramatically. During this same period private pension coverage has been rapidly extended and social security benefits have been increased. While the observation that these two trends occurred over the same period of time does not mean that one caused the other, it does highlight the possibility that the two may be related. Motivated in part by these trends, the National Bureau of Economic Research has undertaken a study of the labor market aspects of pension plans as part of a larger project on pensions in the American economy. For the past two years, economists at several universities have been engaged in analyzing the nature of private pension plans and their potential incentive effects in the labor market. This volume represents the results of their work.

This overview is intended to introduce the reader to the subject, to motivate the work that follows, and to provide a distillation of the major findings of the volume. Trends in pension coverage and concomitant trends in labor force participation of older workers are discussed first. This serves as a background for our work and helps to motivate the issues that are addressed in the volume. Then the characteristics of the common defined benefit pension plans are described, emphasizing those attributes that are likely to affect labor force behavior. The intent is to demonstrate the order of magnitude of the potential incentive effects of these plans without attempting to present empirical estimates of the impacts, but suggesting that the response of workers to pension plan characteristics could be substantial. To introduce the reader to the subject, the discussion and illustrations in this section also emphasize the possible relationship be-

David A. Wise is John F. Stambaugh Professor of Political Economy, John F. Kennedy School of Government, Harvard University, and research associate, National Bureau of Economic Research.

tween pensions and retirement, while the volume papers cover a much wider range of empirical and theoretical issues. The next two sections review the broad range of the volume's empirical findings and theoretical conclusions regarding the impact of pension plans on the labor force, using the first section as background.

1.1 Trends in Pension Coverage and in Labor Force Participation

The rapid growth in pension coverage over the past three decades is documented in figure 1.1. (The material in this section is taken largely from Ellwood, in this volume.) Approximately 50% of the work force now has some form of pension coverage. Perhaps more significant, the rapid increase in pension coverage has been accompanied by a striking increase in the number of retired persons who collect some pension income, as documented in figure 1.2 (Ellwood, in this vol.). Pension coverage in the government sector is now nearly complete, but a large portion of the private sector still remains uncovered after rapid growth and coverage during the 1950s and somewhat slower growth and coverage thereafter (Ellwood, in this vol.).

With large portions of the private sector still uncovered, who is covered and who is not? What are the correlates of private sector coverage? First, there is a striking relationship between private pension coverage and union status. Nearly 80% of all union members report that they are cov-




Fig. 1.1 Private and government pension coverage as a percentage of the civilian labor force (Ellwood, in this vol.).




Fig. 1.2 Pension benefit recipients as a percentage of the civilian population 55+ (Source: Kotlikoff and Smith, 1983).

ered by a pension plan, while only about one-third of all nonunion workers report such coverage. The relationship between pension coverage and union status is evident in all major industrial categories. Pension coverage is also much more common in large firms than in small ones. And in establishments in all size groups, coverage of union members is more common than coverage of those not in unions (Ellwood, in this vol.). In short, the descriptive evidence suggests that union status and firm size are perhaps the two most prominent determinants of pension coverage.

In addition, the older, the more educated, and the wealthier are more likely to be covered by a pension plan. Nonwhite workers, on the other hand, seem just as likely to be covered as whites. A pattern emerges: workers who gain lesser rewards in the labor market typically are less likely to have pension plans. Thus it appears that pensions do little to counterbalance differences among groups in wage earnings (Ellwood, in this vol.). In summary, pension coverage appears to go hand in hand with union representation, large establishment size, and other economic compensation in the labor market.

While pension coverage was increasing, there was at the same time a striking decline in the labor force participation rate of older workers. For example, between 1960 and 1980, the labor force participation rate for men aged 55–65 fell from 87% to 72%. For those over 65, the rate fell from 35% to below 20% (Ellwood, in this vol.). More detailed evidence for the 1970s is shown in table 1.1. In particular, labor force participation

Table 1.1 Labor Force Participation Rates of Men 60 and Older, 1969-79

Year	60 to 64	65 +
1969	75.8	27.4
1971	74.0	26.3
1973	68.9	23.4
1975	65.4	21.9
1977	62.6	19.9
1979	61.1	20.3

Source: Selected issues of *Employment and Earnings*, table A-4.

fell from close to 76% in 1969 to 61% in 1979. Over the same period the rates for men 65 and older fell from 27% to 20% (Hausman and Wise, in this vol.).

1.2 The Structure of Pension Plans and Their Potential Incentive Effects

Having in mind the increase in pension coverage in recent decades and the substantial decline in labor force participation, it is informative to examine the provisions of pensions that might influence labor market behavior. The goal here is not actually to estimate the impact that pensions have on the labor market but rather to give a rough idea of the magnitudes of the incentives created by pensions and to point to the impact that various pension provisions have on these incentives. (The material in this section is taken largely from Kotlikoff and Wise, in this volume.)

Three-quarters of all persons and nearly 85% of all union members participating in private pension plans are enrolled in defined benefit plans where benefits are determined according to a specified formula. The remainder are enrolled in plans where benefits are directly related to contributions made on behalf of (and by) the employee and to the performance of the plan's investment portfolio. Virtually all government pensions are defined benefit plans. Because most workers are covered by defined benefit plans and because these plans are likely to have the most important effects on labor market incentives, the following remarks are confined to this type of plan.

Defined benefit formulas often are quite complex. Benefits are typically determined by years of service times wages in the last years of employment times a percentage figure (typically 1%–2%). In addition, plans typically have vesting provisions and many are integrated with social security. Early retirement provisions also are an important aspect of most plans. Because of the many plan variations and the complex details of the formulas, it is easiest to understand the potential incentive effects of the

plans by considering pension accrual rates implied by a typical plan (Kotlikoff and Wise, in this vol.).

Consider first a worker who at age 30 begins participating in a defined benefit plan with the following characteristics: the plan calculates normal retirement benefits as 1% of average earnings over the last five years of service times years of service. Benefits are reduced by 3% for each year that early retirement precedes normal retirement. “Cliff vesting” occurs after 10 years; the worker is entitled before this time to no benefits and at 10 years to all benefits that accrue according to the formula described above. The early and normal retirement ages are 55 and 65, respectively. For a typical worker, I shall describe the annual increment to pension wealth as a percentage of the wage rate. Underlying the calculations is a representative lifetime age-earnings profile that assumes substantial growth in real wage rates between ages 30 and 50 and very little growth from 50 to 65. Consider first the pension accrual patterns for this plan if wage inflation is 6% and the nominal interest rate is 9% (3% real). The accrual pattern is shown in the top line of figure 1.3. (Kotlikoff and Wise, in this vol.). The other lines on the graph show the accrual profiles under different interest rate assumptions.

Three aspects of the accrual profiles need to be understood. First, there is no pension accrual before the year of vesting, and in the year of vesting there is an increase in pension wealth that varies from approximately 4% to 14% of wage earnings in that year depending on the interest rate. Sec-




Fig. 1.3

Pension increments as a percentage of salary, by age, for a wage stream with 6% inflation discounted at real interest rates of 3%, 6%, and 9% (Kotlikoff and Wise, in this vol.).

ond, the rate of accrual increases slowly at first and then rather sharply until the age of early retirement. Thus pension wealth tends to be accrued during the later years of one's working life. Third, at the age of early retirement there is a sharp drop in the accrual rate, and the rate falls thereafter. The drop occurs because the early retirement reduction is less than an actuarially consistent reduction would be. That is, annual benefits are not reduced enough to offset the increase in the number of years that one will receive benefits if one retires early. In this case, two competing forces are determining accrual rates after age 55. On the one hand, benefits increase due to increases in wage rates with age and because of additional years of service. On the other hand, the worker gives up the option of taking benefits that are actuarially advantageous. After age 55, the rate of accrual continues to decline. Finally, note that the accrual pattern is substantially affected by the nominal interest rate. It is clear that the accrual due to vesting is affected by interest rates, but it can also be seen that if interest rates are high relative to the rate of inflation, then the accrual after age 55 can indeed be negative. In this case pension wealth could actually decline with additional years of work.

It is useful now to contrast a plan with an early retirement option with a plan that has no early retirement option or that uses an actuarially fair early retirement reduction formula. The difference is shown in figure 1.4 (Kotlikoff and Wise, in this vol.). Notice that without the early retirement option, benefits continue to increase to age 65. Indeed, in this case the




Fig. 1.4

Pension increments as a percentage of salary, by age, for plans with an early retirement option versus retirement at 65, assuming 6% wage inflation and a 3% real interest rate (Kotlikoff and Wise, in this vol.).

pension accrual rate is very much tilted toward the later years of a person's working career.

It should be clear from these graphs that the defined benefit formula seems to provide an increasing incentive to work at least to the age of early retirement. After this the incentive provided by the pension formula declines substantially. To the extent that retirement benefits provide an incentive to continue working, it is also clear that the incentive is much greater without the early retirement option than with it.

It is important to emphasize the impact of the interest rate on pension accruals. Suppose, for example, the graph in figure 1.4 were reproduced, assuming zero wage inflation and a 10% interest rate. Under these assumptions the pension accrual rate after the age of early retirement would indeed be negative. If a person worked an additional year, pension wealth would actually be lost at a rate approaching 15% of wage earnings at age 65 (Kotlikoff and Wise, in this vol.). While these assumptions are probably unrealistic, they demonstrate an important point.

Wage inflation is also important. While the general pattern of pension wealth accrual is not greatly altered by wage inflation, the increment to pension wealth at the time of vesting is very greatly affected by inflation. For example, with 2% wage inflation and a 5% interest rate, the increment to pension wealth in the year of vesting would be over 35% of the wage rate in that year. On the other hand, with 10% wage inflation and a 13% interest rate, the increment at the time of vesting would be only about 5% (Kotlikoff and Wise, in this vol.). Thus, while it is sometimes argued that vesting is not an important determinant of labor market behavior, it can be seen that the effect is likely to vary substantially with wage inflation.

Finally, in the majority of plans, pension accrual ceases once normal retirement age is reached. In many other plans, there are limits on additional accrual. Almost no plans make provisions for actuarial increases to compensate for the fact that late retirees will collect benefits over a shorter period. Typically, continued work after the normal age of retirement involves a substantial loss in pension wealth. Thus pensions typically provide a strong disincentive to work after the age of normal retirement.

What is the effect of pensions on job change? Figure 1.5 illustrates the cost in pension wealth of job change (Kotlikoff and Wise, in this vol.). The graph, for ease of exposition, assumes no early retirement option. The figure shows accrual rate profiles for workers joining the pension plan at ages 30, 40, and 50. The figure is constructed under the assumption that workers of the same age receive identical wage compensation, regardless of job change. Thus the diagram indicates the potential loss in accrued pension benefits for workers who switch jobs but receive the same wage compensation in the new job and are covered by the same pension plan. The top line of this graph shows the accrual rate for a person who




Fig. 1.5 Pension increments as a percentage of salary, by age, for an employee beginning work at 31, 41, 51, with no early retirement option, assuming 6% wage inflation and a 3% interest rate (Kotlikoff and Wise, in this vol.).

starts work at age 30 and does not change jobs thereafter (with 6% wage inflation and a 3% real interest rate). A person with one job change would accumulate benefits up to age 41 according to the top curve but then would accumulate benefits according to the curve labeled age 41. Note that no benefits would be accumulated for the first 10 years. The difference in accumulated pension benefits at age 65 reflects the difference in the areas under the two accrual paths. This difference could be very substantial and depends, of course, both on when job changes occur and how frequently they occur. It is important to note that the loss in this example is not due primarily to vesting. It results from the characteristics of many defined benefit plans that, in effect, index pension accumulation to increases in earnings. This indexation is uncoupled when a job change occurs. If the person were to change jobs at age 51, the loss would be greater. The loss would be greater still if the person were to change jobs twice, say at age 41 and then again at age 51. Individual pension accrual patterns with job change would of course depend on whether the worker becomes vested in earlier plans, but because of their relationship to final wage rates, benefits from the last plan are typically much more important than those from earlier plans.

This graph also shows that, even under moderate wage inflation and interest rate assumptions, vesting can be a very important component of compensation. For a worker hired at age 51, the increment in pension wealth in the year of vesting amounts to 40% of salary in that year.

In short, in inflationary times the benefit provisions of most defined benefit pension plans are in fact powerful deterrents to job switching even for the young and even if the alternative job offers identical wages and provides an identical pension plan. Thus the benefit provisions of defined benefit formulas not only serve as a strong deterrent to retirement prior to the early retirement age, they serve as a remarkably strong deterrent to job switching, even at very young ages. In addition, it appears that pension plans tend to encourage retirement after the age of early retirement and provide a very strong incentive to retire at the age of normal retirement. Whether such effects are intended or whether workers recognize them is yet another question.

We have examined the characteristics of a typical defined benefit plan. But what can one say about the average incentive effects implicit in the wide range of stipulations among actual plans? The accrual profiles implied by a very large number of plans described in the Bureau of Labor Statistics' Level of Benefits Survey look on average much like the hypothetical plans described above (Kotlikoff and Wise, in this vol.). The analysis of actual plans suggests that vesting may provide an important incentive for staying with a firm and that pension wealth accruals provide an important incentive for remaining until the age of early retirement. After this age the incentive for continued work provided by typical pension accruals declines rapidly. And it should be kept in mind that after age 65 the average accrual rate would indeed be negative, certainly encouraging retirement. It is also found, based on the analysis of actual pension plans, that job change can reduce pension benefits substantially. For example, consider accrued benefits at age 65 for persons hired at age 31, 41, or 51. Assume in addition that a person does not become vested in another plan prior to joining the firm. Based on the analysis of 700 plans and estimated wage profiles—and assuming 6% wage inflation and a 9% interest rate—the average aggregate of benefits of persons hired at age 51 are only about 50% of the benefits of those hired at 31. Persons hired at 41 would accumulate about 80% of the benefits of persons hired at 31. Benefits if hired at age 51 range from only 31% of the benefits if hired at 31 in retail trade, to 78% in manufacturing. The ratio among those hired at age 41 ranges from 65% in retail trade to 95% in transportation (Kotlikoff and Wise, in this vol.).

Implications that may be drawn from the above discussion are based directly on the structure of pension plans. In addition to the incentive effects of the plans, their structure also seems to provide substantial evidence for long-term contract versus spot market theories of the labor market (Kotlikoff and Wise, in this vol.). The next section summarizes the volume's more detailed empirical analysis of the determinants of pension coverage, the effects of pension on labor force participation, and important implications of the unfunded nature of the Federal Civil Service Re-

tirement System. For convenience, papers that rely more heavily on a formal conceptual approach are summarized in the following section. The issues addressed include the use of pensions to select certain types of workers and to encourage early retirement, the importance of heterogeneity among workers in designing optimal plans, the incentive effects of particular job characteristics, and pensions as insurance. Simulation results on the riskiness of pension plans are also discussed in the section. The arbitrariness of the grouping, however, is revealed by the very substantial overlap in the issues addressed in the papers and their conclusions.

1.3 The Determinants of Pension Coverage and the Impact of Pensions on Labor Force Participation

The descriptive data above suggest a very substantial relationship between union membership and pension coverage and also suggest that typical defined benefit pension plans could have a potentially significant effect on labor force participation. While these data only suggest possible cause and effect, more detailed analysis points to the existence of a causal relationship.

Unions have a substantial effect on pension coverage, even after controlling for other determinants of coverage such as firm size, worker education, and the wage rate (Freeman, in this vol.). Unions increase the probability that establishments or workers will have a pension plan by sizable and statistically significant amounts, with estimates ranging from .17 to .32. Given the existence of a plan, however, the effect of unions on contributions to the plan is not substantial. Although estimates vary widely, the evidence suggests that for a nonunion worker to have the same probability of having a pension as a union worker, the wage rate of the nonunion worker would have to be at least twice as high as the wage rate of the union worker, given like other characteristics. While establishment size is a key determinant of whether a nonunion worker has a pension plan, it has only a modest effect on whether a union worker has a pension plan.

Not only do unions increase the allocation of funds to pensions, they affect the provisions of pensions as well. In particular, union pension plans are more likely to pay benefits on a flat rate dependent on years of service rather than on earnings. It may be argued that paying flat rate benefits is the pension equivalent of standard rate wage policies and reflects the redistributive goal of unions as a political organization. At the same time, because pension benefits are based on years of service, these plans tend to benefit older workers more than younger ones (Freeman, in this vol.).

The evidence on the structure of pension plans presented above suggests that defined pension plans could have a very substantial effect on retirement behavior. Because available data do not allow one to associate par-

ticular individuals with details of their pension plans, however, it is not possible to provide direct evidence on the effects of the characteristics of these plans on retirement. Reliable data that would reveal the potential interaction between wage rates and pension coverage, for example, have not been released. Pension coverage may not be exogenously determined. In addition, it is important to keep in mind that income after retirement is evidently the major reason for pensions, so if they lead to earlier retirement this should not necessarily be interpreted as an undesirable result. It may, however, be in conflict with other goals of society such as extension of the working life with increases in longevity. Nonetheless, the varied evidence that is available suggests a very substantial effect of pension coverage on the reduction in labor force participation of older workers. Summary data show that the proportion of preretirement income that is replaced by pension benefits is considerably higher for those who retire younger than for those who retire at an older age. For example, among those who retire between 50 and 54, 27% of preretirement is replaced by pensions, while for those who retire after age 65, approximately 17% of preretirement income is replaced by pension benefits (Ellwood, in this vol.). In general, those who retire later tend to have lower pension wealth (Taubman, in this vol.). The implication is that increased benefits may encourage workers to retire earlier. On the other hand, persons who want to retire at an early age could simply arrange to acquire higher pension benefits to accommodate their retirement goals. The circumstantial evidence provided by more detailed analyses, however, suggests that this effect probably is not the major explanation.

Compare, for example, the retirement behavior of federal civil service employees, who have very generous retirement benefits, with the retirement behavior of employees in the private sector. Over half of federal employees retire before the age of 60, whereas the comparable figure for the private sector is only 7% (Leonard, in this vol.). This seems clearly related to the provisions of the federal civil service retirement system. Consider the capitalized value of the pension received under the current federal civil service retirement system by an employee who joins the system at age 25 and who receives typical longevity salary increases over his or her lifetime. Suppose that such an employee were to attain a nominal salary of \$25,000 at the age of 58. The value of pension entitlement for this illustrative employee accumulates slowly across his working life, reaching by age 54 an amount equivalent to about \$130,000, given on his or her sixty-fifth birthday. The next year, when the employee qualifies for full retirement, the value jumps to the equivalent of \$323,000. It stays at this level for a short time and then begins to fall, reaching \$248,000 if the employee waits until age 65 before retiring (Leonard, in this vol.). It is clear that federal employees have a substantial, increasing incentive to work until they reach eligibility for full retirement. At this point, the equivalent of their entitle-

ment peaks; if they continue working it starts to fall. Indeed, this illustrative individual would lose \$75,000 in pension benefits if he retired at 65 instead of 55. The discontinuity in the entitlement at age 55 comes from the shift in the entitlement's value as the employee crosses the combination of age and experience that allows early retirement. This may be a rather extreme case because the civil service pension system is much more generous than the typical private pension plan. Nonetheless, the effect seems clear and the direction of the effect is substantiated by other analyses.

This very generous system is achieved only at very substantial cost and budget consequences. The unfunded liability of the federal civil service retirement system amounts to approximately \$500 billion, a net liability approximately one-half the size of the current officially recognized national debt. Labor expenses recognized in the direct expenditures budget considered by Congress would have to be about 22% higher than they currently are to account for full funding of pension obligations occurred in each year. Thus a full accounting of pension obligations would imply considerably higher labor compensation for federal employees than is commonly recognized (Leonard, in this vol.).

Among California teachers in public schools, and after controlling for other individual attributes that may be related to retirement, the age of retirement shows that larger pension benefits are associated with earlier retirement. Analyzing a random sample of California schoolteachers, both retired and still teaching, we find a standard deviation increase in pension wealth is associated with an approximately one-year decrease in the age of retirement.¹

If higher private pension benefits and provisions lead to earlier retirement, then one might suppose that greater social security benefits would also tend to encourage earlier retirement. The evidence on the effect of social security benefits on the age of retirement is indeed consistent with the evidence on the effect of private pension plans on retirement. Analysis of data collected in the Retirement History Survey, that follows for 10 years persons aged 58–63 in 1969, verifies two hypotheses about the likely effects of pension benefits (Hausman and Wise, in this vol.). The first is that, other things equal, persons who can retire with larger benefits are more likely to leave the labor force than those who would receive smaller benefits if they were to retire. The second is that persons who by working another year could increase the discounted value of pension benefits that they would receive in the future are more likely to continue working. Or, if working another year would reduce the present value of future benefits, the individual is more likely to retire early. This loss is analogous to the typical loss in private pension benefits if one continues to work after the age of early retirement, as demonstrated in the graphs above. In short, analysis of the Retirement History Survey shows both a strong effect of the increment in pension benefits on prolonging labor force participation

and a strong effect of cumulated social security wealth on departure from the labor force. Indeed, the data suggest that a large proportion of the dramatic decrease in labor force participation of older workers between 1969 and 1979 can be explained by the large increases in social security benefits over this period (Hausman and Wise, in this vol.).

It can be argued, as mentioned above, that individuals choose jobs and associated pension plans, make savings decisions, and take account of expected social security benefits to provide retirement income consistent with preplanned desired age of retirement. In major part, however, this could not be said to be true of social security benefits, large portions of which must have been unexpected, in particular the very large increases in the early 1970s. Thus the relationship between social security benefits and age of retirement that is observed in the data must have been induced to a substantial degree by unexpected social security benefits and is not a reflection of retirement decisions based on expected benefits.²

Indeed, further analysis of the Retirement History Survey shows that retirees in this sample would receive three to four times as much in benefits as they made in contributions (Hurd and Shoven, in this vol.). Possibly more surprising, the wealthy received the largest transfers in absolute amounts, and in many cases they even received the highest rates of return on social security contributions (Hurd and Shoven, in this vol.).

The description of pension plan provisions presented above also suggested that, at least through the age of early retirement, pensions should tend to reduce job change. Although to date we have not undertaken detailed analyses of the effect of pension plan provisions on job change—again due to limitations of data available to us—summary data suggest that pension provisions do reduce job turnover. For example, whereas 73% of all workers aged 51–55 without pensions have less than 10 years' tenure, only about 33% of workers with pensions have been with their employer for less than a decade (Ellwood, in this vol.). The direction of this relationship holds for every age group. Of course, such data prove nothing. There is a host of possible explanations for these differences, some of which are emphasized in conceptual work referred to below. Still, the data do suggest that pensions may in fact have sizable impacts on job change.

1.4 More Formal Conceptual Results

Conceptual analyses of the effects of pension plans in general reinforce the implications suggested by the structure of the plans described above. In addition, these analyses provide insights that are not so readily discernible based on the structure of the plans.

Analysis of accrued benefits based on plan provisions presented above shows that job change is likely to be associated with a substantial cost in loss of pension benefits. Descriptive data mentioned above also show that

turnover or job change among employees covered by pensions is much less than among those who are not covered. Consistent with this observation, theoretical analysis demonstrates that pensions can be important in reducing turnover among workers who are attracted to the firm (Viscusi, in this vol.). But, in the face of worker uncertainty about future job performance and turnover costs, pensions can also be important in self-selecting more stable employees. (Viscusi, in this vol.). This suggests, of course, that observed relationships between pension coverage and job turnover are likely to result at least in part from self-selection, by workers who would like to avoid job change, of firms with appropriate pension plans.

The structure of pension plans also suggests that most plans tend to encourage retirement after some age, in many instances after the age of early retirement and in almost all cases after the age of normal retirement. It is difficult to make contracts that commit workers to retire at a certain age, and it may be institutionally impractical to reduce wages commensurate with the reduction in productivity of older workers. It can be demonstrated, however, that pensions may be used to help induce appropriate retirement behavior, consistent with economic efficiency (Nalebuff and Zeckhauser, in this vol.). In addition, work heterogeneity has important implications for the design and effects of pension plans. In particular, heterogeneity among workers implies that the concept of actuarial fairness is illusive and much more ambiguous than is commonly recognized (Nalebuff and Zeckhauser, in this vol.). As demonstrated above, most defined benefit plans are apparently actuarially unfair to those who prolong their working lives. However, if workers can estimate their life span, those who expect to live longer may choose to retire later and the apparent pattern of actuarial benefits may be reversed. For example, a person who expects to live five more years may find that the increase in pension benefits that would be gained by working another year is not large enough to offset the reduction of one year in the number of years that benefits would be received. On the other hand, a person who expects to live 10 more years and thus would receive the incremental benefits over a longer period of time may find that it is worthwhile to prolong labor force participation for an additional year. The structure of a pension plan that is optimal for an individual, that is, one that maximizes *ex ante* expected utility, however, must make it actuarially unfavorable for the individual to retire later than some age (Nalebuff and Zeckhauser, in this vol.).

We know, of course, that typical pension plans cover large numbers of workers that may be dissimilar in many respects, in particular with respect to life expectancy. Some proposed pension reform proposals would indeed impose more similar plans on even more heterogeneous groups of workers. In a first-best situation, however, a separate pension plan would

be designed for each homogeneous group of workers. The imposition of the same plan on nonhomogeneous groups of workers may be shown to work to the possible detriment of each of the groups. That is, it could turn out that under the common plan no group is better off than under the plan tailored to its particular characteristics (Nalebuff and Zeckhauser, in this vol.).

Given existing compensation arrangements, heterogeneity among individuals may in general have an important impact on the compensation that individual workers receive. A variety of economic models of compensation suggest that firms should allow workers to choose individually how they would like to have their compensation divided between wages and benefits. More compensation in the form of retirement benefits or health insurance benefits would be offset by reduced wage rates. In this case, there would be no reason for firms to try to attract certain types of workers. In practice, however, even among equally productive workers wages and benefits will not necessarily balance (Bulow and Landsman, in this vol.). Workers may, for example, receive the same wage rather than the same total compensation. Nonetheless, for labor as a whole, increases in aggregate benefits would be offset by decreases in wage compensation. Under these circumstances, firms that offer particular types of benefits should expect to attract workers likely to take advantage of these benefits. An important implication of this possibility is that we can say only a limited amount about the aggregate firm worker implicit labor contract by looking at individual compensation profiles (Bulow and Landsman, in this vol.).

It should be clear that defined benefit pension plans have potentially large incentive effects on labor participation. Whether they are efficient or not is another question. If the pension rule is taken to be exogenous, then many provisions of defined benefit plans are shown to have adverse incentive effects. For example, a large number of plans incorporate a 10-year cliff vesting provision, but complete and immediate vesting would be a necessary condition for fully efficient pension plans (Lazear, in this vol.). On the other hand, defined contribution plans, in contrast to defined benefit plans, always induce an efficient allocation of resources (Lazear, in this vol.).

An alternative to the analysis of the incentive effects of exogenously given pension plans is to consider the form that private pensions are likely to take when employers provide pensions—in contrast to an optimally formulated government pension policy, for example. In particular, workers may like to have insurance against many work-related contingencies: disability, future productivity, change in employment. Because these uncertainties are hard to verify, however, they cannot be the object of insurance directly. A worker may thus want to allocate some part of his current

compensation to provide income in situations in which the income will be more important to him, and thus pensions may serve as insurance (Diamond and Mirrlees, in this vol.).

Finally, uncertainty about receipt of future pension benefits is an important aspect of the equivalent compensation value of the plans. This, of course, is not an attribute of plans that can be gleaned from the structure of the plans as described above. Several economically relevant events during a worker's lifetime, both macroeconomic events and individual-specific ones, are likely to affect the value of pension benefits actually received. Through simulation analysis, it is possible to compare the riskiness of alternative forms of defined benefit pension plans (Green, in this vol.). Results based on this methodology show that among four forms of defined benefit plan, a plan based on a percentage of final salary is least risky when compared to a plan that bases benefits on the highest 10 years of salary, a plan based on the worker's career average salary, and one that determines benefits by years of service only. It may be argued that the extent to which pension plans are risky has important implications for the effect of pension plans on other forms of individual saving (Green, in this vol.). The certainty equivalent value of a pension to the worker is likely to be much less than the value of the liability it represents to the firm. This is because part of the variance in pension wealth viewed by the worker is due to individual-specific risk that from the point of view of the firm is averaged out over many workers. Therefore it may be argued that pensions decrease the value of other forms of savings by workers by less than the increase in required saving by firms necessary to offset their pension liabilities. Thus the riskiness of private pensions may have substantial importance in studying the effect of private pensions on aggregate savings in the national economy (Green, in this vol.).

1.5 Conclusions

Over the past three decades pension coverage has expanded substantially; over the latter part of this period there has been a very dramatic reduction in the labor force participation of older workers. The typical defined benefit pension plan provides an increasing incentive for an employee to remain with the same firm until the age of early retirement; after that the structure of the typical plan provides less incentive to work, and indeed the stipulations of many plans imply actual losses in pension wealth accrual after the age of early retirement. After the age of normal retirement, almost all plans incorporate loss in pension wealth if the employee continues to work. Thus the structure of pension plans alone suggests that the increase in pension coverage may have had a good deal to do with the decline in the labor force participation rate of older workers. More detailed analysis of the relationship between pension coverage and age of retire-

ment appears to confirm this possibility. In addition, empirical analysis suggests that the increases in social security benefits have played a substantial role in the decline in the labor force participation of older workers.

In addition to providing an incentive for older workers to retire, the typical structure of defined benefit pension plans imposes a substantial cost in pension wealth on workers who change jobs. The order of magnitude of these losses may not be widely appreciated, but the evidence available suggests that pension coverage reduces job change.

More formal analysis of pensions reinforces the apparent implications that can be drawn from the structure of the plans themselves. In particular, the use of compensation in the form of pension coverage can provide a mechanism for encouraging older workers to retire when institutional constraints may prevent employers from reducing the wage earnings of older workers in accordance with reductions in labor productivity. Theoretical analysis also demonstrates that pensions can be used to select more stable employees and to discourage job changing among those hired. In addition, conceptual analysis highlights the importance of heterogeneity among employees, suggesting that policies that encourage uniformity among pension plans may not necessarily increase the general welfare of employees.

Notes

1. These findings are from Steven I. Kutner, "Individual Attributes and Pension Acceptance Decisions: A Case Study." Although presented at the conference, the final version of the paper was not finished in time for inclusion in this volume.
2. This point is emphasized by Hurd and Boskin (1981).

References

- Hurd, M., and Boskin, M. 1981. The effect of social security on retirement in the early 1970s. NBER Working Paper no. 659.
- Kotlikoff, L. J., and Smith, D. E. 1983. *Pensions in the American economy*. Chicago: University of Chicago Press (for NBER).

This Page Intentionally Left Blank