Volume Title: Advances in the Economics of Aging

Volume Author/Editor: David A. Wise, editor

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-90302-8

Volume URL: http://www.nber.org/books/wise96-1

Conference Date: May 6-9, 1993

Publication Date: January 1996

Chapter Title: Why Are Retirement Rates So High at Age 65?

Chapter Author: Robin L. Lumsdaine, James H. Stock, David A. Wise

Chapter URL: http://www.nber.org/chapters/c7318

Chapter pages in book: (p. 61 - 82)

2 Why Are Retirement Rates So High at Age 65?

Robin L. Lumsdaine, James H. Stock, and David A. Wise

Age 65 is no longer the typical retirement age. Most employees now retire before 65, and those who are covered by defined-benefit pension plans often retire well before 65. Nonetheless, a large fraction of persons who are still working at 64 retire at 65. For example, at one of the large Fortune 500 firms that are studied in this paper (Firm 3), 48% of men working at 64 retire at 65. In contrast, only 21% of men who work through age 63 retire at 64. Women at this firm show a similar increase in retirement rates, from 18% at age 64 to 41% at age 65. Similar jumps in retirement rates at age 65 are found at other individual firms and more generally in nationwide measures of labor force participation. In each of the six data sets discussed in this paper, the highest retirement rate occurs at age 65.

In a series of earlier papers, Stock and Wise (1990a, 1990b) and Lumsdaine, Stock, and Wise (1990, 1991, 1992, 1994) developed "option value" and stochastic dynamic programming models of retirement. These models have been estimated on several firm data sets. A striking feature of the estimates is the extent to which they track actual retirement patterns that often exhibit sharp jumps in retirement rates at specific ages. Indeed, the models predict very well the retirement rates under special unanticipated "window" plans designed to encourage early retirement. Although in general these models fit most spikes in the data surprisingly well, in particular at ages 55, 60, and 62, they invariably

Robin L. Lumsdaine is assistant professor of economics at Princeton University and a faculty research fellow of the National Bureau of Economic Research. James H. Stock is professor of political economy at the John F. Kennedy School of Government, Harvard University, and a 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.

Financial support was provided by the National Institute on Aging. The authors are grateful to Brigitte Madrian for providing some of the data.

underpredict the age-65 retirement rates of persons who do not retire before that age. In earlier papers we attributed the mismatch between predicted and actual rates to an "age-65 retirement effect," having in mind the influence of custom or accepted practice. This paper considers other commonly proposed explanations for the underprediction of age-65 retirement. It is difficult to directly demonstrate the influence of custom and the like. The spirit of the paper is to rule out other explanations and thus, by implication, to leave the age-65 retirement effect as the remaining possibility. The results suggest that such an effect is the only plausible explanation that cannot be rejected. In particular we conclude that the availability of Medicare at 65 does not explain the age-65 retirement jump.

The age-65 spike is in large part unexplained by our economic models of retirement, and to our knowledge is rarely explained by other models that do not force, by age-specific variables or by other means, a "fit" to the age-65 rate. Exceptions include Gustman and Steinmeier (1986, 1994) and Phelan and Rust (1993), as discussed below. There are a number of economic reasons why individuals might choose to retire at age 65. Social Security treats age 65 as the normal retirement age, and after age 65 the rate of increase in benefits is less than actuarially fair. Kotlikoff and Smith (1983) estimate that 90% of firm pension plans also treat 65 as the normal retirement age, and under many defined-benefit plans there is a strong implicit financial penalty to working past the normal retirement age, as shown by Kotlikoff and Wise (1988), for example. However, measured in terms of expected lifetime benefits, the economic incentive to retire at 65, instead of 64 or 66, for example, is not large enough to explain the age-65 rate. In particular, although our economic models of retirement—which incorporate the financial incentives implicit in the detailed provisions of firm pension plans and Social Security provisions—predict high retirement rates at age 65, these predicted rates typically fall far short of the actual age-65 rates.

In addition, Medicare eligibility begins at age 65. Thus a person not covered by employer-provided retiree health insurance has an incentive to remain in the firm until age 65 to avoid a lapse in medical insurance coverage.

The unexplained age-65 spike is important because it limits our ability to predict the effect of potential policy changes, like the planned increase in the Social Security normal retirement age from age 65 to age 67. Would there then be a spike at 67, or would it remain at 65?

We seek to quantify the age-65 retirement puzzle and to explore potential explanations for it. These include in particular the potential gap in health insurance coverage between retirement and the Medicare eligibility age. We also consider whether family status affects age-65 retirement. And we explore the possibility that our previous results were importantly affected by small samples of older workers. Because so many employees retire early, the number still employed at 65 is typically small.

None of these possibilities explains the age-65 spike, lending indirect support to the "age-65 retirement effect" explanation. To support the plausibility of an age-65 effect, we also consider the possibility that for some employees the utility cost to electing to retire at this "customary" retirement age is small. We conclude that the economic cost is indeed small for some, although for most employees it is quite large. For most employees, choosing to retire at age 65 would impose noticeable economic cost. However, for some it might not be very costly to retire at the "customary" age of 65. To the extent that this is true, the customary effect might not persist in the face of new financial disincentives for age-65 retirement.

2.1 Age-65 Retirement Rates

We review additional evidence in the literature on age-65 retirement effects. In addition, we document the spike in retirement rates at age 65 in six separate data sources, three reflecting the experience of individual firms and three based on nationally representative surveys. As emphasized above, however, it is not solely the jump in retirement rates at 65 that motivates this paper, but rather that the jump is not explained by financial considerations incorporated in formal models.

2.1.1 Previous Literature

Many previous studies have found evidence of an age-65 retirement spike. However, few have successfully fit this spike without explicitly incorporating age or age dummies as explanatory variables. Gustman and Steinmeier (1986) were successful in fitting both the age-62 and age-65 retirement spikes in data from the Retirement History Survey (RHS). However, they modeled the trade-off between labor and leisure as a smoothly increasing function of age and, importantly, did not have detailed firm pension data; thus the Social Security normal and early retirement ages were allowed to play important roles in determining the profile of retirement benefits.

Also using the RHS, Phelan and Rust (1991) calculated a frequency distribution of retirement ages. They considered six different definitions of retirement, including the year that a person first worked less than full-time, the age of first receipt of Old Age, Survivors, and Disability Insurance (OASDI), and a self-reported retirement date. Although the retirement frequency distributions differ for the different definitions, all exhibit a spike at age 65. A spike in the frequency distribution of retirement at 65, while not the same as a spike in the departure rate at that age, implies a spike in the age-65 hazard rate as well. In a subsequent study (1993) Phelan and Rust consider retirement rates of individuals with and without employer-related health insurance and find that the age-65 spike is more pronounced for individuals with health insurance.

Blau (1994) too uses the RHS in his study of labor force dynamics. Using quarterly data, he finds that a substantial fraction of individuals retire in the first quarter after their 65th birthday. He provides simulations of the sensitivity

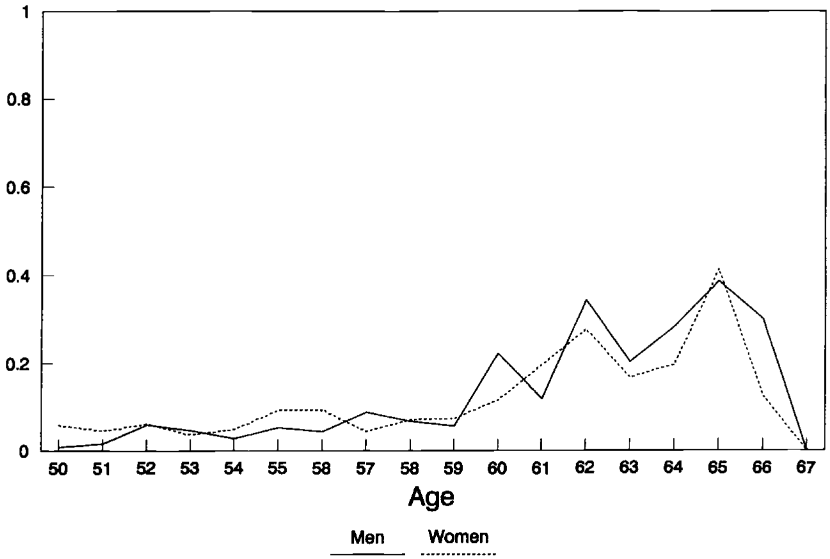

Fig. 2.1 Firm 1 hazard rates: office workers

of the age-65 retirement spike to an individual's level of Social Security benefit. Geweke, Zarkin, and Slonim (1993) also document evidence of a large spike in the probability of application for Social Security benefits in the first quarter after an individual's 65th birthday.

2.1.2 Firm-Specific Data Sets

The first three data sets are from employment records of three firms, here referred to as Firm 1, Firm 2, and Firm 3. For each firm we have data on past wages, years of service, and the details of the firm's pension plan. Depending on the firm, we also have information on occupation and some additional individual attributes. Departure rates for selected groups of employees in each of these firms are plotted in figures 2.1, 2.2, and 2.3, respectively. Although the details of each firm's pension plan differ, their overall characteristics are similar. The early retirement age is 55 and the normal retirement age 65 in each of the firms. The departure rates have generally similar shapes.

Firm 1 departure rates pertain to office workers and are shown by gender. There were 1,354 men and 2,497 women aged 50 and over in 1981.[1] Firm 3 departure rates pertain to all firm employees and are also shown by gender, with 10,221 men and 2,889 women aged 50 and over in 1982.[2] Only 718 obser-

1. Departure rates for salesmen in Firm 1 were analyzed by Stock and Wise (1990a, 1990b). In this firm the date of retirement is inferred from the year in which the employee ceased to receive a paycheck from the firm.

2. In this firm, retirement is determined by the retirement date recorded in the data set.

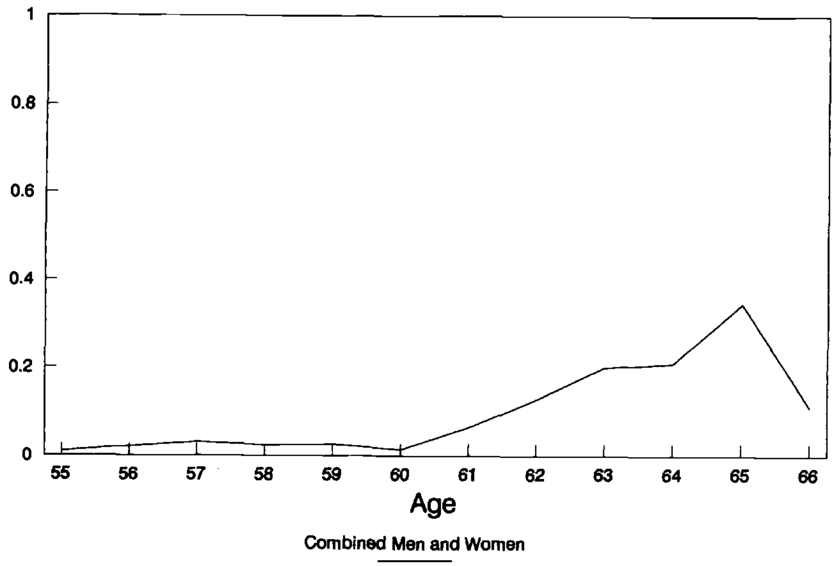

Fig. 2.2 Firm 2 hazard rates

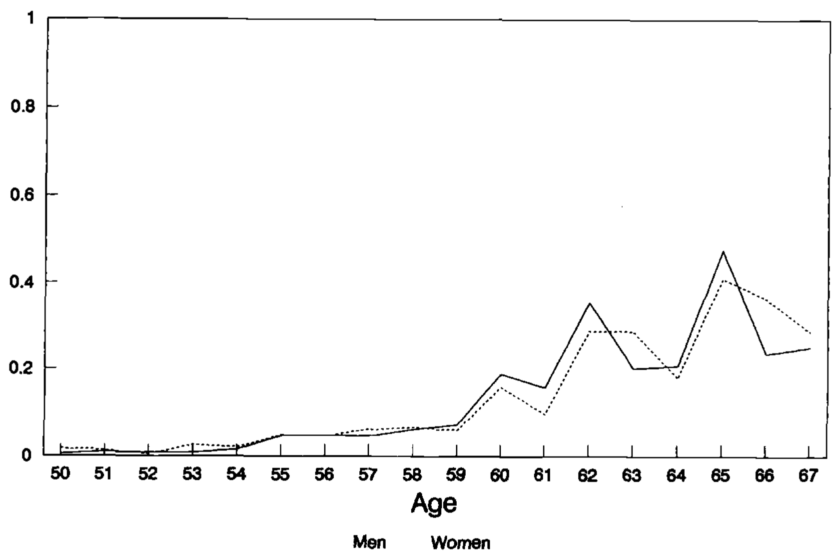

Fig. 2.3 Firm 3 hazard rates

vations of men and women 55 and over in 1979 are available for Firm 2, and the departure rates are shown for men and women combined. All three firms are large Fortune 500 companies with defined-benefit pension plans.[3]

Departure rates in Firms 1 and 3 show spikes at 65.[4] For example, at Firm 1 the retirement rate for 64-year-old men is 28%, while for 65-year-old men it is 39%. For women at Firm 1, the retirement rates at ages 64 and 65 are respectively 20% and 41%. The pattern is similar in Firm 3: the age-64 retirement rate is 21% for men and 18% for women, and rises to 48% for men and 41% for women at age 65.[5] The evidence for Firm 2 is less pronounced, although there is a noticeable spike at 65 and the general pattern of departure rates in Firm 2 is similar to the pattern in Firms 1 and 3.

2.1.3 Cross-Firm Data Sets

The high age-65 retirement rates could be typical in large firms with defined-benefit plans but not representative of the broader population of firms and workers. Evidence from three additional surveys with data on workers from many different firms, however, reveals a similar pattern in the broader population.[6] The first survey is the 1987 National Medical Expenditure Survey (NMES), a survey of approximately fourteen thousand households. The NMES respondents were asked their date of retirement. The departure rates reported here are for respondents who were at least 70 years old and retired.

The second and third data sets have been constructed from different waves (quarterly interviews) of the Survey of Income and Program Participation (SIPP): the 1984 wave with the "Education and Work History" supplemental questions (SIPP-EWH) and the 1984, 1985, and 1986 waves with the "Characteristics of Job from Which Retired" supplemental questions (SIPP-CJR). The retirement definition used in the SIPP-CJR is the date at which an individual left the firm that was providing pension benefits. In the SIPP-EWH, the retirement date is the last date worked. Hazard rates, like those based on the NMES

3. These data have been collected as part of the National Bureau of Economic Research project on the economics of aging. The data for each firm and its pension plan have been discussed in detail elsewhere, and we will not repeat the discussion here. For a discussion of the data for Firm 1 and its pension plan, see Stock and Wise (1990b). The Firm 2 data and its pension plan are discussed in Lumsdaine, Mutschler, and Wise (1992). The Firm 3 data and its pension plan are discussed in Lumsdaine, Stock, and Wise (1994).

4. For Firm 1 there are also noticeable spikes at 62 and 60, corresponding respectively to the Social Security early retirement age and to special provisions of the pension plan.

5. The data for Firms 2 and 3 have exact dates of retirement and birth, while the data for Firm 1 contain only the year of retirement and birth. This introduces some measurement error into the Firm 1 hazard rates, which is not present in the Firms 2 and 3 hazard rates. For Firm 1 the ages are computed to be the age of the individual on January 1; thus some retirements that in fact occur at age 65 (retirements by workers who began the year at 64 and retired that year but after their 65th birthday) are miscounted as age-64 retirements.

6. We are grateful to Brigitte Madrian for graciously providing us with these data, specifically the number of retirees by age for each of the three data sets, broken down by whether or not they have employer-provided retiree health insurance.

data, were computed for respondents who were at least 70 years old and retired. For additional discussion of the data sets and definitions of retirement dates, see Madrian (1993).

Hazard rates for men based on the three national data sets, as well as the three firm data sets, are reported in table 2.1. Age-65 (and age-64) retirement rates based on the national data are very similar to the rates in the individual firms, with the national age-65 rates ranging from 44 to 52%. Thus the age-65 retirement spike seen in the firm data is common nationwide.

Although an age-65 spike is revealed in each of the data sets, there are important differences at other ages. There are more early retirements in the individual firms than in the national data sets. For example, the age-60 retirement rate in Firm 3 is 19% for men, while it is less than 8% in each of the national data sets. This age-60 spike in Firm 3 corresponds to pension plan provisions that encourage retirement at age 60. More generally, the defined-benefit pension plan provisions in each of the firms provide substantial incentives to leave the firm by age 65, and usually before that. Such incentives are not common to all of the respondents in the national data sets. Indeed, only about one-half of workers are covered by employer-provided pension plans and only about two-thirds of these are covered by defined-benefit pension plans, with provisions that are likely to be similar to those in the three firms. The proportion of older workers covered is higher, however.

Another explanation is that the firm data measure separation from the firm, and some of those who leave the firm could remain in the labor force, especially those who leave the firm at younger ages. Thus, notwithstanding the

Table 2.1 **Retirement Rates for Men from Six Data Sets**

Age	Firm 1	Firm 2	Firm 3	NMES	SIPP-EWH	SIPP-CJR
55	0.054	0.000	0.048	0.022	0.026	0.016
56	0.045	0.016	0.048	0.018	0.019	0.027
57	0.089	0.017	0.047	0.025	0.019	0.019
58	0.069	0.019	0.062	0.033	0.029	0.038
59	0.057	0.023	0.073	0.048	0.035	0.054
60	0.222	0.027	0.188	0.063	0.064	0.075
61	0.119	0.069	0.157	0.110	0.107	0.095
62	0.344	0.000	0.354	0.174	0.150	0.170
63	0.203	0.174	0.201	0.160	0.146	0.188
64	0.283	0.357	0.207	0.350	0.274	0.353
65	0.386	0.357	0.476	0.441	0.456	0.518
66	0.300	0.200	0.235	0.324	0.306	0.391
67	0.000	0.667	0.250	0.424	0.378	0.404
68	0.000	1.000	0.100	0.569	0.528	0.523
N	1,354	436	10,221	1,064	979	1,190

Notes: Entries are empirical hazard rates. The data sets are discussed in the text.

firm pension provision incentives for early retirement, early departure from the worker's primary job can be consistent with later departures from the labor force, the concept measured in the NMES and SIPP-CJR data.

Because a large fraction of workers at the three firms retire before age 65, the age-65 retirement rate is based on a relatively small number of workers. In earlier work (Stock and Wise 1990b; Lumsdaine, Stock, and Wise 1992, 1994), we emphasized that, although the retirement models fit departure rates very well in general, they typically underpredict age-65 departure rates. However, because early retirement is less common nationally than in the firms we have studied, the national data sets contain a rather large sample of persons who work until age 65. And a substantial fraction of those who do, retire at 65. For example, according to the SIPP-EWH data, of those retiring between ages 55 and 69, 18% retire at age 65.

2.2 Descriptive Evidence on Early Retirement

We consider several hypotheses that have been suggested to explain the large age-65 retirement rates. First, we consider the possibility that the age-65 jump is driven by health insurance, the availability of Medicare at age 65. Second, using data on 65-year-olds from Firm 3, we estimate simple models of retirement linking the retirement decision to various economic and demographic attributes, such as marital status.

2.2.1 Health Insurance and Medicare Eligibility

Medicare coverage is available at age 65. The extent to which it provides an incentive to remain employed until 65 depends on the nature of health coverage before age 65. For employees not covered by firm health insurance, medical insurance does not affect the retirement decision beyond the need for income to meet health care costs. Most large firms, however, provide health insurance for current employees and typically also provide retiree health insurance. Coverage usually can be extended to the families of employees and retirees. For these workers, Medicare availability provides no additional incentive to postpone retirement until age 65. But many workers are in a third situation: they are covered by employer medical insurance while employed but not after retirement. They may retain the option to continue coverage through the firm group health plan but must pay the full cost of the premium. Employees in this group will have an incentive to postpone retirement until age 65, when they become eligible for Medicare. This is a possible reason for the large age-65 retirement rate.

Some evidence investigating this view has been noted recently. In chapter 4 in this volume, Gruber and Madrian estimate that the effect of continuation coverage laws is an increase in retirement rates at all ages, not just at the age of Medicare eligibility. In a more structural analysis, Phelan and Rust (1993) incorporate risk aversion and find significant differences in the retirement patterns of individuals who have health insurance versus those who do not. However, Gustman and Steinmeier (1994) estimate that, although most of the effect of health insurance on the retirement decision is via retiree health insurance (as opposed to insurance while working), this effect is quantitatively small.

In addition, we offer two types of indirect evidence that may cast doubt on this explanation. The first is based on the data from Firm 3, the firm for which we have the most detailed information about postretirement medical coverage. Retirees from the firm continue to be covered under the firm's group medical plan at no additional cost. Even beyond age 65, the plan reimburses costs not covered by Medicare. Thus lapse of coverage before Medicare eligibility would not have been a consideration in retirement from this firm. Nonetheless, as discussed in the previous section, for both men and women at Firm 3 the departure rates are highest at age 65 and more than double the age-64 departure rate.

The second type of indirect evidence is drawn from the three national data sets. Madrian (1993) has performed a careful analysis of the effect of firm-provided retiree health insurance on retirement, and we draw on her evidence here. Table 2.2 provides NMES and SIPP retirement rates for employees with and without firm-provided retiree health insurance. The striking feature of the data is the qualitative and quantitative similarity of the hazard rates for those with and without firm-provided employee health insurance. Indeed, the NMES and SIPP-EWH age-65 hazard rates are less for those without than with retiree

Table 2.2 Retirement Rates for Men: NMES and SIPP Data Sets, with and without Firm-Provided Retiree Health Insurance

	NMES			SIPP-EWH			SIPP-CJR		
Age	Without	With	Difference	Without	With	Difference	Without	With	Difference
55	.023	.019	.004	.024	.028	−.004	.024	.010	.014*
56	.015	.022	−.006	.022	.013	.009	.023	.030	−.008
57	.023	.030	−.007	.022	.013	.009	.009	.027	−.018**
58	.030	.038	−.008	.031	.026	.005	.030	.045	−.015
59	.048	.048	.000	.032	.040	−.008	.029	.072	−.044***
60	.058	.069	−.012	.061	.070	−.009	.071	.078	−.007
61	.104	.119	−.015	.111	.098	.013	.088	.100	−.011
62	.155	.203	−.048*	.140	.172	−.032	.153	.184	−.030
63	.143	.187	−.044	.155	.126	.029	.178	.197	−.019
64	.305	.424	−.119***	.258	.306	−.048	.355	.350	.005
65	.430	.464	−.034	.444	.483	−.039	.563	.480	.082
66	.294	.390	−.096	.326	.258	.068	.397	.387	.010
67	.427	.417	.010	.381	.370	.012	.409	.400	.009
N	624	440		657	322		493	697	

Notes: Entries are empirical hazard rates, computed as discussed in the text. The difference between hazard rates for employees without and with retiree health insurances is significant at the following level: * = 10%; ** = 5%; *** = 1%.

health insurance, although the differences are not statistically significant. For the SIPP-CJR, the age-65 hazard is higher for those without than with insurance, but the difference is not statistically significant. The age-65 national data hazard rates are also comparable to the rate in Firm 3, which provides postretirement health insurance. At other ages the differences between the national data hazard rates for persons with and without retiree health insurance are typically statistically insignificant. These comparisons suggest that Medicare eligibility is not an important determinant of age-65 retirement. More formal models of the retirement decision based on these national data sets have been estimated by Madrian (1993).

2.2.2 Other Determinants of Retirement at Age 65

We consider next age-65 retirement among Firm 3 employees, for whom Medicare eligibility is not plausibly an important consideration. Because we have historical payroll and demographic data for a large number of Firm 3 employees, along with complete details of retirement benefits offered by the firm, the Firm 3 data provide an opportunity to examine non-Medicare economic and noneconomic determinants of age-65 retirement. The analysis in this section is based on a sample of 203 employees who turned 65 in 1981 and were employed at Firm 3 on their 65th birthday. Of these 203 employees, 40 were women and 163 were men. The combined age-65 retirement rate for these employees was 58.1%.

The departure rates at age 65 for several subgroups of these employees are summarized in table 2.3. Only a few of the differences are statistically and substantively significant. Age-65 married men are significantly more likely to retire than single men. Employees with less than 10 years of service at the firm are less likely to retire, although there are only small differences between the retirement rates for those with 10–19, 20–29, and 30+ years of service. Those with low earnings, and therefore typically with lower pension benefits and lower lifetime wealth, are somewhat less likely to retire at 65 than those with a high income, although this difference is not significant at the 5% level. Age-65 retirement rates do not vary with the job classification.

To determine whether these factors are related to age-65 retirement, after controlling for economic variables, we estimated a series of probit retirement specifications using this sample of 65-year-olds. In addition to the variables identifying the table 2.3 subgroups, the specifications include several economic variables: predicted annual age-65 income, the expected present value of Social Security payments, the expected present value of pension payments, the change in expected Social Security wealth resulting from working an additional year (that is, retiring at age 66 rather than age 65—Social Security accrual), and the change in expected pension wealth resulting from postponing retirement until age 66 (pension accrual).

An additional economic variable is the "option value" of remaining employed an additional year, calculated using the Stock-Wise (1990b) option

Table 2.3 Age-65 Retirement Rates for Different Categories of Worker (data set: 65-year-olds at Firm 3)

Category	*N*	Rate	Category	*N*	Rate	Difference[a]
All	203	.581				
		(.035)				
Women	40	.475	Men	163	.607	−.132*
		(.079)			(.308)	(.088)
Married	142	.641	Single	61	.443	.198**
		(.040)			(.064)	(.075)
Married men	132	.659	Single men	31	.387	.272**
		(.041)			(.087)	(.096)
Married women	10	.400	Single women	30	.500	−.100
		(.155)			(.091)	(.180)
Income ≤ median	101	.535	Income ≥ median	102	.634	−.099*
		(.050)			(.048)	(.069)
<10 yrs. service	41	.463				
		(.078)				
10–20 yrs. service	50	.620				
		(.069)				
20–30 yrs. service	32	.594				
		(.087)				
>30 yrs. service	80	.613				
		(.054)				
Service job	101	.554				
		(.049)				
Technical job	50	.580				
		(.070)				

Notes: Entries are the fraction of 65-year-olds in the indicated group who retire at age 65. Standard errors are in parentheses. Differences are significantly different from zero at the following level: * = 10%; ** = 1%.

[a]Entries in the "difference" column are the difference between retirement rates for the group in the first and second category columns in that row.

value model. It is computed as the difference between the lifetime expected present value of income and retirement benefits were the individual to retire at the optimal age (which could be later than 65), and the present value of income and retirement benefits under age-65 retirement. For this calculation the optimal age of retirement is taken to be the age that maximizes the current expected value of lifetime earnings. Thus this option value is a measure of the monetary opportunity cost of retiring at age 65 rather than later. Of course, for some employees the opportunity cost is zero, that is, age-65 retirement maximizes lifetime wealth.

All of the present values—the option value, Social Security accrual, and pension accrual—were computed using forecasts of future earnings based on the past employee earnings and the firm's historical age-income profile.[7] Pres-

7. The earnings forecasting equation for employees at Firm 3 is discussed in detail in Lumsdaine, Stock, and Wise (1994).

ent values assume a discount factor of 0.90 and incorporate mortality rates calculated from life tables.

Parameter estimates are presented in table 2.4. In each specification, married male employees have a significantly higher tendency than single men to retire at age 65. Consistent with the results in table 2.3, married women are slightly less likely than single women to retire at 65, although the difference is not

Table 2.4 Probit Models of Age-65 Retirement (data set: 65-year-olds at Firm 3)

	Model				
Variable	(1)	(2)	(3)	(4)	(5)
Constant	0.018	−0.343	−1.585	−1.341	−0.644
	(0.195)	(0.386)	(1.006)	(0.927)	(0.512)
Female		0.334	0.117	0.309	0.322
		(0.337)	(0.363)	(0.343)	(0.339)
Married women		−0.261	−0.095	−0.110	−0.161
		(0.464)	(0.471)	(0.472)	(0.475)
Married men		0.691***	0.664**	0.670**	0.675**
		(0.258)	(0.262)	(0.261)	(0.263)
Service job		−0.061	0.286	0.304	0.061
		(0.249)	(0.344)	(0.350)	(0.303)
Technical job		−0.070	0.119	0.157	−0.021
		(0.286)	(0.342)	(0.350)	(0.299)
Option value	1.068	0.569	−1.199	−0.658	0.726
	(1.016)	(1.036)	(1.876)	(1.812)	(1.135)
Social Security present value			3.243		
			(2.596)		
Pension present value			0.623		
			(0.410)		
Social Security accrual				−16.134	
				(15.387)	
Pension accrual				−5.130	
				(3.042)	
10–19 years					0.221
					(0.282)
20–29 years					0.316
					(0.318)
30+ years service					0.269
					(0.302)
−ln likelihood	−137.26	−132.55	−129.56	−129.78	−131.97
Likelihood-ratio test					
Versus model 1		9.42*			
Versus model 2			5.98*	5.54*	1.16

Notes: The sample consists of 203 workers in Firm 3 who were employed by the firm when they turned 65 and who retired either at age 65 or after; the dependent variable is one if retirement is at age 65 and zero otherwise. All monetary values are in $100,000 (1980 dollars). Standard errors are in parentheses. Significant at the following level: * = 10%; ** = 5%; *** = 1%.

statistically significant. Gender, job category, and even job tenure are only weakly related to this retirement decision, both unconditionally and after controlling for the economic variables. Conditional on not retiring before age 65, although the overall age-65 retirement rate is very high, none of the economic variables is a statistically significant determinant of age-65 retirement. In contrast, these variables predict retirement at other ages. Indeed, calculations that underlie these variables play the central role in the structural option value and stochastic dynamic programming retirement models that predict very successfully retirement over all ages. This analysis provides no economic or demographic explanations for age-65 retirement decisions. The only conclusion is that married men are more likely than unmarried men to retire at 65, given that they have worked until then.

2.3 Evidence from a Structural Retirement Model

As explained above, the motivation for this series of analysis is the general failure of structural models of retirement to account for the high age-65 retirement rate, without including age-specific dummy variables or other model specifications that assure a close match to retirement rates at specific ages. We document this failure and explore alternative specifications that might, in a mechanical sense, permit the models to fit the age-65 retirement rate. We first estimate a "base" stochastic dynamic programming model of retirement, which we take as representative of the new generation of retirement models. We then modify this model to incorporate noneconomic reasons for retiring at age 65.

2.3.1 The Stochastic Dynamic Programming Model

The stochastic dynamic programming model used here is described in detail in Lumsdaine, Stock, and Wise (1992). It incorporates aspects of the models of Berkovec and Stern (1991) and Daula and Moffitt (1991). The model is summarized briefly here. It is a simplified version of a fuller stochastic dynamic programming model in which retirement is treated as an absorbing state and annual consumption is set equal to annual income. Rust (1989) discusses the computational complexity in a more general approach.

In this model, if the worker is employed and earns income Y_t in year t, then the systematic component of utility in that year is

$$U_w(Y_t) = Y_t^{\gamma_w}. \tag{1}$$

A retired person receives monetary retirement benefits B_t and the systematic component of utility in year t is

$$U_R(B_t) = [kB_t(r)]^{\gamma}, \tag{2}$$

where the retirement benefits $B_t(r)$ depend on the year r in which the worker retired. The factor k represents a multiplicative increase of utility obtained from receiving payments B_t without having at the same time to work.

The stochastic dynamic programming model predicts retirement by comparing the expected value of retiring today to the expected value of working another year and thereby retaining the option to retire later. The decision is based on the recursive representation of the value function, W_t,

$$W_t = \max\left\{E_t[U_W(Y_t) + \varepsilon_{1t} + \beta W_{t+1}],\ E_t\sum_{\tau=t}^{S}\beta^{\tau-t}[U_R(B_\tau(t)) + \varepsilon_{2\tau}]\right\} \tag{3}$$

with

$$\begin{aligned} W_{t+1} = \max\Big\{ & E_{t+1}[U_W(Y_{t+1}) + \varepsilon_{1t+1} + \beta W_{t+2}], \\ & E_{t+1}\sum_{\tau=t+1}^{S}\beta^{\tau-t-1}[U_R(B_\tau(t+1)) + \varepsilon_{2\tau}]\Big\}, \end{aligned}$$

where β is the discount factor and S is the maximum length of life (taken to be age 95). Discounting incorporates mortality rates.

The errors ε_{1t} and ε_{2t} in (3) represent random shocks to utilities associated with income when working and when retired, respectively. The specification estimated here assumes that the errors, ε_{1t} and ε_{2t}, are independent normally distributed. Details of the estimation method are presented in Lumsdaine, Stock, and Wise (1992).

2.3.2 Data and Results

Our previous retirement model estimates, based on firm data, have used samples in which older employees were underrepresented relative to those in their 50s and early 60s. This is a consequence of high departure rates at younger ages, leaving few employees still employed at age 65; a random sample includes few older employees. Thus a possible explanation of the mismatch between actual and predicted age-65 retirement rates is the disproportionate weight given to the younger of the older employees. To address this possibility, the sample used here includes approximately equal numbers of employees at all ages, with a total of 1,007 women and 1,727 men between age 50 and 69. The balanced sample is made possible by the large number of employees in Firm 3.

The estimated parameters for the stochastic dynamic programming model are presented in table 2.5 for women and in table 2.6 for men. The discount factor β was set to 0.9 in models 1–3; it was estimated in model 4. The parameter values for the base model 1 in each table indicate little curvature in the utility function; the hypothesis that $\gamma = 1$ cannot be rejected at the 10% level for either the men or the women. The parameter k is approximately three for both men and women. Interpreted literally, these results imply that utility is approximately linear in income, with \$1 received while retired worth approximately \$3 of income while working.

Actual and predicted departure rates for this sample are compared in figure 2.4 for women and in figure 2.5 for men. The top panel in these figures shows

Table 2.5 **Parameter Estimates for Dynamic Programming Model (data set: Firm 3, women, N = 1,007)**

	Model			
Parameter	(1)	(2)	(3)	(4)
γ	1.177	1.087	1.057	1.190
	(0.219)	(0.203)	(0.213)	(0.239)
k	2.974	2.973	2.932	8.628
	(0.238)	(0.093)	(0.249)	(0.991)
k_{age65}		2.457	1.075	2.100
		(0.850)	(0.502)	(1.139)
$k_{married\&65}$			−0.947	−0.965
			(1.851)	(0.053)
β	0.9[a]	0.9[a]	0.9[a]	0.969
				(0.022)
σ	0.159	0.160	0.163	0.097
	(0.020)	(0.017)	(0.021)	(0.028)
Summary statistics				
−ln likelihood	303.59	299.43	299.07	297.08
Likelihood-ratio test versus model 1		8.32**	9.04*	13.02**

Notes: Estimation is by maximum likelihood. All monetary values are in $100,000 (1980 dollars). Significant at the following level: * = 5%; ** = 1%.
[a]Parameter values imposed.

Table 2.6 **Parameter Estimates for Dynamic Programming Model (data set: Firm 3, men, N = 1,727)**

	Model			
Parameter	(1)	(2)	(3)	(4)
γ	1.019	1.009	1.009[a]	1.009
	(0.113)	(0.105)		(0.007)
k	3.591	3.576	3.576	3.577
	(0.176)	(0.171)	(0.154)	(0.013)
k_{age65}		−0.574	−0.294	−0.574
		(0.700)	(0.402)	(0.004)
$k_{married\&65}$			−0.250	−0.400
			(0.383)	(0.005)
β	0.9[a]	0.9[a]	0.9[a]	0.903
				(0.004)
σ	0.158	0.158	0.158[a]	0.157
	(0.015)	(0.014)		(0.011)
Summary statistics				
−ln likelihood	588.53	587.75	587.84	587.83
Likelihood-ratio test versus model 1		1.56	1.38	1.40

Notes: Estimation is by maximum likelihood. All monetary values are in $100,000 (1980 dollars).
[a]Parameter values imposed.

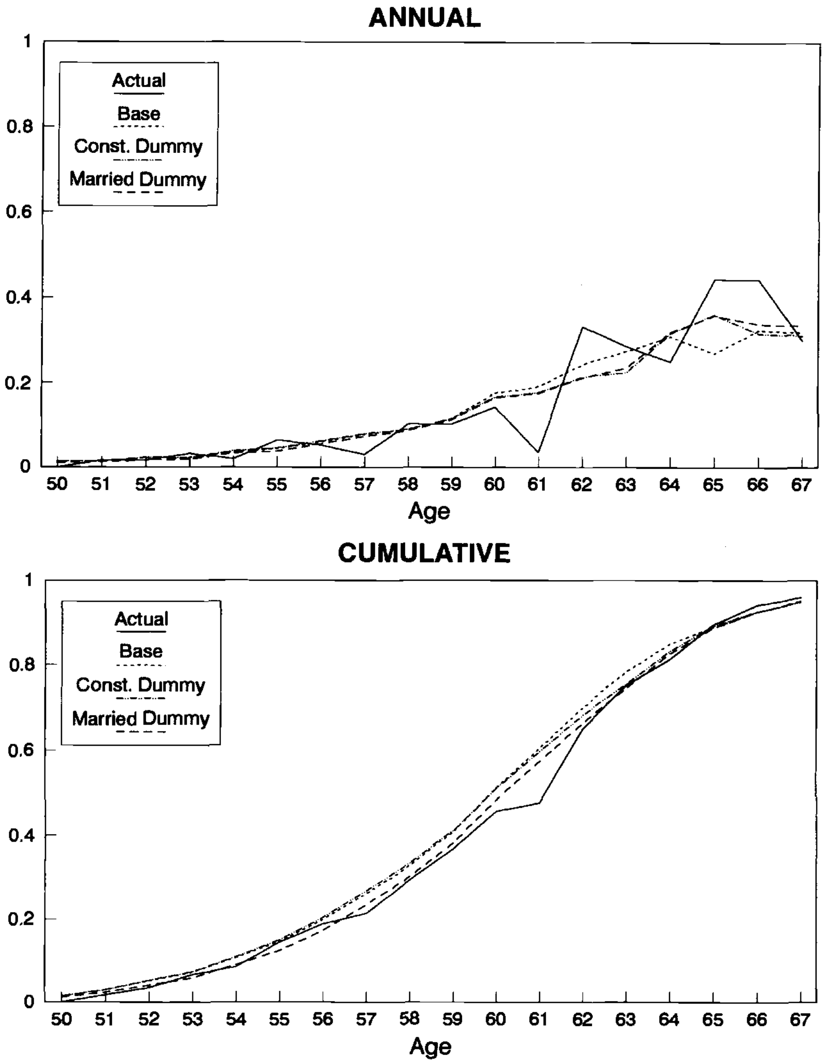

Fig. 2.4 Departure rates: women

annual departure rates, the lower panel cumulative departure. The base model 1 fits the general trends in the hazard rates reasonably well. However, at age 65 in particular, but also at age 62, the models underestimate the high retirement rates. For example, the hazard rate at age 65 predicted by the base model is 27% for women and 34% for men; the actual hazard rates are 44% for women and 61% for men.

The underprediction of age-65 retirement is typical of the results from our

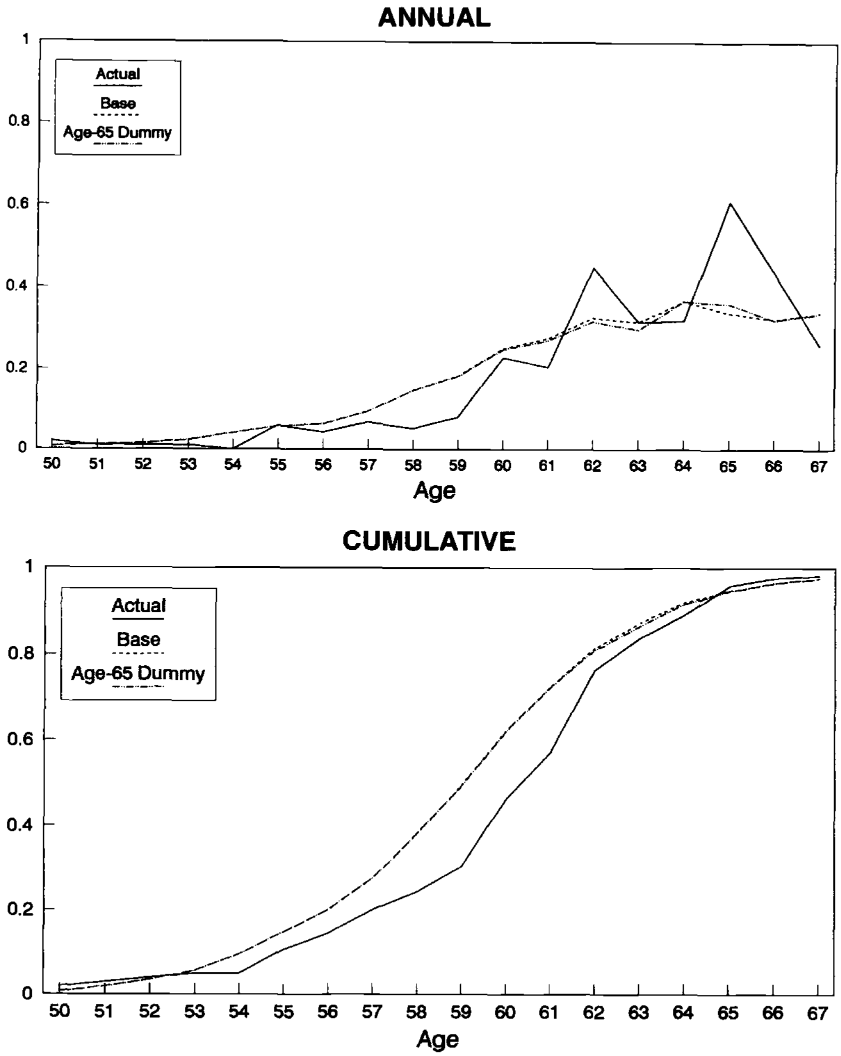

Fig. 2.5 Departure rates: men

prior analyses based on this and other firm data. The predicted departure rates are often 40 to 60% of the actual rates. For example, Lumsdaine, Stock, and Wise (1992) considered two variants of the stochastic dynamic programming model estimated in this paper (one with normal errors, as here, and one with logistic errors) and the Stock-Wise (1990b) "option value" model. All the models underpredicted age-65 retirement rates by amounts comparable to those found here. Thus the use of a random sample with relatively few older employees is not the explanation for the mismatch between actual and predicted age-65 departure rates.

To investigate age-65 retirement in more detail, we examined an additional parameterization that allows employees to behave as though they receive additional utility—in addition to financial inducements at that age—from retirement at 65. Because the analysis of section 2.2 suggests that age-65 retirement varies with marital status, we included a dummy variable for marital status as well. We parameterize k in equation 2 as a function of age and marital status,

$$k = k_0 + \exp\{k_{65} D_{65} + k_{married\&65} D_{married} D_{65}\}, \tag{4}$$

where D_{65} and D_{married} are one if the individual is 65 and married, respectively, and zero otherwise.

Results using the parameterization (4) for women are given in models 2–4 of table 2.5. The coefficient on the age-65 dummy is individually statistically significant, and the coefficients on the two dummies when estimated together are jointly significant. When the marital status dummy is included, its estimated coefficient is approximately equal to the coefficient on the age-65 dummy, and opposite in sign, so the predicted probability of retirement at age 65 is higher only for single women. This accords with the evidence from table 2.3, showing that single women aged 65 have somewhat higher retirement rates than married women.

The predicted model 2 and 3 departure rates for women are also plotted in figure 2.4. At younger ages, the age-65 utility "bonus" has little effect on the predicted rates. However, the bonus results in somewhat lower retirement rates just before age 65 (from 60 to 64) and higher rates at 65. Still, this parameterization (4) "explains" only half of the age-65 mismatch.

The results for men are given in the final three columns of table 2.6, and the hazard rates based on model 2 are plotted in figure 2.5. Unlike the results for women, none of the coefficients on the dummy variables are significant, even though the retirement rates reported in table 2.3 are significantly higher for married than for single age-65 men. We interpret this to mean that differential rates for age-65 men result from different earnings histories—and thus pension and Social Security wealth. The departure rates based on model 2 for men are very similar to those of the base model, consistent with the small numerical value of the coefficient on the age-65 dummy variable.

In summary, for neither men nor women does the base model explain the age-65 retirement spike. For single women the equation (4) parameterization appears to explain some of the difference between actual and predicted age-65 departure rates. For men this parameterization explains essentially none of the mismatch.

2.3.3 The Opportunity Cost of an Age-65 Rule of Thumb

The optimization model of the previous section is a mathematically tractable framework for approximating retirement decisions. Indeed, the model typically predicts retirement rates rather accurately. Age 65 is the exception. One explanation for the "unexplained" retirement at 65 is that the cost of choosing that age, relative to 64 or 66 for example, is small, so that there is no overriding reason to deviate from the age that some may consider "normal." In particular this may be true for employees who choose not to retire substantially before age 65. In this sense age 65 could serve as a customary focal date for retirement, which might be sustained if the opportunity cost of retiring at that focal age is small.

We explore in this section the cost of choosing to retire at 65. That is, we evaluate the opportunity cost of adopting a "retire-at-65" rule of thumb. From an economic perspective, whether such a rule of thumb is plausible depends on the opportunity cost of following the rule. To illustrate the possible magnitudes involved, we measure this opportunity cost by the lifetime utility forgone by choosing to retire at 65 rather than at some other age. For a given individual the expected lifetime utility of retiring at age r is

$$V_t(r) = \sum_{s=t}^{r-1} \beta^{s-t} U_W(Y_s) + \sum_{s=r}^{S} \beta^{s-t} U_R(B_s(r)). \tag{5}$$

The individual does not live beyond age S. In particular the lifetime utility obtained by adopting the age-65 rule of thumb is $V_t(t_{65})$. We compute the maximum lifetime utility as it is computed in the Stock-Wise (1990b) option value model, $V_t(r^*)$, where r^* is the value of r ($r = t, \ldots, t_{74}$) that maximizes $V_t(r)$.

The measure of opportunity cost tabulated here is the fraction of lifetime utility forgone, f_t, by adopting the age-65 rule of thumb relative to choosing to retire at r^*, that is,

$$f_t = \frac{V_t(r^*) - V_t(t_{65})}{V_t(t_{65})} \tag{6}$$

Because γ is insignificantly different from 1, for the computation of the distribution of equation (6) we set $\gamma = 1$. This would make the units of lifetime utility current dollars, except that dollars received while retired are weighted by the leisure parameter k. The estimates of the leisure parameters from model 1 in tables 2.5 and 2.6 for women and men, respectively, are used in the calculations.

Selected percentiles of the estimated distributions of lifetime utility forgone as a result of the age-65 rule of thumb for employees between the ages of 55 and 64 are shown in table 2.7 for employees in our balanced sample. For some employees the retirement date that maximizes the expected value of lifetime utility is 65, so for them $f_t = 0$. These are the employees with the highest predicted age-65 retirement rates. If expected lifetime utility is maximized by retiring at other than age 65, then f_t is positive. The table shows that for at least half the employees the cost of electing to retire at 65 rather than the optimal age is substantial, exceeding 20%. For example, the typical 60-year-old man would gain 23% by electing to retire at the optimal age rather than age 65;

Table 2.7 **Distribution of Fractional Utility Losses from Adopting an Age-65 Rule of Thumb**

	Percentile		
Age	10%	25%	50%
	A. Men		
55	.007	.043	.123
56	.007	.051	.143
57	.009	.055	.157
58	.009	.060	.172
59	.012	.074	.209
60	.013	.081	.232
61	.019	.108	.277
62	.016	.098	.282
63	.009	.086	.297
64	.009	.063	.200
	B. Women		
55	.018	.074	.158
56	.011	.074	.172
57	.007	.073	.188
58	.017	.096	.228
59	.014	.110	.268
60	.022	.131	.303
61	.016	.118	.342
62	.016	.131	.373
63	.025	.150	.400
64	.020	.102	.313

Notes: Entries are the percentiles of the distribution of $[V_t(r^*) - V_t(t_{65})]/V_t(t_{65})$ by age, as defined in the text. For women, the lifetime utilities were computed using $\gamma = 1$, $k = 2.974$, and $\beta = 0.9$; for men, $\gamma = 1$, $k = 3.591$, and $\beta = 0.9$.

the typical woman would gain 30%. Nonetheless, the gain for a minority of employees is small, 1 or 2% at the tenth percentile. Thus these calculations are inconsistent with the possibility that the typical employee would lose little by electing to retire at 65.

The reasoning of Akerlof and Yellen (1985) seems not to help in this instance. They point out that apparently large deviations from optimality when measured in terms of decision variables are consistent with nearly rational behavior, when the individual's objective function is rather flat in a region around the optimum. For most employees the opportunity cost of following an age-65 rule of thumb is large, and for them it would be expensive to retire at a customary but suboptimal date. However, for a minority of employees it seems not to be costly to shift retirement from an optimal date to age 65. These shifts by a relatively small fraction of workers to the customary retirement age could in principle be sufficient to explain the age-65 spike, because the absolute number of retirees at age 65 is small even though the retirement rate is large. An implication of this explanation, and in particular of the large opportunity costs for most workers, is that shifts in Social Security or pension plan provisions that made age-65 retirement less advantageous could overcome an age-65 customary retirement effect.

These computations are based only on future labor income and pension and Social Security benefits; they omit other sources of wealth like home equity and personal financial assets. Postretirement utilities should be augmented in accordance with the amount of personal wealth that the worker does not plan to bequeath. For example, the specification in equation (2) could in principle include the annuity value of personal financial assets. We do not have individual asset data for Firm 3 employees. However, incorporating other wealth into V_t would tend to decrease the proportional lifetime utility losses in table 2.7.

2.4 Conclusions

The high age-65 retirement rate for the small proportion of employees who work until that age is not explained by Medicare eligibility, based on comparison of the retirement rates of employees with and without employee-provided retiree health insurance. This conclusion is supported by the high age-65 retirement rates of Firm 3 employees, who have generous retiree health insurance. Still the age-65 retirement rate for Firm 3 employees is twice as high as the age-64 rate. Nor is the high age-65 rate explained by demographic attributes of employees. Nor is it explained by the use of data sets with small proportions of persons who are still employed at age 65.

Disproportionate age-65 retirement might be explained by an age-65 rule of thumb, which could be rationalized economically if the cost of such a rule were small. But our calculations suggest that the opportunity cost would typically be very large, even for employees who worked until 64.

Thus we are left with the hypothesis with which we started. We are inclined to attribute the unexplained high age-65 departure rates to an "age-65 retirement effect," that is, to the influence of custom or accepted practice.

References

Akerlof, George A., and Janet L. Yellen. 1985. Can Small Deviations from Rationality Make Significant Differences to Economic Equilibria? *American Economic Review* 75: 708–20.

Berkovec, J., and S. Stern. 1991. Job Exit Behavior of Older Men. *Econometrica* 59 (1): 189–210.

Blau, David M. 1994. Labor Force Dynamics of Older Men. *Econometrica* 62 (1): 117–56.

Daula, Thomas V., and Robert A. Moffitt. 1991. Estimating a Dynamic Programming Model of Army Reenlistment Behavior. In *Military Compensation and Personnel Retention: Models and Evidence,* ed. C. L. Gilroy, D. K. Horne, and D. Alton Smith.

Alexandria, VA: U.S. Army Research Institute for the Behavioral and Social Sciences.
Geweke, John, Gary A. Zarkin, and Robert B. Slonim. 1993. The Social Security Acceptance Decision. Department of Economics, University of Minnesota. Manuscript.
Gustman, Alan, and Thomas L. Steinmeier. 1986. A Structural Retirement Model. *Econometrica* 54 (3): 555–84.
———. 1994. Employer-Provided Health Insurance and Retirement Behavior. *Industrial and Labor Relations Review* 48: 86–102.
Kotlikoff, Laurence J., and Daniel E. Smith. 1983. *Pensions in the American Economy.* Chicago: University of Chicago Press.
Kotlikoff, Laurence J., and David A. Wise. 1988. Pension Backloading, Wage Taxes, and Work Disincentives. In *Tax Policy and the Economy,* vol. 2, ed. L. Summers, 161–96. Cambridge: MIT Press.
Lumsdaine, Robin L., Phyllis Mutschler, and David A. Wise. 1992. Pension Provisions, Anticipated Windows, and Retirement. Department of Economics, Princeton University. Manuscript.
Lumsdaine, Robin L., James H. Stock, and David A. Wise. 1990. Efficient Windows and Labor Force Reduction. *Journal of Public Economics* 43: 131–59.
———. 1991. Windows and Retirement. *Annales d'Économie et de Statistique* 20/21: 219–42.
———. 1992. Three Models of Retirement: Computational Complexity versus Predictive Validity. In *Topics in the Economics of Aging,* ed. David A. Wise, 19–57. Chicago: University of Chicago Press.
———. 1994. Pension Plan Provisions and Retirement: Men and Women, Medicare, and Models. In *Studies in the Economics of Aging,* ed. David A. Wise, 183–220. Chicago: University of Chicago Press.
Madrian, Brigitte C. 1993. Post-Retirement Health Insurance and the Decision to Retire. Working paper, Department of Economics, Massachusetts Institute of Technology.
Phelan, Christopher, and John P. Rust. 1991. U.S. Social Security Policy: A Dynamic Analysis of Incentives and Self-Selection. Department of Economics, University of Wisconsin. Manuscript.
———. 1993. How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets. Department of Economics, University of Wisconsin. Manuscript.
Rust, John P. 1989. A Dynamic Programming Model of Retirement Behavior. In *The Economics of Aging,* ed. David A. Wise, 359–98. 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.