

This PDF is a selection from a published volume from the
National Bureau of Economic Research

Volume Title: Social Security Programs and Retirement
around the World: Micro-Estimation

Volume Author/Editor: Jonathan Gruber and David A. Wise,
editors

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-31018-3

Volume URL: <http://www.nber.org/books/grub04-1>

Publication Date: January 2004

Title: Introduction and Summary

Author: Jonathan Gruber, David A. Wise

URL: <http://www.nber.org/chapters/c10698>

Introduction and Summary

Jonathan Gruber and David A. Wise

If they aren't paid, people don't work. This fundamental economic principle is dramatically demonstrated by social security provisions and retirement. The social security programs in most developed countries are financed on a pay-as-you-go basis. Under this arrangement, most countries have accumulated large unfunded liabilities and, in many countries, face looming financial burdens. The aging of the populations in almost all countries is often cited as the reason for the financial burdens faced by the social security programs. Many of the programs are very generous and thus are increasingly costly as the population ages because there is a greater proportion of the population retired and collecting benefits, relative to the fraction of the population that is in the labor force and paying for the benefits. Perhaps just as important, although not as widely appreciated, is that the provisions of the programs themselves typically encourage retirement by reducing pay for work. This penalty on work induces older employees to leave the labor force early and thus magnifies the financial burden caused by population aging.

This volume represents the second stage of a research project to study the relationship between social security provisions and retirement. The first stage of the project is reported in Gruber and Wise (1999). In that volume, we documented the enormous disincentives for continued work at older ages in many countries. The analysis also revealed a strong corre-

Jonathan Gruber is professor of economics at the Massachusetts Institute of Technology, director of the research program on children at the National Bureau of Economic Research, as well as research associate of the National Bureau of Economic Research. 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.

spondence across countries between social security program incentives to retire early and the proportion of older persons that has left the labor force. The weight of the evidence suggested that this relationship was largely causal. The results of the second phase strongly affirm the causal relationship between retirement and social security incentives to quit work.

In this second stage, we turn to country-by-country analysis of retirement behavior based on micro-data. The research teams in each of twelve countries compiled comprehensive, large databases of individuals. The data in each country match information on retirement decisions to the retirement incentives inherent in the social security provisions of each country. Retirement models were estimated based on the micro-data. The results show the enormous effect of social security incentives on retirement, and the uniformity of findings is striking. In every country, the quantitative magnitude of the incentive-program effects on retirement is very large.

The key advantage of this micro-estimation approach is that in each country the effects on retirement of changes in social security provisions can be predicted. To demonstrate the effect of such changes, the country papers in this volume include simulations of the effects of two *illustrative* reforms: One illustrative reform delays the benefit eligibility ages in each country. A second illustrative reform assumes common provisions in each of the countries—reducing retirement incentives in some countries and increasing incentives in other countries. Under the first reform, the simulations show a large reduction in retirement in each country, and a corresponding increase in the labor force participation of older workers in each country. Under the second reform, the simulations show an increase in retirement in some countries and a decrease in other countries, in accordance with the relationship between the current country provisions and the common reform provisions.

In short, the findings confirm the conclusions based on the first stage of this ongoing project and, in addition, illustrate the enormous magnitude of the effect that changes in social security provisions would have on retirement and thus the labor force participation of older people.

Like the first stage, this second stage of the project relies on the analyses of a large group of economists who conduct the analysis for each of their countries. The authors of the individual country papers in this volume are

Belgium: Arnaud Dellis, Raphaël Desmet, Alain Jousten, and Sergio Perelman;

Canada: Michael Baker, Jonathan Gruber, and Kevin Milligan;

Denmark: Paul Bingley, Nabanita Datta Gupta, and Peder J. Pedersen;

France: Didier Blanchet and Ronan Mahieu;

Germany: Axel Börsch-Supan, Simone Kohnz, Giovanni Mastrobuoni, and Reinhold Schnabel;

Italy: Agar Brugiavini and Franco Peracchi;

Japan: Takashi Oshio and Akiko Oishi;

The Netherlands: Arie Kapteyn and Klaas de Vos;
 Spain: Michele Boldrin, Sergi Jiménez-Martín, and Franco Peracchi;
 Sweden: Mårten Palme and Ingemar Svensson;
 The United Kingdom: Richard Blundell, Costas Meghir, and Sara Smith;
 and
 The United States: Courtney Coile and Jonathan Gruber.

The central feature of the project is the presentation of comparable analysis in each of the countries. Each of the country studies follows essentially the same format, although country-specific issues are often discussed as part of the analysis done for that particular country.

Background: The First Stage

The goal of the first stage of the project was to describe the incentives inherent in the social security provisions in the project countries and to relate the incentives to the labor force participation of older workers. Each of the studies in the first volume begins with a description of the historical evolution of labor force participation and then presents data on the current age-specific activities and income sources of men and women. Each of the papers presents data for men and women in nine areas:

1. Labor force participation rates by age interval between 1960 and the present;
2. The proportion of employees covered by the public pension system and the proportion of persons over fifty-five receiving public pensions from 1960 to the present;
3. Replacement rates under the public pension system from 1960 to the present;
4. Current labor force participation rates by age;
5. Labor force status—employed, unemployed, disability, or retired;
6. Proportion receiving various public pensions—such as old age, disability, or survivor—by age;
7. Proportions receiving employer-provided pensions by age;
8. Source of household income by age;
9. Retirement and public pension hazard rates by age.

Each paper then describes the institutional features of the country's social security system, highlighting any interactions with other public and private programs that might also influence retirement behavior. The core of each paper is a detailed analysis of the retirement incentives inherent in the provisions of that country's retirement income system. By making the same analytic calculations and by presenting the same simulations in each of the countries, the individual studies provide a means of comparing the retirement incentives among the countries.

Because it provides the background and the motivation for the continu-

ing project, we summarize here the key results from the first stage of the project. The decline in the labor force participation of older persons is perhaps the most dramatic feature of labor force change over the past several decades. The decline has been striking in all but one of the countries studied here. The labor force participation rates of men aged sixty to sixty-four for the years 1960 to 1996 are shown for each of the eleven countries in figure 1, which for ease of exposition is presented in two panels. (Denmark was added to the project after the first stage was completed.) The decline was substantial in each of the countries, but was much greater in some countries than in others. In the early 1960s, the participation rates were




Fig. 1 LFP trends for men ages 60 to 64: *A*, Japan, Sweden, US, UK, Germany, Belgium; *B*, Japan, Spain, Netherlands, France, Italy

above 70 percent in each of the countries and above 80 percent in several countries. By the mid-1990s, the rate had fallen to below 20 percent in Belgium, Italy, France, and the Netherlands. It had fallen to about 35 percent in Germany and 40 percent in Spain. Although U.S. analysts have often emphasized the “dramatic” fall in that country, the U.S. decline from 82 percent to 53 percent was modest in comparison to the much more precipitous decline in these European countries. The decline to 57 percent in Sweden was also large, but modest when compared to the fall in other countries. Japan stands out with the smallest decline of all the countries, from about 83 percent to 75 percent. Each of the country papers presents completely parallel labor force and other data for men and women, including current labor force participation and labor force departure rates by age, which are key components of the analysis in this volume.

By considering the labor force participation rates by age in each country it is possible to calculate the proportion of persons in a given age interval who were out of the labor force. These unused-productive-capacity measures for all of the countries are shown in figure 2, for men in the age range from fifty-five to sixty-five. For the entire age range from fifty-five to sixty-five, unused capacity ranges from 67 percent in Belgium to 22 percent in Japan. The goal was then to consider how this measure of labor force participation was related to the provisions of the social security programs in the countries.

The incentive measure calculated in this first volume was the implicit social security tax on work. To understand that measure, it is useful to think of wage compensation for working an additional year in two components.




Fig. 2 Unused productive capacity




Fig. 3 Sum of tax rates on work from early retirement age to 69

The first is wage earnings. The second component is the increase in the expected present discounted value of promised future social security benefits, known as the accrual in social security wealth (SSW), which is equal to $SSW_{t+1} - SSW_t$. It is natural to think of this difference as positive, or at least not negative—that is, if a person works for an additional year and thus forgoes one year of benefits, it might be expected that benefits begun one year later would be increased enough to offset the fact that they are received for one fewer years. This is true, for example, for the typical worker in the United States: If a worker forgoes claiming benefits at the earliest possible age (sixty-two) and works another year, subsequent benefits are increased by 6.67 percent to account for the fact that benefits will be received for one fewer year.

In most of the countries studied in this project, however, the accrual is significantly negative. This is a consequence in large part of not increasing benefits enough if the age of benefit receipt is delayed so that benefits are not “actuarially fair.” Thus, what the worker gains in wage earnings is partially, or even largely, offset by a loss in future social security benefits. We call the ratio of this loss to wage earnings (after tax) the social security implicit tax on earnings. In many countries this tax can be 80 percent or more at certain ages. Suppose that the tax rates for each of the years from the early retirement to age sixty-nine are summed: We call this the tax force to retire. It is shown for all the countries in figure 3.

The relationship between the (logarithm of the) tax force to retire and the proportion of men age fifty-five to sixty-five that is out of the labor force (unused capacity) is shown in figure 4. There is a striking correspon-

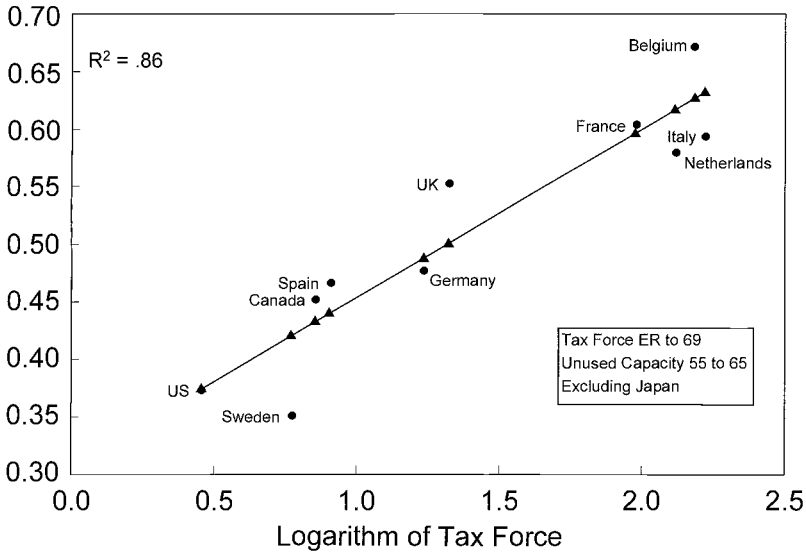


Fig. 4 Unused capacity versus tax force

dence between the two series, showing a clear relationship between the social security tax on work and departure from the labor force. A number of examples in the previous volume, some of which are replicated below, suggest that the relationship is largely causal. We concluded in the introduction to that volume that:

It is clear that there is a strong correspondence between the age at which benefits are available and departure from the labor force. Social security programs often provide generous retirement benefits at young ages. In addition, the provisions of these programs often imply large financial penalties on labor earnings beyond the social security early retirement age. We conclude that social security program provisions have indeed contributed to the decline in the labor force participation of older persons, substantially reducing the potential productive capacity of the labor force. It seems evident that if the trend to early retirement is to be reversed, as will almost surely be dictated by demographic trends, changing the provisions of social security programs that induce early retirement will play a key role. (Gruber and Wise 1999, 35)

The Second Stage Analysis and Issues

The first stage of this project established two key findings: (a) The social security systems in many countries provide enormous incentives to leave the labor force at older ages, and (b) A strong correspondence exists between social security incentives to retire and the withdrawal of older work-

ers from the labor force. This implies that social security incentives to retire are likely an important cause of the low labor force participation of older workers in many countries. The relationships in the first volume, however, do not provide a means of estimating the magnitude of the effect on labor force participation of changes in plan provisions.

Thus, in this second stage of the project, we undertake analysis to estimate how much the retirement age would change if social security provisions were changed, based on within-country analysis of the determinants of retirement. The analysis is based on the micro-data for each country that considers the relationship between retirement and the incentives faced by individual employees. That is, rather than considering system-wide incentives for representative persons (such as those with median earning histories) and comparing these incentives with aggregate labor force participation across countries, we now turn to micro-econometric analyses within countries. The results of these analyses are based on differences in individual circumstances within a given country.

This approach has two key advantages. First, the analyses in this volume show that social security retirement incentives have very similar effects on labor force participation in all countries. In particular, the results strongly confirm that the relationship between labor force participation and retirement across countries is not the result of cultural differences among countries that could yield different norms for work at older ages. That is, the within-country analyses show similar responses to retirement incentive effects, even though the countries differ with respect to cultural histories and institutions.

Second, the analysis of micro-data also allows consideration of several features of social security systems, as well as individual attributes, that may simultaneously affect retirement decisions. In particular, we can consider jointly the age at which benefits are first available and the incentive to retire once benefits are available. The first stage of the project showed that both of these features were important determinants of retirement.

The importance of benefit eligibility ages presents a particular challenge for the analyses in this volume. We believe that much of the effect of social security provisions is likely to be through the choice of benefit eligibility ages, which in some instances may tend to establish social norms for retirement. For example, a common finding of many of the analyses in this volume is that even very detailed models of retirement incentives cannot explain the large jumps in retirement rates at normal and early entitlement ages. As a result, the retirement effects of major system reforms (like raising the early entitlement age) may be greatly understated by comparisons among individuals within a given retirement income system who all face the same eligibility ages. In addition, small private saving may limit the option of most persons to retire before pension benefits are available (many people are liquidity constrained), and such effects may not be cap-

tured by retirement models. We discuss later how we deal with this critical issue.

In addition, as emphasized in the first volume, unemployment and disability benefits often provide early retirement income before the nominal social security normal retirement age. Thus in many countries, to estimate the effect of the plan provisions on retirement, it is necessary to consider all three programs jointly, which the micro-econometric analyses in this volume accomplish.

Considered jointly, we believe that the analyses in this volume provide overwhelming confirmation that the provisions of social security programs play a key role in the determination of retirement decisions. This result complements the conclusion from the first stage of the analysis. In addition, the estimates in this volume allow prediction of the effects on retirement of changes in program provisions as well as the effect of changes on program costs, which will be taken up in the next stage of the project.

As in the first volume, the analysis in each country follows a template so that results can be compared across countries. The micro-analysis in each country is based on a sample of individuals. In some cases, the data come largely from administrative records. In other cases, the data were obtained from special surveys. The coverage is not precisely the same in each country. For example, the data for Italy pertain to private sector workers only, excluding public sector employees. Nonetheless, it has been possible to estimate the same models in each country, even though the population covered by the country data sets may differ in some respects.

In this section, we first describe the incentive measures used in the analysis, as well as other features of the model specifications. The alternative incentive measures are constructed economic variables that describe the financial gain or loss from continuing to work. Then we discuss the method used to obtain estimates when there are multiple routes to retirement. Next we briefly summarize the parameter estimates obtained in the analyses across the twelve countries.

In the next section, we discuss a key empirical regularity that strongly influences the analyses undertaken in this volume: the correspondence between benefit eligibility ages and retirement decisions.

In the concluding section we describe the simulations undertaken to summarize the implications of the estimated models, and discuss the simulation results. The simulations describe the effect of *illustrative* policy changes. The goal is to provide an understanding of the nature of the findings, focusing on selected portions of the analyses described in detail in the country papers. The simulations demonstrate the implications of the retirement model estimates.

As part of this discussion, we take some care to explain the different methods of simulation that are used in the analysis and why certain features of the simulation methods are of key importance. A central aspect of

the analysis is experimentation with various approaches to estimation and to simulations based on the estimates. The aim is to determine the most reliable methods to use in the subsequent stages of the project.

Some of the discussion is necessarily somewhat technical. We believe, however, that it is important to present an overview of the critical features of the analysis so that readers can approach the individual country papers with a broad understanding of the issues and rationale behind the approach taken in the country analyses. Thus, in the text of this introduction, we have explained the main features of the analysis and have included some additional, more technical detail in an appendix.

The Estimation Models: Incentive Measures and Control Variables

The goal of the analysis in the country papers is to estimate the probability of retirement based on the provisions of the country's social security system—which provides differential incentives depending on individual employee circumstances—and on other individual attributes of employees. The focus of the analysis is the plan provisions. In particular, the way the incentives to retire, inherent in plan provisions, are in fact related to the retirement choices that individuals make. There are several ways that the incentives might be measured. All of the measures describe the financial gain or loss from continuing work. The specifications used in the country papers are summarized in the following table. The most important variable in each specification is the incentive measure, which is noted across the top of the table. In addition to these measures, the models control for various other variables (covariates).

The specifications are summarized in table 1. Each of the specifications includes SSW. The expectation is that, all else equal, persons with greater SSW are more likely to retire. In principle, total wealth should be controlled for, but in most countries the data do not provide measures of other forms of wealth.

The focus of the analysis is on forward-looking measures of the incentive for retirement, or for continued work. A natural starting point is a

Table 1 Estimation Method (Incentive Measure) and Variables

Variables	Estimation Method (Incentive Measure)					
	Single Year		Peak Value	Option Value	Option Value	Option Value
	Accrual	Accrual				
SSW	X	X	X	X	X	X
Linear age	X		X		X	
Individual age indicators		X		X		X
Earnings	X	X	X	X	X	X
Sector	X	X	X	X	X	X
Demographics	X	X	X	X	X	X

measure that looks ahead only one year, the single-year accrual measure. This measure captures the effect of another year of work on future benefits. Thus, as a basis for comparison, the country analyses present the single-year accrual incentive measure.

However, it has been shown in other contexts,¹ as well as in the first volume of this project, that the financial gain from continuing to work may vary from year to year. That is, the gain from working one more year may be large, for example, but once that single additional year is worked, the gain from working one more year may be small or even negative. Likewise, the gain from one more year may be small, but might then be followed by a year of large gain. In this case, a person who decides to retire based on looking forward just one year, would forego the gain in pension wealth that would be gained by continuing to work for two years. Thus a key principle of the approach followed here is that the estimation should account for the pension accrual not just in the next year, but many years into the future.

The benchmark approach for considering the entire future path of accruals is the option value model.² To summarize, this model evaluates the expected present discounted value of incomes for all possible future retirement ages and then measures the value of retirement today versus the value at the optimal date (perhaps today, but more likely in the future). If looking ahead suggests gains from work at some time in the future, there is an incentive for the person to remain in the labor force to take advantage of these gains.

The relationship between the measures can be explained briefly. As previously described, the social security accrual from one year to the next is given by:

$$(1) \quad SSW_{t+1} - SSW_t$$

That is, this measure describes the change in promised future social security benefits from working one additional year.

A simplified version of the option value measure at age t can be described by

$$(2) \quad \begin{aligned} & \text{Simplified } OV_t(r^*) \\ &= \left(\begin{array}{l} \text{discounted} \\ \text{future wages} \\ \text{through age } r^* \end{array} \right) + \left[\left(\begin{array}{l} \text{discounted} \\ \text{benefits if} \\ \text{retire at } r^* \end{array} \right) - \left(\begin{array}{l} \text{discounted} \\ \text{benefits if} \\ \text{retire at } t \end{array} \right) \right] \\ &= \left(\begin{array}{l} \text{discounted} \\ \text{future wages} \\ \text{through age } r^* \end{array} \right) + [\text{Peak Value}]. \end{aligned}$$

1. For example, see Lazear (1983), Kotlikoff and Wise (1985, 1989), and Coile and Gruber (2000a).

2. See Stock and Wise (1990a,b).

In this formulation, a person considering whether to retire at age t considers the present value of benefits if he retires now (at age t) with the benefits if he retires at some later time. If the person retires at some later age he will gain from future wage earnings and from any gain in future pension benefits. The gain in wage earnings is represented by the first bracket and the gain in pension benefits by the difference between the terms in the second bracket. The age at which the total of the two components is the greatest is denoted by r^* . The option value prescription is that the person will continue to work if this option value is positive.

Notice that the option value approach as set out above combines both of the components of compensation from working: one component is wage earnings, the other is the change in promised future social security benefits. We label this second component the peak value.³ It includes only social security benefits and not wage earnings. The peak value occurs at the age that gives the greatest discounted value of social security benefit. That age need not be at r^* , although for simplicity the two are assumed to be the same in this description. A more precise discussion of the differences in the measures is presented in the appendix.

As previously emphasized, a crucial issue in the analyses in this volume is identification—that is, determination of the separate effect of each variable on retirement, as distinct from each of the other variables. Determining the effect of plan incentives on retirement is a key goal, but other individual attributes also influence the decision to retire. For example, persons are more likely to prefer retirement to work as they age. A linear age variable will potentially capture this effect, but only if preferences for leisure evolve linearly with age. Individual wage differences may also *proxy* for differences in the preference for work versus retirement. A wage-earning covariate may help to control for this form of heterogeneity among individuals. But both age and wage earnings also determine in part the value of the incentive measures. Thus, including age and wage covariates may make it more difficult to accurately determine the effect of the incentive measures or to accurately isolate the program incentive effects from the effect of worker heterogeneity. To put it another way, the importance of controlling for differences in taste regarding work may suggest the inclusion of the wage and age variables, separate from their incorporation in the option value. But there is a countervailing consideration: Much of the estimated effect of these variables is likely to reflect the influence of financial incentives and not individual heterogeneity. Thus, the full effect of the program incentives may be understated when the separate age and earnings controls are included.

The issue of identification also arises in considering the option value as compared with the peak value incentive measures. If individual heterogeneity were not a concern, the option value measure would be the most

3. As proposed by Coile and Gruber (2000a,b).

parsimonious incentive measure to use, as it captures the full financial incentive on retirement of both future wage earnings and retirement benefits combined. But, to the extent that wages proxy for the taste for work, the option value variation across individuals may reflect in part this wage proxy for heterogeneity, rather than the financial retirement incentive. The peak value measure recognizes this possibility by measuring the retirement incentive by the future stream of retirement benefits only, without including the future stream of wage earnings. But to the extent that future wage earnings have an important incentive influence on retirement, the peak value approach understates the full effect of financial incentives on retirement.

Perhaps the most important identification issue arises when age-specific variables are included to estimate the retirement effect of program eligibility ages. It is common to find that the retirement rate at certain ages is larger (or smaller) than would be predicted on the basis of an incentive measure alone. For example, in the United States, the retirement rate at age sixty-five is noticeably higher than is predicted based on financial incentives alone. Perhaps this is due to a customary retirement age effect: Since age sixty-five is the normal retirement age, many employees may think that age sixty-five is the age to retire. In addition, in virtually no instance in any country do employees typically retire before some form of retirement benefits are available. The retirement rate at the early retirement age—the age at which a person is first eligible for benefits—is typically substantially greater than would be predicted on the basis of financial measures alone. This empirical regularity likely reflects a liquidity constraint; most employees have not saved enough to retire without receiving public social security or employer-provided pension benefits. To capture this effect, some of the specifications allow an indicator variable for each age. These indicator variables allow retirement to jump or decline at each age, after controlling for the financial incentive measures. But the inclusion of these age indicator variables in particular raises the identification question: Here, the issue is whether the effect of the financial incentive measures can be distinguished from the effect on retirement of plan eligibility ages. This is a critical issue and is discussed in more detail later.

There is no right answer to these identification dilemmas. The country analyses in this volume follow what is perhaps a conservative as well as a flexible approach. All of the specifications control for background variables, including sex, education, industry of employment, and both current and average lifetime earnings. This is the conservative part. The flexibility is reflected in the different incentive measure specifications, each estimated using linear age and then again using indicator variables for each age. In this way the sensitivity of the findings to different incentive measures and to the controls that are included in the analysis can be assessed.

One additional note on estimation: The option value model as set out by Stock and Wise, and in several subsequent applications, was estimated by

maximum likelihood methods to obtain the relevant behavioral parameters. In this cross-country context, however, we concluded that that approach would likely pose numerical complexities that would best be avoided for this comparative analysis over a large number of countries. Thus the estimation undertaken in the papers in this volume is typically based on a regression counterpart to the option value model as well as two other approaches, as previously explained. In some countries, however, the option value parameters in the Stock and Wise specification have been estimated by a grid search, and in other instances the Stock and Wise option value parameter estimates have been used to calculate the option value that is used in the regression. In at least one country, the option value model was estimated by maximum likelihood.

Multiple Retirement Options

In some countries, like the United States, social security is the single public program that provides retirement benefits to the vast majority of retirees. The only retirement decision is then at what age to choose to retire under this program. (A small fraction of persons retire under the public disability program and many employees are covered by employer-provided pension plans.) In other countries, however, there are two or more programs under which a person can retire. Germany is a good example.

Figure 5 shows paths to retirement for men in Germany between 1960 and 1995. The figure shows clearly the changes in the pathways to retirement after the 1972 reform, which is discussed later. Here, we draw attention to the multiple paths to retirement. All persons are eligible to retire at age sixty-five, the social security program normal retirement age, but only a small proportion of employees work until that age. A large fraction of

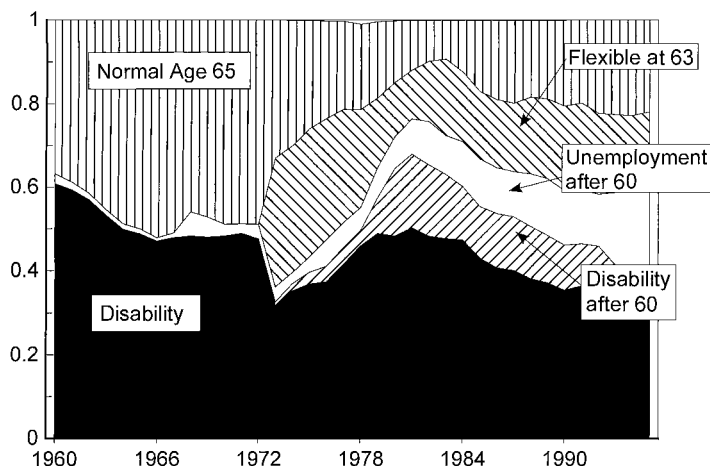


Fig. 5 Germany: Pathways to retirement for men, 1960 to 1995

Source: Data provided by Axel Börsch-Supan.

employees retire under the disability program before age sixty. Others retire under the social security disability program available after age sixty. Some can retire under the social security system unemployment program after age sixty. Still others are eligible to retire at age sixty-three under the flexible retirement program that allows persons with long service to retire at that age—essentially the early retirement age. In addition to these programs, liberal interpretation of unemployment plan provisions allows persons to retire with unemployment benefits before age sixty. Whether a given person is eligible for a program depends on specific plan provisions, like eligibility for flexible retirement at age sixty-three. Eligibility for other programs, like disability or unemployment, is uncertain.

In Germany, the approach is to assign (or predict) eligibility probabilities for each of the programs at each age, depending on the empirical probabilities of retirement under each of the programs at that age. Then the incentive measures are weighted averages, with the weights given by the probabilities. This instrumental variable method is described in more detail in the appendix, using the situation in Denmark as an example.

Parameter Estimates

We do not attempt in this introduction to provide a detailed discussion of the estimates. Rather, we rely on the simulations based on the parameter estimates to indicate the implications of the estimated models. The simulations are later discussed. Here we highlight the strikingly common finding in virtually all the country papers: The retirement incentives inherent in most social security programs are strongly related to early retirement.

The estimation results are summarized in table 2. For each incentive specification, the table shows the sign and the statistical significance (at the 5 percent level) of the estimated effect of the incentive measure. The table also shows the sign and the statistical significance of SSW. For each incentive measure, the sign and significance level are shown when linear age is used and when the age-specific indicator variables are used.

The results in table 2 are striking. In ten of the twelve countries, almost all of the estimated incentive measure effects are negatively related to retirement and significantly different from zero. (With respect to the following discussion, it is also notable that the sign and significance of the incentive measures rarely depends on whether age indicator variables are used in the specification.) In two countries—Italy and Spain—the peak value and option value effects are typically not significant and sometimes of the wrong sign.⁴ Also in these two countries, the single year accrual effect is negative and significantly related to retirement in four of the six cases.

4. In the United Kingdom, the option value incentive measures are significant when a “bootstrap” method, which accounts for repeated observations on the same person, is used to calculate standard errors. Also in the United Kingdom, both the peak value and the option value incentive measures are very significant—under conventional standard error estimates—when cohort indicator, instead of age indicator, variables are used.

Table 2 Summary of Parameter Estimates, by Country and Specification

Country	Estimation Method (Incentive Measure)					
	Single Year Accrual		Peak Value		Option Value	
	Linear Age	Age Indicators	Linear Age	Age Indicators	Linear Age	Age Indicators
Belgium	ACC:-*	ACC:-*	PV:-*	PV:-*	OV:-*	OV:-*
	SSW:-*	SSW:-*	SSW:-*	SSW:-*	SSW:-	SSW:-
Canada	ACC:-*	ACC:-*	PV:-*	PV:-*	OV:-*	OV:-*
	SSW:+*	SSW:+*	SSW:+*	SSW:+*	SSW:+*	SSW:+*
Denmark	ACC:+*	ACC:-*	PV:-*	PV:-*	OV:-*	OV:-*
	SSW:+*	SSW:+	SSW:+*	SSW:+*	SSW:+*	SSW:+*
France	ACC:-*	ACC:-*	PV:-*	PV:-*	OV:-*	OV:-*
	SSW:-*	SSW:-*	SSW:-*	SSW:-	SSW:-*	SSW:-*
Germany	ACC:-*	ACC:-*	PV:-*	PV:-*	OV:-*	OV:-*
	SSW:-	SSW:-	SSW:-	SSW:-	SSW:-	SSW:-*
Italy	ACC:-*	ACC:-*	PV:+	PV:-	OV:+	OV:-
	SSW:-*	SSW:-	SSW:+	SSW:-	SSW:+	SSW:+
Japan	ACC:-*	ACC:-*	PV:-*	PV:-*	OV:-*	OV:+
	SSW:+	SSW:-	SSW:-	SSW:-	SSW:+	SSW:-
The Netherlands	ACC:+*	ACC:+*	PV:-*	PV:-*	OV:-*	OV:-*
	SSW:+*	SSW:+	SSW:+*	SSW:+*	SSW:+*	SSW:+*
Spain	ACC:-*	ACC:+	PV:-*	PV:+	OV:-	OV:+
	SSW:+*	SSW:+*	SSW:-	SSW:+	SSW:+	SSW:+
Sweden	ACC:-*	ACC:-	PV:-*	PV:-*	OV:-*	OV:-*
	SSW:+*	SSW:+*	SSW:+*	SSW:+*	SSW:+*	SSW:+
United Kingdom	ACC:-*	ACC:-	PV:-	PV:-	OV:-*	OV:-*
	SSW:+*	SSW:+*	SSW:+*	SSW:+*	SSW:+*	SSW:+*
United States	ACC:+*	ACC:+	PV:-*	PV:-*	OV:-*	OV:-*
	SSW:+*	SSW:+	SSW:+*	SSW:+	SSW:+	SSW:-*

Notes: See text for explanation of abbreviations. Regarding the United Kingdom, the option valued estimates are significant when standard errors accounting for repeated observations for the same person are used and when cohort indicator variables, instead of age indicators, are used. Both the option value and the peak value incentive measures are very significant.

Spain: The indications in this table pertain to the Regimen Especial Trabajadores Autonomos (RETA) employee group.

*Statistically significant at the 5 percent level.

The estimated effect of SSW, however, is often not statistically different from zero and in many cases is of the wrong sign. In many countries, it is likely easier to identify the effect of the incentive measures than the effect of wealth levels. Because of program provisions, there is much more variation in the incentive measures than in SSW.

Thus, overall, the results from these twelve separate analyses seem amazingly consistent to us. The incentives inherent in retirement income programs are clear determinants of individual retirement behavior. The estimates themselves strongly suggest a causal interpretation of the cross-country results presented in our first volume. The results point to an im-

portant relationship between incentive effects and labor force participation, independent of cultural differences among countries. The magnitude of the implied effects are also vary comparably across countries, as shown by the simulations discussed later.

Eligibility Ages and Retirement: A Key Empirical Regularity

The effect on retirement of the changes in benefit eligibility ages perhaps presents the most difficult prediction challenge of the project. Thus, we give special attention to a consistent empirical regularity in retirement that highlights this challenge. In each country, retirement rates are strongly related to particular eligibility ages prescribed in country-specific plan provisions. Perhaps most importantly, retirement rates increase sharply at ages of first eligibility for benefits. The age of first eligibility may differ from person to person and varies by program (e.g., social security, disability, or unemployment) in many countries. In the absence of eligibility for benefits, retirement is rare in each of the countries. One way to see this relationship is to consider retirement hazard rates. The hazard rate shows the proportion of persons employed at a given age who retire over the subsequent year. The empirical regularity between hazard rates and eligibility ages across countries is shown in some detail in the first volume of the project. Here, we show additional country examples that help to motivate, in particular, the simulation and estimation methods used in this volume.

United States

Labor force departure rates for men in the United States are shown in figure 6. The hazard rates are close to zero before age fifty-four and then in-

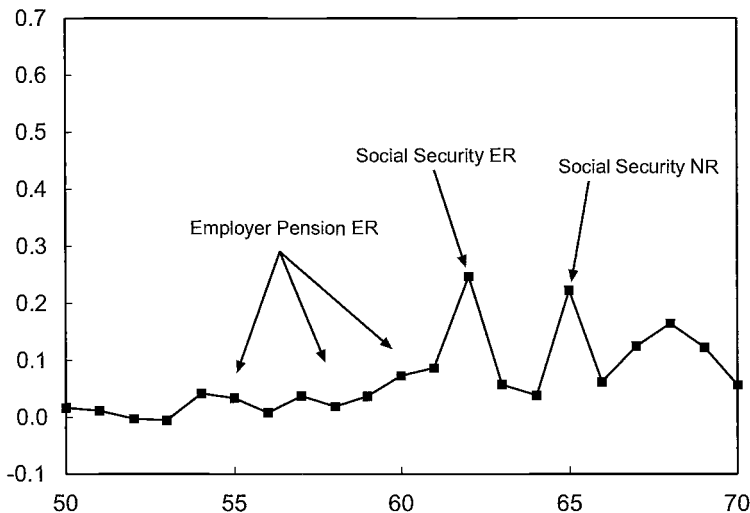


Fig. 6 Hazard rates for men in the United States

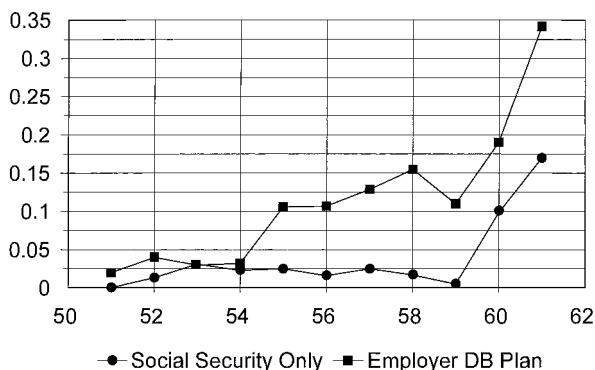


Fig. 7 U.S. retirement hazards for men, social security only and employer defined benefit

crease gradually through age sixty-one. All employees are eligible for social security early retirement benefits at age sixty-two, and there is a sharp jump in the hazard rate at that age. The rise in departure rates at age fifty-four can be attributed not to the social security program, but to eligibility for early retirement under employer-provided pension plans.

The importance of eligibility can be seen clearly in figure 7. This figure shows hazard rates by pension plan coverage. The hazard rates for men who have no private pension coverage and are covered only by the social security program are indicated by the line with round markers. The hazard rates for these persons are very close to zero until the social security early retirement age, when they are first eligible for retirement benefits. At that age, there is a sharp increase in the departure rate. The important feature of the pattern is that there is essentially no retirement before that age.⁵

The other line in the figure shows the hazard rates for men who are covered by an employer-provided defined benefit pension plan. The early retirement age under these plans is often age fifty-five and is rarely over age sixty. The hazard rates are very low before age fifty-five. But for these employees, there is a sharp jump in the hazard rate at age fifty-five, when many in this group are first eligible for benefits. And then another jump at sixty-two, when social security benefits are first available. If the early retirement age for the U.S. social security program were raised from sixty-two to sixty-five, for example, these data strongly suggest that the jump in the hazard rate at age sixty-two would no longer occur at that age, but would shift to age sixty-five instead. The critical question is whether the hazard rates would remain close to zero until age sixty-two.

5. These rates are based on labor force participation rates of Health and Retirement Study respondents. The precise age of departure from the labor force is obscured by the two-year interval between survey waves, and thus the jump in the hazard rate does not match the early retirement age of sixty-two exactly.

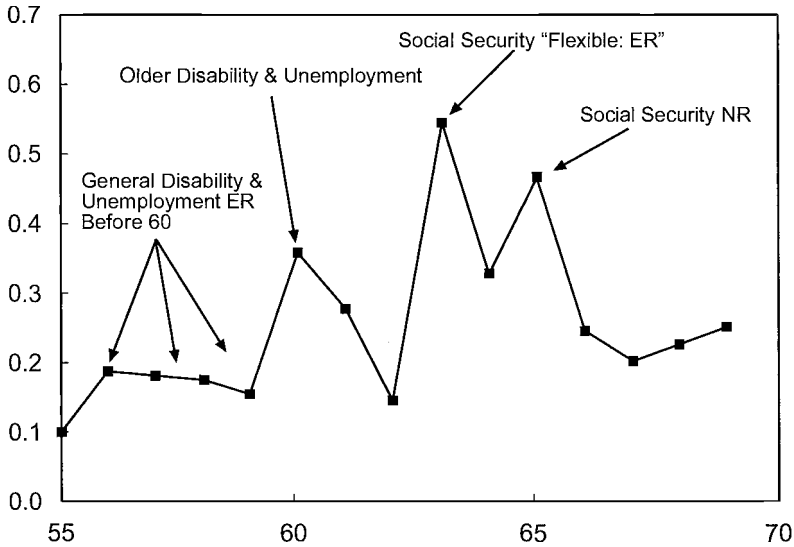


Fig. 8 Hazard rates for men in Germany

Germany

The German social security system provides additional and perhaps better examples. Figure 8 is analogous to figure 6 for the United States and shows hazard rates, with respect to departure from the labor force, for men in Germany. The detailed provisions of the 1972 social security legislation (effective until 1998) are mirrored in the retirement rates by age. In particular, there is a jump in the hazard rate at each important eligibility age. The ages of key plan provisions are also noted on the figure so that the correspondence between provisions and retirement is easily seen. Men who are disabled or unemployed at age sixty and have a certain number of years of employment under the social security system are eligible for early retirement at that age. There is a corresponding large jump in the retirement at that age. Men who have been employed for thirty-five years are eligible for early retirement at age sixty-three, and there is a corresponding jump in the retirement rate at that age. The normal retirement age is sixty-five and all men are eligible for benefits at that age. Again, there is a corresponding spike at that age as well. (By age sixty-five, however, fewer than 29 percent of men are still in the labor force in Germany.) In addition, even before age sixty, liberal interpretation of disability and unemployment plan provisions effectively serves to provide early retirement benefits, so that many men are eligible for effective early retirement before age sixty. There is a corresponding jump in the hazard rates between fifty-five and fifty-six.

Like the U.S. data, the German data also suggest that if the eligibility

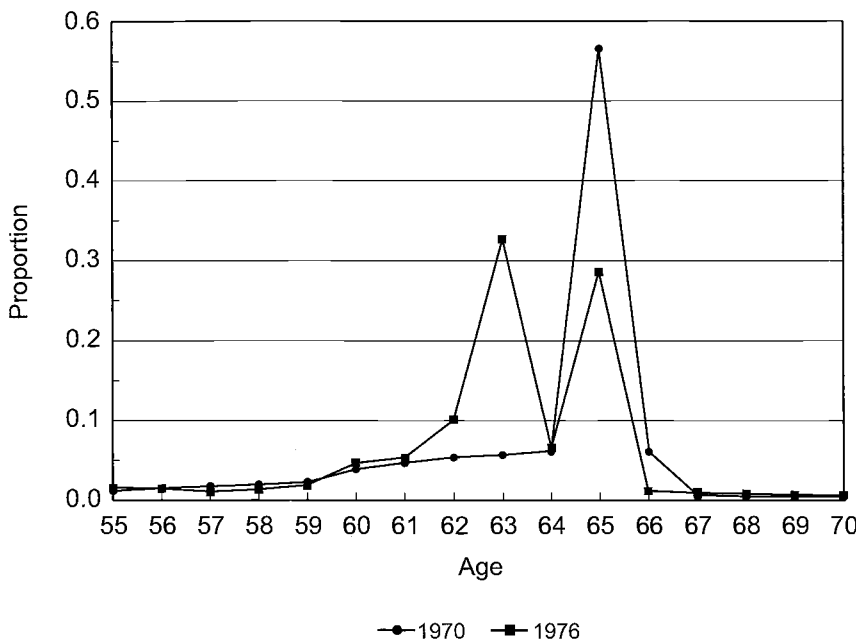


Fig. 9 Germany: Retirement ages, ages before and after 1973 reform

Source: Data provided by Axel Börsch-Supan.

ages were increased, for example, the observed jumps in the hazard rates would also shift upward. This sort of pattern is exhibited in all countries and is shown in detail in the first volume of the project.

The German system provides additional examples that reinforce the importance of plan-specific eligibility ages. First, the provisions of the German social security program were changed in 1972, as mentioned above. Second, the provisions for men and women are different. The difference in the pre- and post-reform hazard rates, as well as the difference in the hazard rates for men and women, highlight the point. Figure 9 shows the distribution of retirement ages for men under the social security system provisions in 1970 and in 1976. Before the 1972 reform, retirement under the social security program was essentially only possible at the normal retirement age of sixty-five, and there is a correspondingly large spike in the distribution at that age. The 1972 reform provided for early (flexible) retirement at age sixty-three, and in 1976 there is a large concentration of retirement at that age. Notice that retirement at ages before sixty and after sixty-five was essentially unaffected by the change in plan provisions. The change in hazard rates was essentially confined to the ages affected by the legislation. (These data do not cover retirement under the disability and




Fig. 10 Germany: Retirement ages 1973, 1976, and 1980

Source: Data provided by Axel Börsch-Supan.

unemployment programs before age sixty. Retirement under these programs, however, is evident in figure 8 and in figure 5.)

Figure 10 is like figure 9, but adds the distribution of retirement ages in 1980. After the 1972 reform, men quickly took up retirement at the early retirement age of sixty-three—as seen in 1976. Over the next few years, retirement through lenient disability and unemployment rules was increasingly taken advantage of, and by 1980 a large retirement concentration at age sixty (through these programs) is evident. Apparently, the ease of retirement at the earlier age was not at first recognized.

In addition, the eligibility ages for men and women differ. The effect of the differences can be seen in figure 11. This figure shows the distribution of retirement ages, under the social security system, for men and women in 1995. For both men and women there is a concentration of retirement at age sixty-five, the normal retirement age. For both men and women, there is also a concentration at age sixty, but for different reasons. Age sixty is the early retirement age for women, but for men, the concentration at age sixty is due to eligibility for disability and unemployment benefits at that age. Women are also eligible for these benefits at age sixty, if they have enough years of employment. The key feature of this figure is the retirement behavior at age sixty-three. For men there is a concentration at age sixty-three, the early (flexible) retirement age. But this option is not available for




Fig. 11 Germany: Retirement ages for men and women in 1995

Source: Data provided by Axel Börsch-Supan.

women, and there is no concentration at age sixty-three for women (or, too few women have enough employment years to retire at that age).

The United Kingdom

The U.K. program also has different provisions for men and women. Men can begin to receive benefits under the public social security program at age sixty-five; women can begin to receive benefit at age sixty. These differences are clearly reflected in the retirement patterns of men and women, as shown in figure 12. This figure shows the labor force “survival probabilities” for men and women who do not have an occupational pension. The proportion of women employees still in the labor force drops by about 20 percentage points (from 60 to 40 percent) at age sixty, but there is essentially no decline for men at this age. On the other hand, there is a 20 percentage point drop (from about 40 to about 20 percent) for men at age sixty-five, when they can receive benefits.

Thus within-country differences in labor force departure rates by gender, by pension plan coverage, and over time show clearly that retirement is strongly influenced by eligibility ages. It seems clear that differences in labor force departure rates among countries are also strongly influenced by differences in eligibility ages. We would like the estimation and simulation methods used in the analysis to capture the essence of the empirical regularity between benefit eligibility ages and retirement and, at the same time,




Fig. 12 U.K. survival rates for men and women

capture the effect of the other plan incentive measures as previously defined. We give particular attention to this issue later.

Simulations: Method and Results

Simulation Method

Perhaps the best way to judge the implications of the country estimates is to consider the simulations based on the estimates. Thus, the main focus of this introduction is on the two illustrative simulations that are performed for each of the countries. The first simulation predicts the effect of delaying *all* program benefit eligibility ages by three years. In countries in which disability, unemployment, or other retirement pathways are important, the eligibility age for *each* of the programs is delayed by three years. The second simulation is intended to predict the effect of the same reform (the “common reform”) in each country. Under the common reform, the early retirement age is set at age sixty and the normal retirement age at sixty-five. Benefits taken before age sixty-five are reduced actuarially by 6 percent for each year before age sixty-five. Benefits taken after age sixty-five are increased by 6 percent for each year the receipt of benefits is delayed. In addition, the replacement rate at age sixty-five is set at 60 percent of (projected) age-sixty earnings.

The simulations are summarized in the table 3. The country papers show simulations done in nine different ways. For each reform, simulations are

Table 3 Simulations

Simulation Method	Simulation and Estimation Method					
	Three-Year Delay in Eligibility Ages			Common Reform—ER at 60, NR at 65, replacement rate of 60% (of age 59 earnings) at 65, 6% reduction before 65, 6% increase after 65		
	Accrual	Peak Value	Option Value	Accrual	Peak Value	Option Value
S1: Without age indicators	X	X	X	X	X	X
S2: With age indicators but without increment age effects	X	X	X	X	X	X
S3: With age indicators and increment age effects	X	X	X	X	X	X

Notes: ER = early retirement age; NR = normal retirement age.

based on the accrual, the peak value, and the option value methods. For each of these estimation methods, three simulation approaches are used. The three simulation methods are based on different uses of the age indicators. In every case the simulations are done by recalculating the relevant incentive (accrual, peak value, or option value) to correspond to the program change. In addition to the incentive measure per se, the SSW measure and wage earnings are also recalculated to correspond to the change in the program. Then the retirement rates are reestimated using the new measures.

The key feature that distinguishes the three methods is the use of age indicators. The first simulation method (S1) does not use age indicators at all, either in the estimation or in the simulation. Only a linear measure of age is used as a control variable. In this case, only the incentive measures (and the related variables) are recalculated to simulate the effect of the reforms. The second simulation method (S2) uses age indicators in the estimation, but does not use the age indicators in the simulation. The third method (S3) uses age indicators in the estimation and, in addition, uses adjusted age indicators to simulate retirement under the program changes.

Method S2 likely minimizes the simulated effect of the program changes. The effect of the incentive measures is estimated conditional on the individual age indicators, but in predicting the effect of program changes, the simulations account only for the change in the incentive measures and do not account for the age effects. This will clearly understate the true effect of the program changes, assuming that there are important program eligibility age effects.

Simulation method S3 may typically yield the largest simulated effect of the program changes. The estimated age indicator effects, as well as the

program incentive effects, are used to predict the effect of the program changes. For example, for the three-year eligibility delay, the age indicator for a given age is taken to be the estimated age indicator three years prior to the given age. The age-sixty indicator, for example, is taken to be the estimated age-fifty-seven indicator. The result is that under the three-year eligibility delay, the projected retirement rate at age sixty is *approximately* the same as the current program age-fifty-seven retirement rate. The spike at the early retirement age under the current program, for example, shows up there years later under the reform. This approach assumes that *all* of the estimated age effects can be attributed to the eligibility-age program provisions. (The ages include the age at which persons are eligible for one or more programs, as well as the normal retirement age.)

Method S1 perhaps provides a middle ground. In this case, the estimation method does not explicitly allow for increases in retirement at given eligibility ages, and thus these effects are not allowed to influence the simulated effect of program changes. These effects will only be reflected in the simulations if they are captured by the estimated incentive effects. Recall that in our terminology, the incentive measures are the option value, the peak value, and the single-year-accrual financial measures; they do not include the eligibility age effects per se, which may also reflect additional incentives. Changing the early entitlement age by three years, for example, will change the incentive measures and SSW at every age, and the effect of this change is captured by the S1 simulations. But any other eligibility-age effects, such as social norms, liquidity constraints, or other reasons to retire at given eligibility ages, are not captured by the S1 simulations.

As the simulation results reported afterwards show, method S3 most often yields the largest estimated effect of program changes—but this is not always the case. In several instances, method S1 yields larger effects than method S3.

In the next section of this introduction, we discuss simulation results based on the option value (OV) estimation results and using the S1 and the S3 simulation methods, marked by an **X** in the table. Without undue complexity, this allows us to describe the general features of the results and to direct attention to the most important issue in estimation and simulation—the use of age indicators. We also focus solely on the results for men for expositional convenience, but results for women as well as men are presented in each country's chapter.

Simulation Results

We begin by considering the results for the three-year delay in program eligibility. We first show results based on S3. Then, for the three-year delay simulation, we compare results based on S1 and S3. Next, we consider the predicted effects of the common reform, and then we compare results under the common reform and under the three-year eligibility delay. Before

proceeding to cross-country comparisons, however, we briefly explain why country-to-country differences in simulated results should be expected.

Differences Across Countries

Although the overall simulated effects of the illustrative reforms are large in all countries, the magnitude of the effect differs from country to country. There are several reasons for the differences: The first and the single most important reason is that the current programs differ substantially among countries, and thus the effect of given reforms should differ as well. A second reason is that the data files upon which the estimates are based differ from country to country; in a few countries, the data pertain to only a portion of the workforce. A third reason is that there may be differences across countries in individual responses to a given incentive. A fourth reason, related to the second, is that the precise calculation of the incentive measures may differ somewhat from country to country. A fifth reason is that the precise implementation of the simulations may differ among the countries. It is, of course, not possible to apportion the quantitative effect of each of these reasons.

The models described above are used to predict the effect of program incentives (and other variables) on hazard rates, or the likelihood that a person in the labor force at a given age will leave the labor force at that age. The simulations begin with base hazard rates, that is, the predicted hazard rates under the current program. Then new hazard rates are predicted based on the provisions of the illustrative reform. These hazard rates are then used to predict the proportion of persons out of the labor force at given ages, and these proportions are used to determine the proportion of persons out of the labor force in given age ranges, such as fifty-six to sixty-five. Thus any of the reasons for differences that affect the base hazard rates (under the current program) or the predicted hazard rates (under the illustrative reforms) will lead to different results among the countries.

Three-Year Delay in Eligibility

Turning to cross-country comparisons, we begin with simulations of the effect of a three-year delay in eligibility. Results for men aged fifty-six to sixty-five are shown in figure 13. The figure shows the out-of-the-labor (OLF) percentage for the base case and under the eligibility delay. The eligibility delay estimates are based on S3. In all countries there is a noticeable reduction in the proportion of men out of the labor force when the eligibility ages are increased by three years.

The comparison among countries may be confounded, however, by the wide variation across countries in the age at which retirement begins. Thus, the change in the proportion out of the labor force may vary more among countries at younger ages than over the entire fifty-six-to-sixty-five age range. To help to standardize for this effect, we define the first age at which

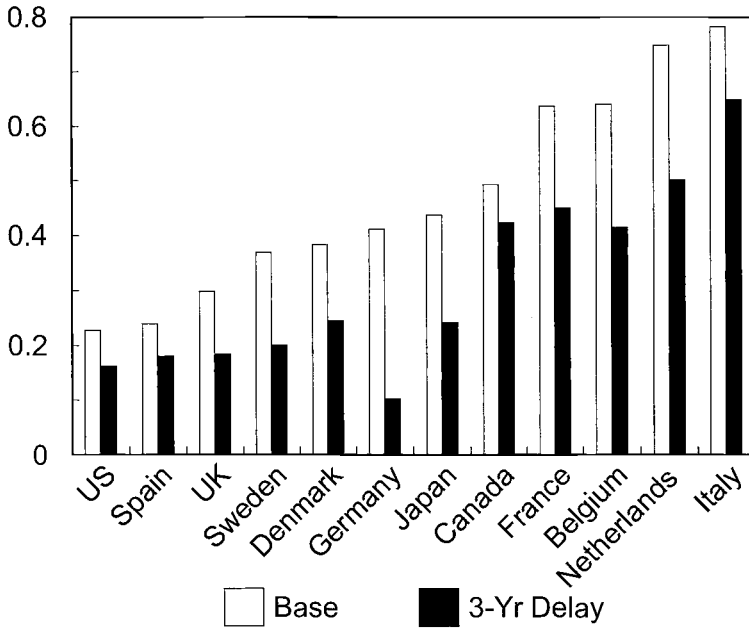


Fig. 13 OLF ages 56–65: Base versus three-year delay, OV-S3

at least 25 percent of men are out of the labor force, which we call the “25 percent age.” Then we consider the five years beginning with the “25 percent age”—“25 percent age plus four years.”

The results for the “twenty-five percent age plus four years age” range are shown in figure 14, in which the “25 percent age” is shown at the top of the bars for the base case. The “25 percent age” ranges from age fifty-three in Italy to age sixty-two in Spain and the United States. Within the “twenty-five plus four years” age range, the OLF proportion is currently (under the base case) between 40 and 50 percent in ten of the twelve countries. Within this more standardized age range, there is typically a greater reduction in the OLF percent—when eligibility is delayed—than for the fifty-six-to-sixty-five age range.

The percent reduction in the OLF proportion for the “twenty-five percent age plus four years” age range is shown in figure 15. The average reduction is 47 percent, with a range from 14 percent in Canada to 77 percent in Germany.

Figure 16 is a comparison of results under S1 and S3. The figure shows reductions in the OLF proportion for the “twenty-five percent age plus four years” age range. Recall that under S1, predictions are based on changes in the incentive measures alone, while under S3, age-specific indicators are used as well.

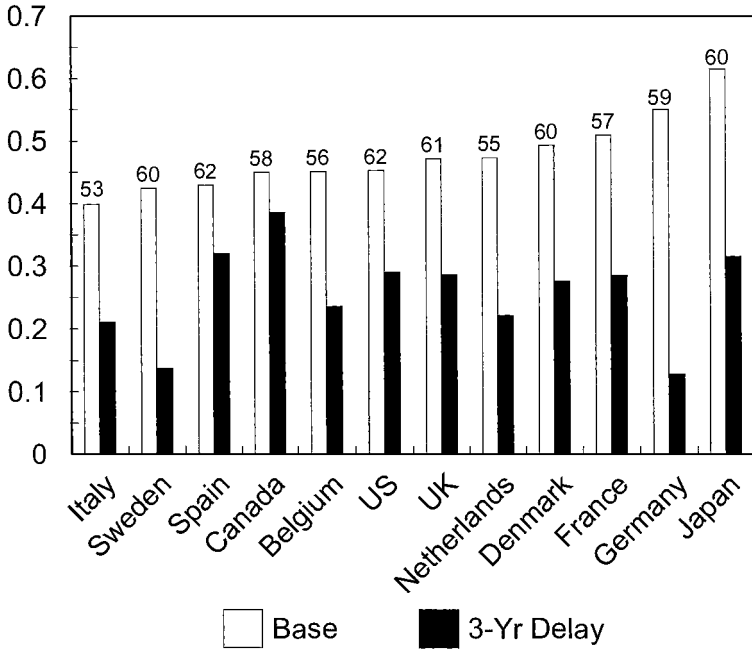


Fig. 14 OLF 25 percent age plus four years: Base versus three-year delay, OV-S3




Fig. 15 OLF change 25 percent age plus four years: Base versus three-year delay, OV-S3




Fig. 16 OLF 25 percent age plus four years: Three-year delay, OV-S1 and OV-S3

There are two notable features of this figure. The first is that the overall reductions are large under either method, and in many of the countries, the two methods yield quite similar results. The average reduction is 47 percent under S3 and 28 percent under S1. Under S3, the reduction is at least 34 percent in eleven of the twelve countries. Even using S1, the reduction is greater than 23 percent in six of the twelve countries (and in two of these countries the reduction is greater under S1 than under S3). The reduction is between 12 and 18 percent in four countries. Only in Italy and Japan is the estimated effect under S1 quite small. For the first six countries—Spain, the Netherlands, France, Canada, Germany, and Denmark—the two methods yield rather similar results. In the first eight countries, the reduction under S1 is at least 47 percent of the reduction under S3. Thus on the whole, the reduction in the OLF proportion is large under either approach.

The second notable feature of the figure is the similarity across countries in the reduction under S3. The reduction is between 34 and 55 percent in nine of the twelve countries (in Germany and Sweden the reductions are 77 and 68 percent, respectively). This similarity reflects the similarity in the estimated age effects at program eligibility ages. In all countries, there are spikes in the hazard rates at these program eligibility ages similar to those for the United States and Germany (shown in figures 6 and 8).

In summary, for almost every country, the estimates under either method show very large reductions in the OLF proportion when program eligibility ages are raised. The reductions, however, are typically larger under S3, which allows age-specific variables to capture the effect of benefit eligibility on retirement. Based on the strong empirical regularity between retirement rates and program eligibility ages, as discussed previously and in the first volume, we believe that S3 likely provides the more tenable predictions of the long-run reductions in the OLF proportions (although responses to *increases* in eligibility ages may not parallel responses to *reduction* in eligibility ages). However, even under the more restrictive S1—which does not directly allow for eligibility age effects—the predicted effect of the delay in eligibility ages is large in almost all countries.

Common Reform

We turn now to simulation results for the common reform. Recall that the common reform has four key features: (a) It sets the normal retirement age at sixty-five; (b) it sets the early retirement age at sixty; (c) benefits are reduced actuarially if they are taken before age sixty-five; and (d) the replacement rate at the normal retirement age is set (approximately) at 60 percent of age-fifty-nine earnings. Figure 17 shows the OLF proportion for

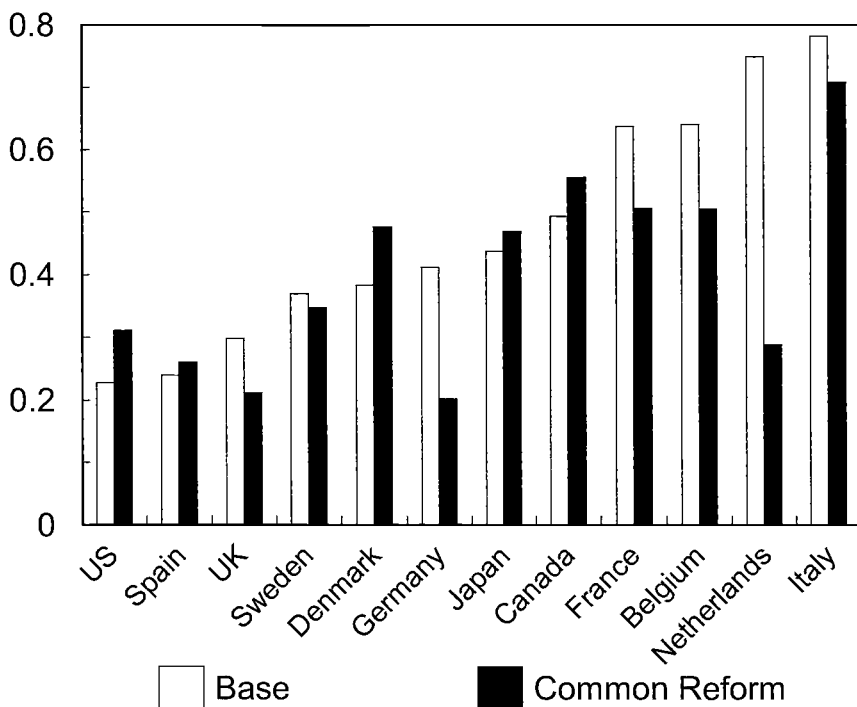


Fig. 17 OLF ages 56–65: Base versus common reform, OV-S3

the fifty-six-to-sixty-five age range under the base case and under the common reform, based on S3. Only in Germany and in the Netherlands is the OLF proportion reduced substantially. In five of the twelve countries, the common reform increases the OLF proportion. But there appears to be no clear pattern between the base-proportion OLF and the reduction in the OLF proportion under the common reform. The variation in the OLF proportion, relative to the base, is reduced under the common reform, but substantial variation across countries remains.

For the “twenty-five percent plus four years” age range, however, there is a strong pattern regarding the change in the OLF percent. The OLF proportion under the base and under the common reform for this age range is shown in figure 18. In this figure, the countries are ordered by the “25 percent age.” For the “twenty-five percent plus four years” age range it is clear that the greatest reductions in the OLF proportion under the common reform are realized in the countries with the lowest “25 percent age” rates. The change in the OLF proportion in the “twenty-five percent” age plus four years age range is shown in figure 19. For the six countries with a “25 percent age” less than sixty, the average *reduction* in the OLF proportion is 44 percent. For the six countries in which the “25 percent age” is sixty or more, there is, on average, a 4 percent *increase* in the OLF proportion.

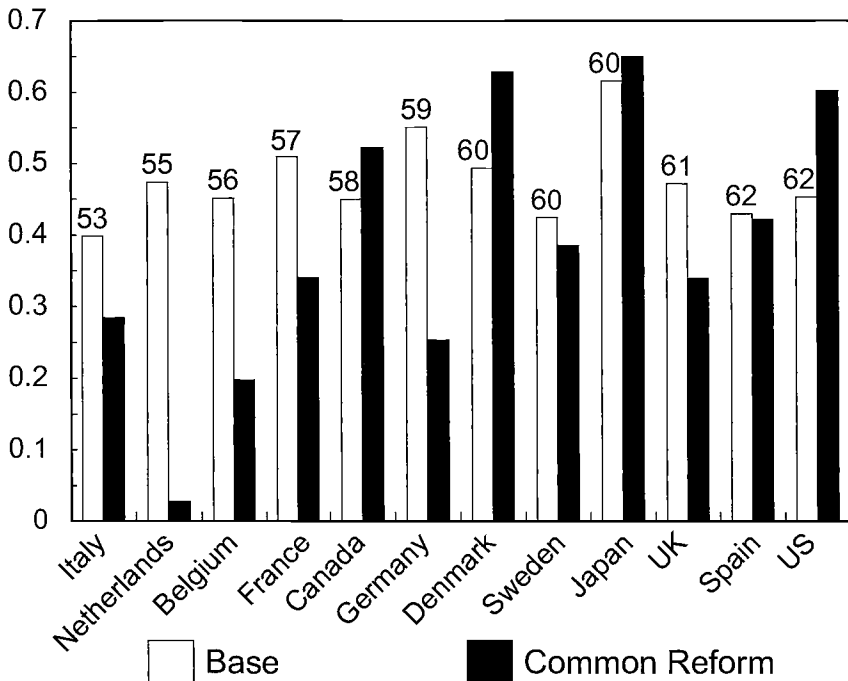


Fig. 18 OLF 25 percent age plus four years: Base versus common reform, OV-S3




Fig. 19 OLF change 25 percent age plus four years: Base versus common reform, OV-S3

The systematic pattern of these results shows a strong correspondence with intuition. For the countries with the lowest “25 percent age,” the common reform represents a substantial increase in the youngest eligibility age, and the actuarial reduction means that benefits at this age are much lower than under the base country plans. Thus, for these countries, the OLF proportion should decline under the reform, which is the case for every country but Canada. But for the countries with a “25 percent age” of sixty or greater, the common reform may reduce the earliest eligibility age—as in the United States—and may provide a greater incentive to leave the labor force. In addition, the 60 percent replacement rate at the normal retirement age represents an increase for some countries, like the United States, and a reduction in the replacement rate for other countries. Consequently, in three of these six countries, there is an increase in the OLF proportion under the common-reform simulation, and on average there is an increase in the OLF proportion. The seemingly anomalous result for Canada is explained by the fact that Canada is the only country in which the “25 percent age” is below the nominal social security entitlement age; the “25 percent age” is fifty-eight, while the social security entitlement age is 60. In




Fig. 20 OLF change 25 percent age plus four years: Base versus common reform, OV-S1 and OV-S3

addition, Canada has relatively low benefits at the early retirement age (age 60). Thus the common reform significantly increases benefit levels, providing an additional inducement to retirement.

The simulated changes under the common reform based on S1 and S3 are compared in figure 20. In each of the countries, both methods either predict a reduction or an increase in the OLF proportion. Overall, the magnitude of the simulated changes based on the two methods is rather similar as well. The most apparent exceptions are Italy, Canada, and Germany. Both methods, on average, show reductions in the OLF proportion in the six countries with the lowest “25 percent age” rates and small changes in the OLF proportion for the six countries with the highest “25 percent age” rates. The differences between the groups are more muted, however, under S1. Based on S3, the average OLF change is –44 percent for the first six countries and 4 percent for the last six, as noted previously. Based on S1, the OLF change is –21 percent for the bottom six countries and –4 percent for the top six. Most of the difference between the methods is accounted for by the differences for Italy, Belgium, Germany, and Canada. (The anomalous result for Canada is already explained.) Again, on the whole, the two methods suggest similar results. Like the simulated

effects of the three-year eligibility delay, we believe that S3 is likely to give the most reliable long run predictions.

Comparing the Three-Year Eligibility Delay and the Common Reform

Finally, figure 21 compares the results for the three-year eligibility delay with the common-reform results. The figure is based on S3. It shows the percent change from the base under the two reforms for the “twenty-five percent plus four years” age range, which is shown as the label on the common-reform bars. Recall that the three-year delay reduces the proportion out of the labor force in all countries. The average reduction in the OLF proportion is 47 percent, and there is little difference in the “25 percent age” (–49 percent for the six countries with the lowest “25 percent age” rates and –45 percent for the six countries with the highest “25 percent age” rates).

The results under the common reform, however, should depend on the base program provisions. As shown in figure 19, for the first six countries the average reduction in the OLF proportion is –44 percent and for the last six countries the average increase is 4 percent. In particular, under the common reform, benefits are not available in any country until age sixty. For

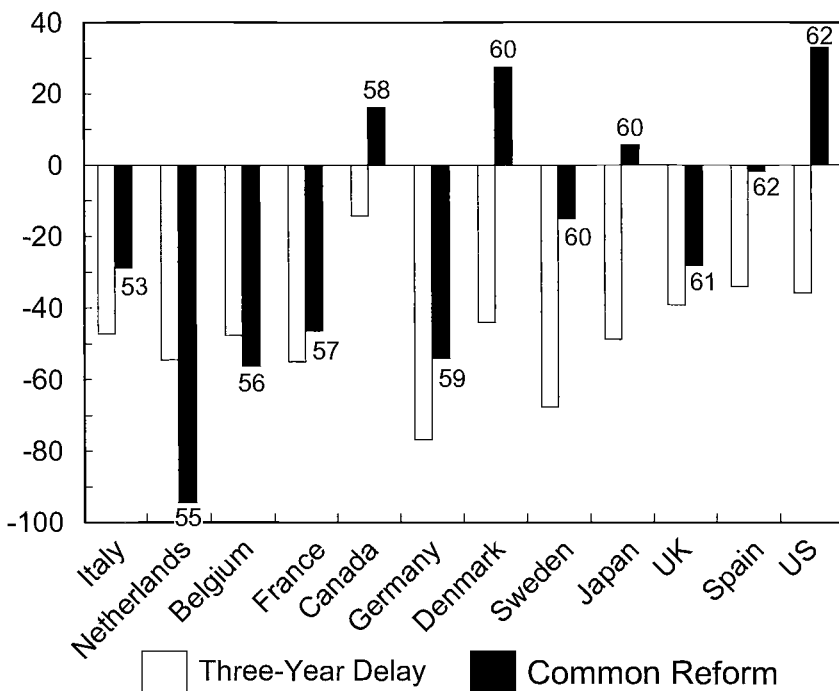


Fig. 21 OLF change 25 percent age plus four years: Three-year versus common reform, OV-S3

many countries, current benefits are available well before age sixty. (The “25 percent age” helps to identify the early eligibility countries.) In most of these countries—Italy, the Netherlands, Belgium, France, Canada, and Germany—the delay in eligibility should tend to reduce the OLF percent, and this is the result in all of these countries, with the exception of Canada. In addition, the actuarial reduction in benefits if they are taken before the age-sixty-five normal retirement age should tend to further reduce the OLF percent in most of these countries. Moreover, the common reform represents a reduction in the replacement rate in some of the countries. In other countries, such as Spain and the United States, current benefits are not available until age sixty or later. In these countries the common reform could increase the OLF proportion.

Overall, the relative effects of the two reforms are plausible, lending credence to the estimation approach. Under the three-year eligibility delay, which should reduce retirement ages in all countries, there is a reduction in the OLF proportion in each countries, and in many countries, the reduction is very large. But for the common reform, the effects should depend on country-specific program provisions, as the simulations show.

Conclusions

Our introduction to the first volume of the project concluded with a striking graph showing a strong relationship across countries between social security program incentives to retire and the proportion of older persons out of the labor force (figure 4 of this introduction). From the weight of the evidence, we judged that the relationship was largely causal. The strong response of retirement decisions to within-country changes in program provisions over time, and to different provisions for different groups at a point in time (shown in figures 6–12), also point to a casual relationship between program provisions and retirement.

The results of the country analyses reported in this volume confirm the strong causal effect of social security program retirement incentives on labor force participation. But perhaps more importantly, the results in this volume show the large magnitude of these effects. Across twelve countries with very different social security programs and labor market institutions, the results consistently show that program incentives accord strongly with retirement decisions. The magnitude of the estimated effects varies from country to country, but in all countries, the effects are large.

The magnitude is illustrated most clearly by the simulations reported in each country’s paper, and we have emphasized the simulations in this introduction. Considering the *average* across all countries, a reform that delays benefit eligibility by three years would likely reduce the proportion of men aged fifty-six to sixty-five out of the labor force between 23 and 36 percent, perhaps closer to 36 percent in the long run. For the “twenty-five per-

cent plus four years” age range, the average reduction would likely be between 28 to 47 percent, and perhaps closer to 47 percent in the long run. The effects are much larger than this in some countries, and in virtually every country, the effects are large regardless of the estimation method.

On the other hand, an illustrative common reform—with early retirement at age sixty, normal retirement age sixty-five, and actuarial reduction in benefits between sixty-five and sixty—has very disparate effects across the countries, depending on the provisions of the current program in each country. For the countries in which the current modal retirement age is younger than sixty, this reform typically implies a reduction in retirement incentives, and under this reform, the simulated proportion of older persons out of the labor force declines substantially in most countries. But for countries in which the modal retirement age is sixty or older, this reform may represent an increase in retirement incentives, and the proportion of persons out of the labor force may increase, on average, in these countries. The strong correspondence between the simulation results and a priori expectations lends credence to the estimation procedures used in the country papers.

In short, the results in this volume provide an important complement to the first volume. The results leave no doubt that social security incentives have a strong effect on retirement decisions, and the estimates show that the effect is similar in countries with very different cultural histories, labor market institutions, and other social characteristics. While countries may differ in many respects, the employees in all countries react similarly to social security retirement incentives. The simulated effects of illustrative reforms reported in the individual chapters make it clear that changes in the provisions of social security programs would have very large effects on the labor force participation of older employees.

In the next stage of the project, we will use the estimation results and simulation methods developed in this stage to estimate the financial implications of changes in program provisions.

Appendix

Incentive Measures

In this appendix, we review the relationship between the two forward-looking incentive measures—the option value and the peak value. Under the option value formulation, the value at age t of retirement r is given by

$$V_t(r) = \sum_{S=t}^{r-1} \beta^{S-t} E_t(Y_S^\gamma) + \sum_{S=r}^S \beta^{S-t} E_t(kB_S(r))^\gamma$$

using the Stock-Wise specification. Here Y is future wage income and B is social security benefit income, which depends on the retirement age r . For simplicity, the probabilities of being alive to collect the income or the benefits have been suppressed. The gain from postponing retirement to r , versus retiring at age t , is given by

$$V_t(r) = \sum_{S=t}^{r-1} \beta^{S-t} E_t(Y_S^\gamma) + \sum_{S=r}^S \beta^{S-t} E_t(kB_S(r)^\gamma) - \sum_{S=t}^S \beta^{S-t} E_t(kB_S(t)^\gamma).$$

If r^* is the retirement year that gives the maximum expected gain, the option value is given by

$$\begin{aligned} \text{OV}_t(r^*) &= \sum_{S=t}^{r^*-1} \beta^{S-t} E_t(Y_S^\gamma) + \left[\sum_{S=r^*}^S \beta^{S-t} E_t(kB_S(r^*)^\gamma) - \sum_{S=t}^S \beta^{S-t} E_t(kB_S(t)^\gamma) \right] \\ &= \left(\begin{array}{c} \text{discounted utility} \\ \text{of future wage} \end{array} \right) + \left[\left(\begin{array}{c} \text{discounted utility} \\ \text{of benefits if} \\ \text{retire at } r^* \end{array} \right) - \left(\begin{array}{c} \text{discounted utility} \\ \text{of benefits if} \\ \text{retire at } t \end{array} \right) \right]. \end{aligned}$$

Considering this equation, we can see that there are two ways to calculate the option value used in the analyses in this volume: One way is to use prior estimated values for the utility parameters γ , β , and k . The second is to set assume a value for β and to set $\gamma = k = 1$.

Multiple Pathways to Retirement and Combining Programs

In the United States, the social security program is the principle public program route to retirement. Only a small fraction of older persons enter retirement through the disability program. In some European countries, however, there are several public programs that provide routes to retirement. The case in Germany is discussed in the text and illustrated in figure 5. Thus, in considering the incentive to retire, it is important to recognize that retirement incentives under several programs may matter. The key question is which program, or programs, a person could choose to enter, out of those available to a given person. For example, who could retire under the disability program? In some instances, administrative provisions limit the universe of persons who might be eligible. In other instances, a large fraction of persons could be eligible, but which persons are eligible is unknown. Thus, the incentives facing a given individual must be estimated probabilistically. We would like to have the probability that each person is eligible for each program. Suppose that the incentive measure under each possible program is calculated for each person for each age. Then for each age, these probabilities could be used to obtain a weighted incentive measure, in which the weights are the probabilities that the person is eligible for each program. This is essentially an instrumental variable approach.

In principle, eligibility probabilities should be estimated for each person

for each age, depending on administrative rules as well as individual attributes. During the course of this part of the project, several different approaches were tried in various countries. Eligibility for disability is a good example. Based on administrative rules, it might be assumed that every person is eligible for disability beginning at some age, or it might be assumed that eligibility probabilities correspond to actual empirical take-up rates by age and other variables (where the take-up rate is estimated based on personal attributes). Here, it is implicitly assumed that the take-up rate for a person with a given set of attributes represents, on average, the eligibility rate for persons with that set of attributes. There is no correct way to do this without knowing eligibility for a sample of persons and then being about to predict eligibility. In few, if any, countries was this an option. Thus, for the purposes of the estimates in this volume, we have elected to assign weights based on empirical take-up rates. In this case the disability-eligibility probability will typically increase with age, for example. (Deviations from this method are noted in the individual chapters.)

To explain the procedure we use the situation in Denmark, which is likely the most complicated of the country situations. Here are the programs in Denmark, together with the eligibility age and information to determine eligibility, are shown in table A.1. An important calculation is the probability that a person is eligible for social and disability pension (SDP). To obtain the probability of SDP eligibility, the approach is to use actual take-up rates by age, year, and sex cells, by disability level.

Suppose the calculation pertains to pension-based SSW. The goal is to obtain a weighted measure based on the probability a person is eligible for a specific program at a given age. At sixty-seven and beyond, a person is only eligible for the old age pensions (OAP) and possibly the public employees pension (PEP). Persons who retire under other programs convert

Table A.1

Program	Eligibility Age	Determine Eligibility Based On
PEW (post-employment wage)	60 to 66	Age and insurance fund information
TBP (transitional benefits program)	55 to 59; unemployed at ages 50 to 59 during 1994 to 1996	Age and insurance fund and unemployment information in two years prior to age
PEP (public employees pension)	60 to 69	Employer pension contributions over required period
SDP (social and disability pension)	18 to 66	Probability: Based on observed participation rates by age-year-gender in each of the three levels of the disability program
OAP (old age pensions)	67	Age = 67

to OAP and start receiving benefits under that program at age sixty-seven. Before age sixty-seven, a person *could* be eligible for the post employment wage (PEW) or the transitional benefits program (TBP), but not both. Assuming that we know for sure whether a person is eligible for these programs, these are the potential sources of wealth:

SSW_{OAP} , which is available to *all* persons;

SSW_{PEW} , which could be available between ages sixty and sixty-six;

SSW_{TBP} , which could be available between ages fifty-five and fifty-nine or fifty and fifty-nine;

SSW_{PEP} , which is available to public employees who meet certain criteria; and

SSW_{SDP} , which could be available even before age fifty.

Now SSW_{OAP} can be thought of as a base that is available to everyone. The question is then what else *is* or *might* be available. Assuming that a person is not eligible for either PEW or TBP (or PEP), but that the person is eligible for SDP with probability p , then the weighted average SSW would be

$$SSW = SSW_{OAP} + p * \max [0, (SSW_{SDP} - SSW_{OAP})].$$

That is, with probability p the person has more than SSW_{OAP} (or, with probability p the person would have SSW_{SDP} and with probability $1 - p$ the person would have SSW_{OAP} only). The formula as set out above accounts for the (unlikely) possibility that SSW_{SDP} is lower than SSW_{OAP} , in which case the disability option would be ignored. (SSW_{SDP} can be collected beginning at an age much younger than sixty-seven, so it will almost surely be greater than SSW_{OAP} , which can only be received beginning at age sixty-seven.)

After age sixty, a person could be eligible for PEW, for example (if you retire before you are eligible for this program you never get these benefits). In that case, SSW is

$$SSW = \max(SSW_{OAP}, SSW_{PEW}) \\ + p * \max [0, (SSW_{SDP} - \max(SSW_{OAP}, SSW_{PEW}))].$$

That is, the certain amount in this case is $\max(SSW_{OAP}, SSW_{PEW})$. Once again, the maximum will almost surely be SSW_{PEW} , since the person can take benefits SSW_{PEW} , which incorporates benefits SSW_{OAP} beginning at age sixty-seven. With probability p , the person could be eligible for more, assuming that SDP would provide more.

A similar procedure is used to estimate SSW in two consecutive years, and from SSW in those two years, the social security accrual from one year to the next can be calculated. The peak value and option value measures are obtained in a similar fashion, but in these cases, wealth measures must be calculated for all ages into the future.

References

- Coile, Courtney, and Jonathan Gruber. 2000a. Social security incentives for retirement. In *Themes in the economics of aging*, ed. David Wise, 311–341. Chicago: University of Chicago Press.
- Coile, Courtney, and Jonathan Gruber. 2000b. Social security and retirement. NBER Working Paper no. 7830. Cambridge, Mass.: National Bureau of Economic Research, August.
- Gruber, Jonathan, and David A. Wise, eds. 1999. *Social security and retirement around the world*. Chicago: University of Chicago Press.
- Kotlikoff, Laurence J., and David A. Wise. 1985. Labor compensation and the structure of private pension plans: Evidence for contractual versus spot labor markets. In *Pensions, labor, and individual choice*, ed. David A. Wise, 55–87. Chicago: University of Chicago Press.
- . 1989. *The wage carrot and the pension stick*. Kalamazoo, Mich.: W. E. Upjohn Institute for Employment Research.
- Lazear, Edward. 1983. Pensions as severance pay. In *Financial aspects of the United States pension system*, ed. Zvi Bodie and John Shoven, 57–85. Chicago: University of Chicago Press.
- Stock, James H., and David A. Wise. 1990a. The pension inducement to retire: An option value analysis. In *Issues in the economics of aging*, ed. David A. Wise, 205–24. Chicago: University of Chicago Press.
- . 1990b. Pensions, the option value of work, and retirement. *Econometrica* 58 (5): 1151–80.