Pensions, the Option Value of Work, and Retirement

James H. Stock, David A. Wise

Econometrica, Volume 58, Issue 5 (Sep., 1990), 1151-1180.

PENSIONS, THE OPTION VALUE OF WORK, AND RETIREMENT

By James H. Stock and David A. Wise[1]

The effects of firm pension plan provisions on the retirement decisions of older employees are analyzed. The empirical results are based on data from a large firm, with a typical defined benefit pension plan. The "option value" of continued work is the central feature of the analysis. Estimation relies on a retirement decision rule that is close in spirit to the dynamic programming rule but is considerably less complex than a comprehensive implementation of that rule, thus greatly facilitating the numerical analysis.

Keywords: Firm, pension, retirement, option value.

The typical firm pension plan presents very large incentives to retire from the firm at an early age, often as young as 55. Although the labor supply effects of Social Security provisions have been the subject of a great deal of analysis, much less attention has been directed to the implications of firm pension plans. Yet the retirement inducements in the provisions of firm plans are typically much greater than the incentives inherent in Social Security benefit formulae, as demonstrated by Kotlikoff and Wise (1985, 1987). Indeed, the provisions of most firm plans are at odds with the planned increase in the Social Security retirement age; private plans encourage early retirement. This paper presents a new model of retirement and uses it to estimate the effects of pension plan provisions on the departure rates of older salesmen from a large Fortune 500 firm. An important goal is to develop a model that can be used to predict the effects on retirement of potential changes in pension plan provisions. The analysis is based on longitudinal personnel records from the firm.

The option value of continued work is the central feature of the model. Pension plan provisions typically provide a large bonus if the employee works until a certain age, often the early retirement age, and then a substantial inducement to leave thereafter. Employees who retire later may do so under less advantageous conditions. If the employee retires before the early retirement age, the option of a later bonus is lost. Continuing to work preserves the option of retiring later, hence the terminology: the "option value" of work.

The provisions of firm pension plans that have motivated our work are described in the next section. The option value model is described in Section 2. Results are presented and the model fit is discussed in Section 3. Simulations of illustrative potential changes in pension plan provisions are presented in Section 4. A summary and conclusions are in Section 5.
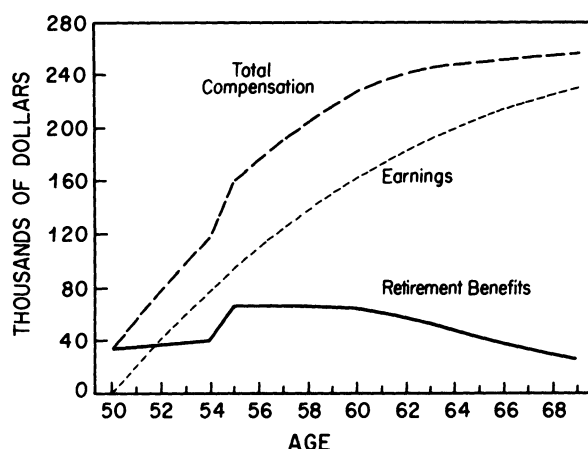
FIGURE 1.—Present discounted values of future earnings and retirement benefits, as a function of date of retirement.

1. BACKGROUND

A. Firm Pension Plan Provisions

Approximately 50 percent of American workers are covered by private pension plans. Of these, approximately 75 percent are covered by defined benefit plans. These plans promise the employee a benefit at retirement that is typically based on age, years of service, and his final salary (or an average of earnings in the last few years of employment). Within this general framework, the benefit formulas of most plans provide a large incentive to remain with the firm until some age and then a substantial incentive to leave the firm at some later age. The specific provisions of firm plans, however, vary enormously. Thus the incentives for retirement or departure from the firm vary widely among firms. The incentives of plan provisions and their variation among plans are described in detail by Kotlikoff and Wise (1985).[2] Because the incentives vary so greatly among plans, to analyze the effects of plan provisions on retirement, it is necessary to account for the precise provisions of an employee's plan. It is also critical to have information on past and current earnings in the firm. For these reasons, we rely on firm personnel records for this analysis.

The easiest way to understand the incentive effects of pension plans is to consider the relationship between age and total compensation—including wage earnings, the accrual of future pension benefits, and the accrual of future Social Security benefits.

Figure 1 shows the future wage earnings and retirement benefits of a 50 year old representative employee drawn from our data set, graphed against potential future ages of retirement. (The underlying annual data are shown in Table I.)

[2] See also Bulow (1981), Lazear (1983), Clark and McDermed (1984), Fields and Mitchell (1985), Frant and Leonard (1987).

TABLE I

EARNINGS, PENSION BENEFITS, AND SS BENEFITS
FOR A REPRESENTATIVE INDIVIDUAL

Age	Earnings Forecast	Annual Pension Benefits		Social Security
		Excl SS Offset	Incl SS Offset	
50	22317	—	2764	4533
51	22389	—	2900	4723
52	22327	—	3010	4914
53	22330	—	3221	5102
54	22271	—	3271	5288
55	22172	9522	6251	5465
56	22024	10884	7047	5640
57	21832	12465	8006	5771
58	21593	14173	9035	5798
59	21310	16271	10394	5822
60	20981	18546	11861	5842
61	20610	19739	12647	5863
62	20196	20989	13473	5885
63	19742	22309	14346	6305
64	19250	23652	15219	6538
65	18721	—	15756	6757
66	18158	—	16809	6546
67	17564	—	17847	6349
68	16943	—	18862	6167

Notes: All values are in 1980 dollars. Income forecasts were computed using the estimated income forecasting equation shown in the Appendix.

The curve labeled earnings is the present value, discounted to age 50, of future earnings; thus the slope is the discounted annual wage rate. Based on earnings forecasts (shown in Table I) this individual will have slightly declining real wage earnings over the next 15 years, with more rapidly declining earnings in his late 60's.[3] (Forecasted future earnings are based on the experience of other employees in the firm, and on the past earnings of this individual. The estimation procedure is described below.)

The retirement benefits curve shows the present value—discounted to age 50 —of expected pension plus Social Security benefits, by age of retirement.[4] The slope of this curve indicates the annual accrual of retirement benefits. The accrual of firm pension benefits is negative for this individual after age 60. The top curve shows total compensation, the sum of wage earnings, and the accrual of retirement benefits. For example, if this person were to work until age 60 and then retire, between 50 and 60 he would earn approximately $150,000 (in present value terms). Then he would be entitled to pension and SS

---

[3] The earnings figures in Table I are in constant 1980 dollars. These values are discounted at a 5 percent real interest rate to compute the present values used in Figure 1.

[4] The SS benefits are assumed to be indexed, and thus constant in 1980 dollars. Pension benefits are assumed constant in nominal terms. (The pension and SS figures in the table are explained below.) The present values plotted in Figure 1 are computed assuming a 5 percent inflation rate and a 5 percent real discount rate.

benefits with a present value at age 50 of $65,000. After age 62 or 63, total compensation from working an additional year is essentially zero. The sharp kinks in the total compensation curve are due to the discontinuous accrual of pension benefits.

The pension accrual is the result of several important provisions of the firm's pension plan. The plan "normal" retirement age is 65, the "early" retirement age is 55, and employees are vested (entitled to future benefits) after 10 years of service. "Normal" retirement benefits are determined by a standard formula of the form $B = k \cdot$ (Years of Service) $\cdot$ (Final Wage), where the final wage is an average over the last few years of employment. But the important features of the pension accrual graphed in Figure 1 are the result of three key exceptions to this standard formula. The first important provision pertains to early retirement benefits. If a person leaves this firm before age 55, early retirement benefits can be taken beginning at age 55, but the benefit will be the normal retirement benefit—that would be received at age 65—actuarially reduced (by an amount specified in the plan) to age 55. That is, the annual benefit is reduced just enough to offset the fact that the benefits will be received for 10 more years; the expected present value of benefits is the same whether they begin at 55 or at 65, after taking account of future life expectancy. If, however, the person stays in the firm until age 55, early retirement benefits can be taken immediately and the early retirement reduction factor is much less than actuarial. Thus the expected present value of future benefits is greater if the benefits are taken beginning at age 55 than if they are taken beginning at 65.

The second important provision of the plan is a Social Security offset after age 65; pension benefits are reduced depending on the person's Social Security benefits. But the offset is not applied to benefits received between 55 and 65, the normal retirement age. The magnitude of the effect of this provision can be seen in Table I. For example, if the person were to retire at age 60, his pension benefit including the SS offset would be $11,861. But the offset is not applied until age 65. Thus he would receive $18,546 per year, shown in column 2, until age 65; after 65 he would receive $11,861. If he were to retire at age 65 he would receive $15,756 per year, including the SS offset. These provisions mean that there is a large incentive to stay in the firm until 55. After 55, the incentive is reduced.

For the person represented in the graph, there is a sharp reduction in the accrual of pension benefits at age 60, due to the third important feature of the plan. If the person has 30 years of service at age 60, he is entitled to full normal retirement benefits. That is, by continuing to work he will no longer gain from fewer years of early retirement reduction, as he did before age 60. These and other plan provisions are described in detail, together with additional examples, in Kotlikoff and Wise (1987).

Finally, SS benefits are available beginning at age 62. The SS benefit shown in the last column of Table I is the benefit that this person would receive at age 62 if he were to retire before then, discounted to age 50, or the benefit that he would receive during his first year of retirement if he were to retire at age 62 or

later. The cumulative values graphed in Figure 1 account for the indexing of SS benefits to the Consumer Price Index.

B. *Prior Emphasis on Social Security Provisions*

The foregoing description suggests that the incentive effects inherent in this firm's pension plan provisions are much more important than those resulting from Social Security provisions. Indeed, this is typically the case. Yet most prior research on retirement behavior has been directed to the effects of Social Security provisions. Recent examples are Blinder, Gordon, and Wise (1980), Burkhauser (1980), Hurd and Boskin (1981), Gustman and Steinmeier (1986), Burtless and Moffitt (1984), Burtless (1986), Hausman and Wise (1985). With few exceptions (Hurd and Boskin (1981) and, to some extent, Hausman and Wise (1985)), these studies suggest only a modest effect of Social Security provisions on retirement behavior. In contrast, there has been very little work relating retirement behavior of covered workers to the retirement incentives provided by their pension plans. Exceptions are Fields and Mitchell (1982), Kotlikoff and Wise (1989), Burkhauser (1979), Hogarth (1988). The apparent reason for this lack of attention has been the absence of appropriate data.

Figure 1 suggests three key requirements for analysis of the retirement effects of pension provisions. First, the data must include the precise provisions of the individual's pension plan, together with detailed information on prior earnings. Second, the estimation method must account for sharp jumps or drops in pension accrual in future years. For example, in considering whether a person will leave the firm at age 50, it is critical to account for the large "bonus" that he will get if he remains until age 55; consideration only of total compensation at age 50 is not sufficient. Third, the estimation method needs to account for the fact that individual circumstances, such as the level of wage earnings, change over time. Such changes in turn affect future pension and Social Security accrual. The combination of firm data and the estimation method proposed here is intended to meet these requirements.

C. *Prior Estimation Methods*

To date, in addition to least squares regression, two basic approaches have been used to analyze retirement behavior. The first is the method of estimation developed to analyze the choices of individuals who face discontinuous or kinked budget constraints. The second approach is the continuous time failure rate or hazard model. Since retirement is typically a discrete outcome, but also has a time dimension (age) which not only characterizes retirement but may also affect the desire for it, it is natural to describe retirement within the context of a continuous time qualitative choice model. These two approaches are described briefly in turn.

The adaptation of nonlinear budget constraint analysis to retirement may be called the "lifetime budget constraint" approach. The central feature of this method is a lifetime budget constraint analogous to the standard labor-leisure budget constraint, but with annual hours of work replaced by years of labor force participation, and annual earnings replaced by cumulative lifetime compensation. The optimal age of retirement is determined by a utility function defined over years of work (post-retirement years of leisure) and cumulative compensation. A careful application of this approach to retirement is by Burtless (1986), who analyzed the effects of changes in Social Security benefits on retirement.[5]

While appealing in many respects, this procedure has an important drawback. It implicitly assumes that individuals know with certainty the opportunities—like wage rates—that will be available to them in the distant future. Although it is plausible to assume that an individual knows his wage rate for the purposes of estimating annual labor supply, the simple extension of this idea to construct a lifetime budget constraint is not as plausible. How much does a 50-year-old person know about his wage at 67? Concomitant with this assumption, the method makes no allowance for updating of information about future opportunities as the individual ages.

The hazard model approach as implemented to date is essentially a reduced-form technique designed to capture the effects on retirement of movements in variables such as Social Security wealth. Implementations of the hazard model have not been as "forward looking" as the nonlinear budget constraint specifications. It is natural under this specification, however, to update information as individuals age. For example, if an individual has not retired at age t, it is convenient to describe the probability of retirement by age $t + 1$ in terms of variables such as annual wage earnings and private pension accruals up to age t and in terms of these values in the period t to $t + 1$; but it is not natural to consider values of these variables in future years. Thus in Hausman and Wise (1985), for example, changes from the current period to the next, in earnings, pension wealth, and the increment to pension wealth, are allowed to affect the decision to retire in the next period, but these values several years hence are not. On the other hand, it is easy within this framework to allow a flexible specification. In particular, different forms of monetary compensation can be entered separately with no increase in computational complexity. And possibly more important for retirement, unexpected shocks, like sudden changes in earnings, enter the analysis very naturally.

2. THE OPTION VALUE MODEL

The model proposed here incorporates the advantages of both of the approaches described above. It allows updating of information, as does the

---

[5] An analogous model was used by Venti and Wise (1984) to describe the rent paid by low income families faced with discontinuous budget constraints. Earlier papers that develop these techniques are Hausman and Wise (1980) and Burtless and Hausman (1978).

traditional hazard model, but also considers potential compensation many years in the future, as does the nonlinear budget constraint approach. Antecedents of our work begin with Lazear and Moore (1988), who argue that the option value of postponing retirement is the appropriate variable to enter in a regression equation explaining retirement.[6] Our model is close in spirit to the stochastic "dynamic programming" model of Rust (1989).[7] A "dynamic programming" model of employment behavior has also been proposed by Berkovec and Stern (1988).[8] Neither Rust nor Berkovec and Stern, however, have information on private pension plan provisions, the focus of our analysis. The retirement decision rule that we propose as an approximation to individual behavior is much simpler than the dynamic programming rule. A concomitant of this assumption is also much simpler econometric implementation than the burdensome calculations imposed by the dynamic programming rule. These simplifications reduce the computational requirements substantially while retaining the key forward-looking features of the dynamic programming approach. Of course both of these models are theoretical abstractions. The important consideration is which decision rule is the better approximation to the calculations that govern actual individual behavior. The answer to this question will have to await further analysis. We show that the rule we assume predicts individual choices well, but we have not compared the predictive validity of the decision rule assumed here with the dynamic programming rule. The primary distinction between the two approaches is explained below.

The key ideas of the model can be summarized briefly. It is intended to capture an important empirical regularity, the irreversibility of the retirement

---

[6] Indeed it was their work and analysis of military retirement rates by Phillips and Wise (1987) that motivated us to pursue this approach.

[7] Rust's (1989) model poses substantially greater numerical complexity than ours and has not yet been estimated for retirement. In principle, he observes not only the individual's retirement age, but subsequent consumption decisions as well. Thus his model allows the individual to optimize over age of retirement and future consumption jointly. The individual's decision is modeled as the solution to a stochastic dynamic programming problem. As in our case, the individual's expectations are conditioned on current known variables such as income. The idea is to recover the parameters of a utility function specified in terms of these choice variables. In practice, though, he uses income to describe consumption (Rust (1988)), with a value function similar to ours, specified in terms of income. To simplify the solution to the dynamic programming problem in his model, he assumes that random unobserved individual components are independent over time, whereas we allow such terms in our model (representing differences among individuals in health status, desire for leisure, and the like) to be correlated. In short, Rust has described a solution to a more complicated choice than ours, but with uncorrelated errors, whereas ours is a solution to a less complex problem, but with correlated errors.

[8] Berkovec and Stern consider transitions among three employment states over time. To simplify the solution to their optimization problem, they assume that disturbance terms are uncorrelated over time, except for additive individual and job specific random effects. Their analysis is in terms of individual attributes like education, race, health status, and age. Government benefits like Social Security are not explicitly modeled, whereas these benefits, as well as firm pension benefits play the central role in our analysis. We estimate a discount, or weighting factor, whereas they obtain estimates of other parameters conditional on an assumed discount rate. Age itself is used explicitly to estimate retirement. As will be emphasized below, age is not a direct determinant of retirement in our model. This has important implications if the model is to be used to predict the effect on retirement of changes in firm pension plan or Social Security provisions.

decision. Although it is not uncommon to work—at least part-time—after "retirement," it is rare to return to the firm from which one has retired. The model focuses on the opportunity cost of retiring or, equivalently, on the value of retaining the option to retire at a later date. It has two key aspects. The first is that a person will continue to work at any age if the expected present value (in utility units) of continuing work is greater than the expected present value of immediate retirement. In effect, the person compares the best of expected future possibilities with the value of retiring now. The second is that the individual reevaluates this retirement decision as more information about future earnings—and thus future retirement benefits—becomes available with age. For example, a decline in the wage between ages 56 and 57 will cause the individual to reassess future wage earnings, and thus future pension benefits and Social Security accrual as well. Thus retirement may seem more advantageous upon reaching 57 than it was expected to be at age 56. Retirement occurs when the value of continuing work falls below the value of retiring.

Because the model is somewhat complex in its details, it is useful to know whether a simpler model could capture the important features of this one; in particular, whether a model that is easier to implement could predict retirement outcomes as well as the more complex model. It is shown in the working paper version of this paper (Appendix A of Stock and Wise (1988a)) that a simplified version of the model developed here has an almost direct hazard model counterpart. Indeed, as is shown in that paper, the proportional hazard model can be interpreted in terms of utility maximization, contrary to a common misperception.[9] Hazard models are very simple to estimate. Unfortunately, the hazard model is obtained only after imposing strong restrictions on several important features of the option value model.

In addition to the general rationale for the option value model, the precise specification as set forth in this paper is guided by two considerations: first, by the features of the firm data that are used in estimation, and second, by the primary goal of the model, to predict the result of changes in firm pension plan provisions. In particular, some individual attributes that might be expected to affect retirement behavior—such as assets other than pension and Social Security wealth—are unknown to us.[10] Our retirement decision function is therefore based on wage earnings and retirement benefit income. In the

---

[9] For a similar argument, see Rust (1987). In that paper, the hazard of a bus engine being replaced, as a function of bus mileage, is derived from an underlying dynamic programming problem.

[10] While assets other than retirement annuity wealth (the present value of firm pension and Social Security benefits) should in principle affect retirement, prior analysis shows that their effect is small relative to Social Security wealth, as demonstrated in Hausman and Wise (1985), for example. In addition, prior work has shown that a large majority of the elderly have very little wealth other than housing and firm pension and Social Security annuities (e.g., Diamond and Hausman (1984), Hurd and Shoven (1983), Hurd and Wise (1989)) and that housing wealth is typically not consumed as the elderly age (Merrill (1984), Venti and Wise (1988, 1989a, 1989b), Feinstein and McFadden (1989)). Indeed, liquid wealth is reduced very little as the elderly age (e.g., Venti and Wise (1989a), Bernheim (1987)). Thus there is substantial evidence that the typical retired person is living largely from Social Security and pension benefits.

absence of additional information, we propose an error structure that is intended to capture the effects of persistent unobserved individual attributes.

A. *The Model*

Consider an individual at the beginning of year t, who has not yet retired. Looking ahead, he will receive wage income Y_s in year s as long as he continues to work; if he is retired in year s, he will receive real retirement benefits B_s. (We adopt the convention that if s is the first calendar year during which the person has no wage earnings, he is assumed to have left the firm during the previous year, at the age that he was on January 1 of year s.) Let r denote the first full year of the individual's retirement (that is, the first year in which the individual has no wage earnings). As described above, these benefits will depend on the person's age and years of service at retirement, and on his earnings history; thus we typically write the benefits as $B_s(r)$.

To develop a decision function relative to retirement, suppose that the individual indirectly derives utility $U_w(Y_s)$ from the real income earned while working and utility $U_r(B_s(r))$ from the pension benefits received while retired. Suppose that in deciding whether to retire the individual weights future income (or utility) by the discount factor β, and that with probability one he will die by year S. If he retires at age r, the weighted, or discounted, value received over the remainder of his life is:

$$(2.1) \qquad V_t(r) = \sum_{s=t}^{r-1} \beta^{s-t} U_w(Y_s) + \sum_{s=r}^{S} \beta^{s-t} U_r(B_s(r)).$$

Thus the value function $V_t(r)$ depends on future earnings and retirement benefits, which in turn depend on the age r at which he retires.

The individual must choose either to work during year t, so that $r > t$, or to retire, so that $r = t$. We assume that he makes the decision by comparing the expected value he would receive were he to retire now, at $r = t$, with the greatest of the expected values from possible retirement dates $r > t$ in the future. Let $E_t(\cdot)$ denote the individual's expectation about future circumstances, based on information available to him at the beginning of year t. (With this convention, real income y_t earned during year t is not known at the beginning of year t.)

The expected gain, in year t, from postponing retirement to age r is then given by

$$(2.2) \qquad G_t(r) = E_t V_t(r) - E_t V_t(t).$$

In the firm that provided our data, retirement is mandatory at age 70. Thus we assume that the individual considers potential retirement dates between $t + 1$ and the year of his seventieth birthday, t_{70}. Let $r^* > t$ be the future retirement year yielding the highest expected value, that is,

$$(2.3) \qquad r^* \text{ solves} \quad \max_{r \in \{t+1, t+2, \ldots, t_{70}\}} E_t V_t(r).$$

The individual retires if there is no expected gain from continued work, that is, if $G_t(r^*) = E_t V_t(r^*) - E_t V_t(t) \leqslant 0$. Otherwise he postpones retirement. In short, we adopt the following decision rule: the employee

(2.4)   continues to work at t if: $G_t(r^*) = E_t V_t(r^*) - E_t V_t(t) > 0$.

We assume that the utility derived indirectly from annual income has a constant relative risk aversion form, with additive individual disturbance terms distributed independently of income and age. Specifically,

(2.5a) $U_w(Y_s) = Y_s^\gamma + w_s,$

(2.5b) $U_r(B_s) = (kB_s(r))^\gamma + \xi_s,$

where w_s and ξ_s are individual-specific random effects that vary over time. They are intended to capture several unobserved determinants of retirement. For example, w_s and ξ_s could reflect individual preferences for work versus leisure. Or, they could reflect evolving health status. They could reflect differences among individuals in unobserved wealth and other variables that may affect retirement decisions. Given the nature of our data, they are also likely to reflect the fact that for some persons the alternative to continued work in the firm is not retirement, but another job, an issue that we return to below. We presume that, for a given individual, there should be considerable persistence in these random effects over time. For example, a disability that affects the burden of working, and thus corresponds to a negative w_s, is likely to yield a negative w_{s+1} as well. Such persistence is captured by assuming that the random individual effects follow a Markovian or first order autoregressive process:

(2.6a) $w_s = \rho w_{s-1} + \varepsilon_{ws}, \qquad E_{s-1}(\varepsilon_{ws}) = 0,$

(2.6b) $\xi_s = \rho \xi_{s-1} + \varepsilon_{\xi s}, \qquad E_{s-1}(\varepsilon_{\xi s}) = 0,$

for $s = t + 1, \ldots, S$. We give particular attention in the empirical work to the case with $\rho = 1$, with the individual effects evolving according to a random walk. We adopt the convention that at time s the individual knows w_s and ξ_s, but not their values at $s + 1$ and subsequent ages; future forecasts of w and ξ are based on (2.6).

With the parameterization (2.5), $G_t(r)$ in (2.2) becomes

$$
\begin{aligned}
(2.7) \quad G_t(r) &= E_t \sum_{s=t}^{r-1} \beta^{s-t} \left[(Y_s^\gamma) + w_s \right] + E_t \sum_{s=r}^{S} \beta^{s-t} \left[(kB_s(r))^\gamma + \xi_s \right] \\
&\quad - E_t \sum_{s=t}^{S} \beta^{s-t} \left[(kB_s(t))^\gamma + \xi_s \right] \\
&= E_t \sum_{s=t}^{r-1} \beta^{s-t} (Y_s^\gamma) + E_t \sum_{s=r}^{S} \beta^{s-t} (kB_s(r))^\gamma \\
&\quad - E_t \sum_{s=t}^{S} \beta^{s-t} (kB_s(t))^\gamma + E_t \sum_{s=t}^{r-1} \beta^{s-t} (w_s - \xi_s) \\
&= g_t(r) + \phi_t(r)
\end{aligned}
$$

where $g_t(r)$ and $\phi_t(r)$ distinguish the terms in $G_t(r)$ containing the random effects, w and ξ, from the other terms.

If whether the person is alive in future years is statistically independent of his earnings stream and the individual effects w_s and ξ_s, then $g_t(r)$ and $\phi_t(r)$ become

$$(2.8) \qquad g_t(r) = \sum_{s=t}^{r-1} \beta^{s-t}\pi(s|t)E_t(Y_s^\gamma) + \sum_{s=r}^{S} \beta^{s-t}\pi(s|t)\left[E_t(kB_s(r))^\gamma\right]$$
$$- \sum_{s=t}^{S} \beta^{s-t}\pi(s|t)\left[E_t(kB_s(t))^\gamma\right]$$

and

$$(2.9) \qquad \phi_t(r) = \sum_{s=t}^{r-1} \beta^{s-t}\pi(s|t)E_t(w_s - \xi_s),$$

where $\pi(s|t)$ denotes the probability that the person will be alive in year s, given that he is alive in year t. Given the Markov assumption (2.6), $\phi_t(r)$ can be written as

$$(2.10) \qquad \phi_t(r) = \sum_{s=t}^{r-1} \beta^{s-t}\pi(s|t)\rho^{s-t}(w_t - \xi_t)$$
$$= K_t(r)\nu_t,$$

where $K_t(r) = \sum_{s=t}^{r-1}(\beta\rho)^{s-t}\pi(s|t)$ and $\nu_t = w_t - \xi_t$. The simplification results from the fact that at time t the expected value of $\nu_s = w_s - \xi_s$ is $\rho^{s-t}\nu_t$, for all future years s. Thus the individual random component $\phi_t(r)$ depends only on the random effect at time t (together with β and ρ). The term $K_t(r)$ cumulates the deflators that yield the present value in year t of the future expected values of the random components of utility. The further r is in the future, the larger is $K_t(r)$. That is, the more distant the potential retirement age, the greater the uncertainty about it, yielding a heteroskedastic disturbance term. This heteroskedastic property is apparently an important determinant of the model's ability to predict departure rates accurately for both younger and older employees.

Combining (2.7)–(2.10), $G_t(r)$ may be written simply as

$$(2.11) \qquad G_t(r) = g_t(r) + K_t(r)\nu_t.$$

B. *The Probability of Retiring*

We consider the probability that an employee in our sample retires during the period in which firm employees are tracked. The analysis is conditional on being in the sample, or, equivalently, not having left the firm between the date of hire and the first year of the data. As described below, it is impractical in the context of our model to compute the probability that an employee remains in the firm until this date. Thus, our analysis pertains to the effect of the pension plan

provisions on the departure decision of older employees who are in the firm at t. We consider first the probability of retirement during the first year of our data, i.e., using cross-sectional data for a single year. Retirement over several consecutive years is then considered. Finally, we return to discussion of the conditional aspect of our analysis.

1. *Retirement probabilities for a single year:* The year in which an individual retires is a random variable, R. Let $\Pr[R = t]$ be the probability that an individual retires in year t. From (2.4), an employee will retire in year t if $G_t(r) \leqslant 0$ for all $r \in \{t + 1, \ldots, T\}$. Thus:

$$
\begin{aligned}
(2.12) \quad \Pr[R = t] &= \Pr\big[G_t(r) \leqslant 0 \; \forall \; r \in \{t + 1, \ldots, T\}\big] \\
&= \Pr\big[g_t(r) + K_t(r)\nu_t \leqslant 0 \; \forall \; r \in \{t + 1, \ldots, T\}\big] \\
&= \Pr\big[g_t(r)/K_t(r) \leqslant -\nu_t \; \forall \; r \in \{t + 1, \ldots, T\}\big].
\end{aligned}
$$

Alternatively, the final expression in (2.12) can be written:

$$
(2.13) \quad \Pr[R = t] = \Pr\big[g_t(r_t^\dagger)/K_t(r_t^\dagger) \leqslant -\nu_t\big]
$$

where $r_t^\dagger$ is the value of r that solves

$$
(2.14) \quad \max_{r \in \{t+1, \ldots, T\}} g_t(r)/K_t(r).
$$

Because the individual either retires at t or he does not, $\Pr[R > t] = 1 - \Pr[R = t]$.

2. *Retirement probabilities for multiple years:* The data set contains data on individual retirement decisions for several consecutive years. To analyze these data requires computing the probability that the individual retires in year τ. In general, suppose that the retirement status is observed for years $t, \ldots, T$. An individual retires in year $\tau \in \{t, \ldots, T\}$ if there is no earlier age when he considers it optimal to retire, and if it is optimal to retire in year τ based on equation (2.4). If it had not been optimal to retire in year t, there would have been, at time t, at least one future r with $G_t(r) > 0$. This would occur if and only if it were true for the $r^\dagger$ that maximized $g_t(r)/K_t(r)$, evaluated at year t. That is, it requires that $g_t(r_t^\dagger)/K_t(r_t^\dagger) > -\nu_t$. The same would have to be true for every year t through year $\tau - 1$. In year τ, however, retirement is optimal, so that $g_\tau(r_\tau^\dagger)/K_\tau(r_\tau^\dagger) \leqslant -\nu_\tau$. Thus

$$
\begin{aligned}
(2.15) \quad \Pr[R = \tau] = \Pr\big[&g_t(r_t^\dagger)/K_t(r_t^\dagger) > -\nu_t, \ldots, \\
&g_{\tau-1}(r_{\tau-1}^\dagger)/K_{\tau-1}(r_{\tau-1}^\dagger) > -\nu_{\tau-1}, \\
&g_\tau(r_\tau^\dagger)/K_\tau(r_\tau^\dagger) \leqslant -\nu_\tau\big].
\end{aligned}
$$

Equation (2.15) can be used to compute the probability that $R = \tau$ for $\tau = t, \ldots, T$. The remaining possible outcome is that the individual does not retire during the period of the data. The probability of this outcome is

$$
\begin{aligned}
(2.16) \quad \Pr[R > T] = \Pr\big[&g_t(r_t^\dagger)/K_t(r_t^\dagger) > -\nu_t, \ldots, \\
&g_{T-1}(r_{T-1}^\dagger)/K_{T-1}(r_{T-1}^\dagger) > -\nu_{T-1}, \\
&g_T(r_T^\dagger)/K_T(r_T^\dagger) > -\nu_T\big].
\end{aligned}
$$

The retirement model thus reduces to a multinomial discrete choice problem, with dependent error terms ν_s. Thus far, the only assumption about the individual effects is that they are Markov. Empirical implementation, however, requires additional distributional assumptions. We assume that ν_s follows a Gaussian Markov process, with

$$(2.17) \quad \nu_s = \rho \nu_{s-1} + \varepsilon_s, \quad \varepsilon_s \text{ i.i.d. } N(0, \sigma_\varepsilon^2),$$

where the initial value, ν_t is i.i.d. $N(0, \sigma_\nu^2)$ and is independent of ε_s, $s = t+1, \ldots, S$. The covariance between ν_τ and $\nu_{\tau+1}$ is $\rho \, \text{var}(\nu_\tau)$, and the variance of ν_τ for $\tau \geqslant t$ is $\rho^{2(\tau-t)}\sigma_\nu^2 + (\sum_{j=0}^{\tau-t-1} \rho^{2j})\sigma_\varepsilon^2$. In the random walk case, with $\rho = 1$, the covariance between ν_τ and $\nu_{\tau+1}$ is $\text{var}(\nu_\tau)$, and the variance of ν_τ for $\tau \geqslant t$ is $\sigma_\nu^2 + (\tau - t)\sigma_\varepsilon^2$.

In this model, there are two equivalent ways to see that uncertainty about the future is reduced as the planning horizon is shortened, presumably as the person approaches typical retirement ages. First, there are fewer future random components of utility to cumulate in the $K_t(r)$ term (see equation (2.10)). Second, the uncertainty about the value of future random effects is reduced—the Markov assumption yields decreasing $\text{var}(\nu_\tau)$ as the planning horizon is shortened. In particular, in a given calendar year, the uncertainty about the retirement decisions of younger persons is greater than the uncertainty about older employees. This property plays a key role in providing the flexibility that allows the model to fit the departure behavior of younger as well as older employees.

In summary: conditional on $\{g_s(r_s^\dagger)/K_s(r_s^\dagger)\}$, $s = t, \ldots, \tau$, the probability that year τ is the first year of retirement is given by (2.15), while (2.16) gives the probability that the person does not retire during the years $t, \ldots, T$. These probabilities are evaluated by computing the appropriate integrals over a multivariate normal density, where the error term follows the Markov process (2.17). The unknown parameters of the model are γ, k, β, and the variance parameters σ_ν^2, σ_ε^2, and ρ.

3. *Retirement conditional on being in the sample:* Strictly speaking, the probabilities $\Pr[R = t]$ in (2.13) and $\Pr[R = t]$ in (2.15) are appropriate only if the individual first *considers* retirement in year t, so that remaining in the firm until year t is unrelated to retirement decisions after that time. If this is not true, the appropriate cross-sectional probability is the conditional probability, $\Pr[R = t|R > t - 1] = \Pr[R = t]/\Pr[R > t - 1]$. If t_0 is the year in which the employee first considered the retirement decision—possibly the first year of employment—then

$$(2.18) \quad \Pr[R = t|R > t - 1]$$

$$= \{\Pr[g_{t_0}(r_{t_0}^\dagger)/K_{t_0}(r_{t_0}^\dagger) > -\nu_{t_0}, \ldots, g_t(r_t^\dagger)/K_t(r_t^\dagger) \leqslant -\nu_t]\}/$$

$$\{\Pr[g_{t_0}(r_{t_0}^\dagger)/K_{t_0}(r_{t_0}^\dagger) > -\nu_{t_0}, \ldots, g_{t-1}(r_{t-1}^\dagger)/K_{t-1}(r_{t-1}^\dagger) > -\nu_{t-1}]\}.$$

Thus (2.18) is the conditional probability appropriate to adjust for potential self-selection effects among employees in the sample at t.

Implementation of (2.18) (and its counterpart for the multi-year probabilities) poses three major practical difficulties in the context of our model. First, it requires making a judgment about t_0, the first year in which retirement is contemplated. It could, for example, be age 50. It could be the year of entering the firm. Or, the year of vesting. The structure of our model does not permit its being treated as a simple "fixed-effect" for each individual. Second, using (2.18) for distant t_0 is likely to add to specification bias; departure of employees under 50 is more likely to be for another job, not to retire. Information on this transition is not available to us. Thus no "correct" treatment is evident. Third, given t_0, evaluation of (2.18) would involve computing a $(t - t_0 + 1)$-dimensional integral with the dependent errors $\{v_s\}$, increasing substantially the computational burden of this approach. We therefore use (2.13), (2.15), and (2.16), and refer to the results as quasi-maximum likelihood estimates.

C. Evaluation of $g_t(r_t^\dagger)/K_t(r_t^\dagger)$

To determine $g_t(r_t^\dagger)/K_t(r_t^\dagger)$ requires evaluation of the expectations $E_t(Y_s^\gamma)$ and $E_t(kB_s(r_t^\dagger))^\gamma$ for $s \geq t$. In the empirical work, the conditional expectation of the first of these terms is approximated by the conditional expectation of its second order Taylor series expansion around the mean of a stream of earnings forecasts computed for each individual.[11] The pension and Social Security benefits depend on the entire earnings stream of the individual through his last year of work. The expectation $E_t(kB_s(r))^\gamma$ was approximated by $(k\hat{B}_s(r))^\gamma$, where $\hat{B}_s(r)$ is the pension benefit calculated using the mean earnings forecasts for the individual through year $r - 1$, based on observed earnings through year $t - 1$.[12]

The income forecasts for each individual were generated by a second order autoregression. The autoregression was estimated using the individual earnings histories of all salesmen employed at least three years, with earnings converted to 1980 dollars using the Consumer Price Index. The parameters of the forecasting model depend on age, A_t, years of service, S_t, and an interaction term, with

$$(2.19) \quad \Delta \ln Y_t = \delta_0(A_t, S_t) + \delta_1(A_t, S_t) \, \Delta \ln Y_{t-1} + \delta_2(A_t, S_t) \, \Delta \ln Y_{t-2} + e_t.$$

[11] $E_t(Y_s^\gamma) \doteq \{1 + (1/2)\gamma(\gamma - 1)E_t[(Y_s - E_t Y_s)/E_t Y_s]^2\}(E_t Y_s)^\gamma$. This term is evaluated assuming that $E_t[(Y_s - E_t Y_s)/E_t Y_s]^2 = (s - t)\mathrm{var}(e_t)$, where $\mathrm{var}(e_t) = \mathrm{SEE}^2$ from the $\Delta \ln Y_t$ forecasting equation given in the Appendix.

[12] In principle, the expectation could be evaluated using Monte Carlo methods to determine the conditional distribution of $B_s(r)$, which could be used to evaluate $E_t(kB_s(r_t))^\gamma$. The Monte Carlo procedure would entail computation of the benefits that an individual would receive as a nonlinear function of the realized income stream, where the future part of the income stream is drawn from an estimated conditional distribution of future income streams. However, the pension and Social Security calculations are quite cumbersome. Beyond its substantial computational advantages, a justification for the approximation that we use is that the benefits calculations involve the entire earnings history of the individual. Thus the unknown elements in these calculations are small, at least for values of r in the near future.

The estimated equation exhibits regression toward the mean; $\delta_1 + \delta_2 < 0$ for typical values of A_t and S_t in the sample. The estimated parameters of equation (2.19) are shown in the Appendix.

D. *Comparison with the Dynamic Programming Decision Rule*

Because the two decision rules are closely linked in spirit, to guide future research the key distinction between the retirement rule we assume and the dynamic programming rule is explained by means of a simple example. With reference to equations (2.5) through (2.9) above, assume that $\gamma = k = \beta = \pi(s|t)$ = 1. Assume also that the individual's time horizon ends at age 70—the firm mandatory retirement age, and assumed in the example to be the age of death. No pension benefits are received if a person works until age 70. In addition, assume that decision makers are either 69 or 68 years old. In the latter case, the time horizon is two years; in the first case the time horizon is only one year. At age 69, a person either continues working or retires. At age 68, a person faces three potential future choices: (i) retire at 68 and thus be retired at age 69 as well, (ii) continue to work until age 69 and then retire at age 69 ($t + 1$), or (iii) work at age 68 and at age 69 (and retire at age 70—$t + 2$). For a person who is still working at age 69, the two decision rules coincide. An employee continues to work at age 69 if:

$$(2.20) \quad E_t V_t(t + 1) > E_t V_t(t), \quad \text{equivalent to}$$

$$E_t(Y_t + w_t) > E_t(B_t(t) + \xi_t),$$

where $V_t(t + 1)$ in this case indicates the utility of retiring at age 70.

At 68, the two decision rules differ. According to the "option value" rule adopted in this paper, a person continues to work at age 68 if:

$$(2.21) \quad \text{Max}[E_t V_t(t + 1), \quad E_t V_t(t + 2)] > E_t V_t(t), \quad \text{equivalent to}$$

$$E_t(Y_t + w_t) + \text{Max}[E_t(B_{t+1}(t + 1) + \xi_{t+1}), \quad E_t(Y_{t+1} + w_{t+1})]$$
$$> [E_t(B_t(t) + \xi_t) + E_t(B_{t+1}(t) + \xi_{t+1})],$$

where E_t indicates the expectation at age 68. Based on the dynamic programming rule, a person continues to work at age 68 if:

$$(2.22) \quad E_t(Y_t + \omega_t) + E_t \max[V_{t+1}(t + 1), \ V_{t+1}(t + 2)] > E_t V_t(t),$$

equivalent to

$$E_t(Y_t + \omega_t) + E_t \max[(B_{t+1}(t + 1) + \xi_{t+1}), \ (Y_{t+1} + \omega_{t+1})]$$
$$> [E_t(B_t(t) + \xi_t) + E_t(B_{t+1}(t) + \xi_{t+1})].$$

The decision rule that we use considers the maximum of the expected values of the two options at age 69 (plus first period earnings); the second rule considers instead the expected value of the maximum of these two options. The expected

value of the maximum of two random variables is greater than the maximum of their respective expected values. Thus to the extent that the maximum of the expected values is large relative to the expected value of the maximum, our rule yields a lower value of postponing retirement, for a given set of parameters. As emphasized above, however, the relevant issue is which rule is used by individuals. They also may undervalue the option of postponing retirement. In any case, to evaluate the maximum of two random variables requires integration, one dimensional in the two period example. The dimension of integration increases by one with each additional time period in the planning horizon, as many as 20 in our analysis. Thus the dynamic programming rule introduces substantial numerical complexity, which the simpler rule does not entail. The decision rule based on the maximum of expected values, together with the random walk (or, more generally, the Markov) assumption, yields the convenient separability property shown in equations (2.7) through (2.10). Such separability is not possible under the dynamic programming rule. An alternative approach is to approximate the dynamic programming rule by simplifying the covariance structure—to reduce the dimension of integration. This is the approach taken by Rust (1989) and Berkovec and Stern (1988), for example.

The difference in the probability of retirement under the two rules depends on the variance of the random components; that is, how much new information is likely to differ from current information. If all future values are known with certainty, the two rules are equivalent. If future Y and B are known with certainty and if the difference between $Y_t + \text{Max}[B_{t+1}(t+1), Y_{t+1}]$ and $B_t(t) + B_{t+1}(t)$ is large relative to the innovations in w and ξ—the shifts in preferences are small relative to the known gap between retirement benefits and earnings—then the two rules also yield similar decisions. The two rules differ to the extent that the variance of the unknown component is large relative to the predictable differences in the values of Y and B across ages.

3. RESULTS

The option value model was estimated using a sample of 1500 salesmen 50 years of age or older on January 1, 1980, selected at random from the firm data. All persons in the sample are men performing similar jobs. To facilitate the forecasting of earnings, the sample was restricted to persons who had at least three years of service before 1980, the first year of our retirement analysis.

Initial estimates were obtained based only on retirement decisions in one year, 1980—whether 1981 was the first full year of retirement, by our dating convention. Expected pension benefits are based on the provisions of the firm plan. Social security benefits were computed according to the provisions in the Social Security Administration (1982), based on individual wage histories at the firm.[13] Estimates based on the 1980 retirement decisions are reported first, followed by estimates based on three consecutive years.

---

[13] To obtain pension and Social Security forecasts for persons who joined the firm before 1969—the first year of our data—backward predictions based on a time-reversed version of the estimated earnings equation in the Appendix were used to estimate earnings before 1969.

TABLE II

PARAMETER ESTIMATES BASED ON RETIREMENT DECISIONS
IN ONE YEAR, 1980[a]

A. *Models Without Earnings and Retirement Benefit Terms*:

				$\mathscr{L}$
1. Constant Only				-579.58
2. Age Dummy Variables				-500.12

B. *Estimated Models, with* $\rho \equiv 1$:

	γ	k	β	$\sigma\ (\times 10^5)$	$\mathscr{L}$
3.	1^b	1^b	0.90^b	.179	-544.26
				(.007)	
4.	1^b	1.76	0.734	.136	-519.77
		(.13)	(.073)	(.010)	
5.	0.584	1.18	0.90^b	.119	-507.68
	(.073)	(.25)		(.015)	
6.	0.632	1.25	0.781	.099	-506.86
	(.088)	(.28)	(.121)	(.018)	

[a] The sample size is 1500. A person was counted as having retired if he had no earnings in 1981; that is, he was retired by January 1, 1981. $\mathscr{L}$ is the log likelihood value.
[b] Parameter value imposed.

A. *One Year*

1. *Parameter Estimates*

Quasi-maximum likelihood parameter estimates are shown in Table II.[14] Estimated parameters of several variants of the option value model are shown in the second panel of the table. Estimates in this table were obtained under the assumption that the random individual effects follow a random walk; ρ is set to 1. (Unrestricted estimates are reported in the next section.) The first panel reports estimates based on the assumption that all employees have the same constant probability of retiring (model 1), or that all persons of the same age have the same retirement probability (model 2).

Estimates of the parameters of the option value model are shown in the last row of the table, specification 6. The estimate of γ is 0.63, suggesting that, in deciding whether to retire, individual valuation of income is nonlinear in future earnings. Interpreted as a risk aversion parameter, it means that the certainty equivalent of \$10,000 with probability .5 and \$20,000 with probability .5 is \$14,685, suggesting that these employees are essentially risk neutral. Logarithmic utility is clearly rejected by the data; the Wald t statistic testing $\gamma = 0$ is 7.18.

The estimated value of k, 1.25, indicates that, in deciding whether to retire, a dollar of income while retired is given more weight than a dollar of income while working. The ratio of the utility of retirement to the utility of employment is $[1.25(B/Y)]^{.632}$, which is 1 when $B/Y = 0.80$. That is, a person would exchange a dollar with work for 80 cents not accompanied by work. The weight

[14] Two algorithms were used to maximize the Gaussian likelihood: a quadratic search method, based on the Davidon-Fletcher-Powell algorithm, and a modified simulated annealing algorithm. The latter algorithm is a random search routine that was used to assure that a global, rather than a local, maximum was found.

given to current versus future income in the retirement decision is indicated by β, estimated to be .781. Each of these estimates seems quite plausible to us.

Whether the parameters are considered precisely measured depends on alternative hypothesized values. For example, these results alone do not reject the possibility of $k = 1$, which would suggest that wage earnings and retirement benefits are given equal weight in the retirement decision. Nor do the data strongly reject the possibility of $\beta = 1$, which would suggest that future real payments are given the same weight as current payments. Although a narrow interpretation of the parameters of the model would treat β as a general measure of individuals' pure rate of time preference, independent from the decision to which it applies, it is probably more realistic to think of it as a weight specific to the retirement decision. Under either interpretation, a priori judgments about its value surely vary widely. It cannot be observed and can typically be estimated only indirectly. It is, however, estimated directly in the option value model, although the estimated value is undoubtedly sensitive to the model specification.

Estimated parameters of simplified versions of the model are shown as specifications 3 through 5. A comparison of the likelihood values of these specifications shows that estimation of the model parameters improves the model fit substantially. Thus there is substantial information in the option value measures. Indeed, the full option value model (specification 6) fits the data almost as well as the specification based on a full set of age dummy variables, as reported in model 2. With age dummy variables, the estimated average departure rate for each age is made to match the actual rate. The option value specification does not impose such a match; age does not enter the specification directly.

2. The Model Fit

The model fit is demonstrated by comparing actual and predicted retirement rates, shown in Table III and in Figure 2. Both the predicted annual hazard rates and predicted cumulative retirement rates appear to be quite close to the actual values. Figure 2 shows the sample rates, together with a 95 percent confidence band (around the sample rates) relative to the actual population rates, and the predicted rates. All of the estimated annual hazard rates, except those for 62 and 65 fall well within the 95 percent confidence band. Of employees in the firm at age 50, only 5 percent would still be in the firm at age 64, according to the model estimates. Of the small number of 64 year olds in the firm, about 58 percent leave at 65; we predict that only 29 percent leave. This is apparently the result of a "customary retirement age" effect that is not associated with particularly advantageous monetary gain.

To prove an external check of the predictive validity of the model, parameter estimates based on 1980 retirement decisions were used to predict 1981 departure rates. Departure rates at younger ages were somewhat lower in 1981 than in 1980 and were somewhat higher at older ages. Actual versus predicted cumulative departure rates, based on actual versus predicted departure rates by

TABLE III

PREDICTED VERSUS ACTUAL RETIREMENT RATES BY AGE,
BASED ON THE SINGLE-YEAR MODEL, 1980[a]

Age	Number of Observations	Annual Retirement Rates		Cumulative Rates	
		Actual	Predicted	Actual	Predicted
50	108	0.037	0.057	0.037	0.057
51	132	0.030	0.052	0.066	0.105
52	121	0.041	0.046	0.105	0.146
53	107	0.047	0.031	0.147	0.173
54	107	0.037	0.020	0.179	0.190
55	126	0.087	0.119	0.250	0.286
56	129	0.116	0.129	0.337	0.378
57	114	0.123	0.160	0.419	0.478
58	111	0.126	0.156	0.492	0.560
59	118	0.153	0.194	0.570	0.645
60	102	0.206	0.207	0.658	0.719
61	71	0.197	0.247	0.726	0.788
62	70	0.471	0.339	0.855	0.860
63	49	0.286	0.365	0.896	0.911
64	19	0.474	0.385	0.945	0.945
65	12	0.583	0.286	0.977	0.961
66	4	0.750	0.306	0.994	0.973

[a] The retirement rates were computed for the 1500 persons used to obtain the estimates reported in Table II. The predicted retirement rates are based on model 6.
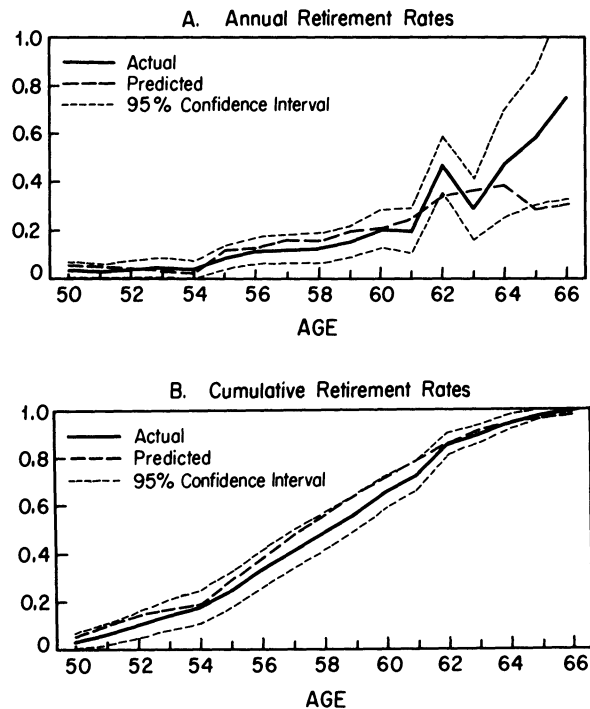

FIGURE 2.—Actual vs. predicted retirement rates.

age, provide a summary of the results. At ages 55, 60, and 62, they are as follows:[15]

	Age 55 Actual Predicted		Age 60 Actual Predicted		Age 62 Actual Predicted	
1980	.250	.286	.658	.719	.855	.860
1981	.237	.248	.688	.760	.899	.890
Change	−.013	−.038	.030	.041	.044	.030

Thus the model predicts rather closely the cumulative departure rates in each year, and also appears to capture the change in departure rates between 1980 and 1981. The actual changes were apparently due to changes in expected future earnings or to differences in the distribution of seniority by age, both of which enter the option value calculations.[16]

A better test of the predictive validity of the model is possible for clerical and other nonmanagerial administrative employees. These employees were offered a special "window" plan in 1982. Under the window plan employees were eligible for up to one year's salary, in addition to the normal pension benefits, if they retired in 1982. The retirement rates of employees eligible for the special bonus increased as much as two-fold during 1982; in 1983 retirement rates returned to their 1980 and 1981 levels. The model, estimated for clerical and other nonmanagerial employees using 1980 data, predicts closely the extremely large 1982 increase in retirement rates under the window plan. These results are presented in detail in Lumsdaine, Stock, and Wise (1989).

B. *Three Consecutive Years*

1. *Parameter Estimates*

Estimates in this section are based on the same sample of employees used to obtain the single-year estimates reported above, but those who don't leave the firm in the first year are followed for two more years. Four outcomes are possible: an employee retires in the first, the second, or the third year, or he does not retire during the three-year period. Estimated parameters of three versions of the model are shown in Table IV. The first estimates pertain to the model specification as described in Section 2-B-2, with $\rho \equiv 1$. In this case, the only difference between the multiple- and single-year versions of the model is that there are two error variances in the multiple-year version: σ_v^2, the variance in the first of the observation years (1980), and σ_ε^2, the variance of the innovation ε in the relationship $v_s = v_{s-1} + \varepsilon_s$. The estimated parameters are close to the single-year estimates (model 6 in Table II), although the estimated values of γ and k are somewhat larger. The estimated value of k is significantly

---

[15] The 1981 comparison is based on 1305 observations.
[16] Real earnings of firm employees were declining over this period.

TABLE IV

PARAMETER ESTIMATES BASED ON RETIREMENT DECISIONS
IN THREE CONSECUTIVE YEARS, 1980–1982

A. Models Without Earnings and Retirement Benefit Terms:								$\mathscr{L}$
1. Constant Only								-1536.31
2. Age Dummy Variables								-1385.45

B. *Estimated Models*:

	γ	k	k_0	k_1	Parameters β	ρ	σ_ν	σ_ε	$\mathscr{L}$
3.	0.769	1.665	—	—	0.786	1^a	0.108	0.081	-1359.32
	(0.055)	(0.123)			(0.049)		(0.012)	(0.011)	
4.	0.813	1.596	—	—	0.761	0.716	0.234	0.186	-1348.62
	(0.054)	(0.107)			(0.069)	(0.028)	(0.037)	(0.029)	
5.	0.729	—	0.950	4.87	0.777	0.778	0.202	0.153	-1336.30
	(0.084)		(0.103)	(0.31)	(0.094)	(0.107)	(0.103)	(0.067)	

^{a} Parameter value imposed.

greater than 1, however, suggesting that a dollar of income without work is indeed given greater weight in the retirement decision than a dollar of income with work. The ratio of the utility of retirement to the utility of work is one when $B/Y = 0.60$. In addition, the estimated value of β is 0.79 and is significantly less than 1, suggesting that future income is given substantially less weight than current income in the retirement decision; the real rate of time preference is rather large. The option value model based on three consecutive years fits the data substantially better than a model with an indicator variable for each age, as can be seen by comparing the likelihood values for specifications 2 and 4 (or 3) in Table IV.

The base error variance is essentially the same as the single-equation estimate. The standard deviation of the innovation (σ_ε) is only slightly smaller than the base standard deviation (.081 versus .108, or \$8,100 versus \$10,800 in utility weighted dollars). Thus, the uncertainty about an individual's valuation of the option value of continued work in future years—which stems from uncertainty about future values of ν—is much greater than the uncertainty about the option value of continued work today. This contributes to greater uncertainty about current departure decisions when comparison is made with more distant future retirement ages, that is, when departure rates of younger employees are considered. (See the discussion following equation (2.17).) On the other hand, because ν_s and ν_{s-1} are dependent, the conditional variance of ν_s given ν_{s-1} is smaller than the unconditional variance of ν_s. A measure of the persistence in the individual disturbance is indicated by the correlation between the ν's in the first and second periods of the sample; with the random walk specification, it is given by $\sigma_\nu^2/(\sigma_\nu^2)^{1/2}(\sigma_\nu^2 + \sigma_\varepsilon^2)^{1/2}$. This correlation is .800 based on the model 3 estimates in Table IV.

The possibility of less persistence is allowed by estimating ρ. Parameter estimates based on this specification are shown in model 4 in Table IV. The estimated value of ρ is .716. Judging by the likelihood values in models 1 and 2,

the χ^2 statistic relative to the hypothesis that ρ is 1 is 21.40 with one degree of freedom. Thus the strict random walk assumption is clearly rejected. However, the estimated variances increase so that the correlation between the first and second disturbance terms does not change much. In this case, it is given by $\rho\sigma_\nu^2/(\sigma_\nu^2)^{1/2}(\rho^2\sigma_\nu^2 + \sigma_\varepsilon^2)^{1/2}$. Its value, based on the model 4 estimates in Table IV, is .669, compared with a correlation of .800 based on the strict random walk assumption. Consistent with this observation, predicted average departure rates based on the two models are similar, although predictions based on the more general model are noticeably closer to the actual rates.

Unlike other empirical retirement models, age is not a variable in the option value model; it enters only indirectly through the survival probabilities $\pi(s|t)$, the wage earning forecasts, and the firm pension plan and Social Security rules. A general test of the extent to which retirement behavior is *not* determined by the monetary variables in the option value model is the gain in the model fit when age itself is added. We implement such a test by parameterizing k as a function of age, allowing the relative value of income without work to income with work to depend on age, independent of the income variables in the model. In addition, this parameterization is a way to recognize that the alternative to work at the firm may be another job, instead of retirement, and thus that the systematic portion of the model may undervalue the "retirement" option for some employees, especially at younger ages.[17]

The model 5 estimates in Table IV are based on the specification $k = k_0 (\text{Age}/55)^{k_1}$. Although this specification suggests that the value of retirement increases with age, independent of the option value measures, the other parameters of the model change little when k is parameterized. The estimates of k_0 and k_1 imply that k is .95 at age 55, 1.45 at 60, and 2.14 at 65.

2. The Model Fit

Like the single-year model, the fit of the three-year version can be evaluated by comparing predicted versus actual departure rates. The results are shown in Table V. It is analogous to the comparison presented in Table III, based on the single-year model. The annual retirement rate in Table V, for persons of a given age, is the average of the rates over the three estimation years. The cumulative figures are based on these average annual rates. Three aspects of the results stand out: First, the model fits the data well. Second, the version of the model with ρ estimated fits the data somewhat better than the version with $\rho \equiv 1$. This

---

[17] Recall that the Markov specification implies a heteroskedastic disturbance with larger variance the greater the difference between the current age and future contemplated retirement ages. Thus the variance of the individual effects is larger for younger employees. One of the unobserved determinants of departure that the random component captures is valuations of the "retirement" alternative that differ from the average. It is this aspect of the specification that allows the model to fit departure rates at younger ages, as shown in Table III for the single-year model. The mean predicted value of "retirement" may be much lower than the actual value for persons whose alternative to continuing to work at the firm is another job, not retirement. The larger random variance for younger persons allows the model to fit the data for this group—that is to predict a small likelihood of retirement for persons less than 55, for example.

TABLE V

PREDICTED VERSUS ACTUAL RETIREMENT RATES OVER A THREE-YEAR PERIOD,
BY SPECIFICATION AND AGE, BASED ON THE THREE-YEAR MODEL, 1980–1982[a]

Age	Number of Observations	Annual Retirement Rates			Cumulative Rates		
		Actual	Pred $-\rho = 1$	Pred $-\rho$ est	Actual	Pred $-\rho = 1$	Pred $-\rho$ est
50	108	0.037	0.054	0.034	0.037	0.054	0.033
51	132	0.034	0.052	0.033	0.070	0.103	0.065
52	121	0.029	0.052	0.039	0.097	0.150	0.102
53	107	0.047	0.040	0.041	0.139	0.184	0.138
54	107	0.034	0.019	0.034	0.169	0.200	0.168
55	126	0.128	0.152	0.126	0.275	0.321	0.272
56	129	0.101	0.130	0.127	0.348	0.409	0.365
57	114	0.099	0.147	0.151	0.412	0.496	0.461
58	111	0.098	0.153	0.164	0.470	0.573	0.549
59	118	0.159	0.169	0.188	0.555	0.645	0.634
60	102	0.219	0.194	0.212	0.652	0.714	0.711
61	71	0.163	0.211	0.235	0.709	0.775	0.779
62	70	0.494	0.325	0.334	0.853	0.848	0.853
63	49	0.292	0.297	0.315	0.896	0.893	0.899
64	19	0.391	0.297	0.318	0.937	0.925	0.931
65	12	0.572	0.221	0.249	0.973	0.941	0.948
66	4	0.450	0.230	0.249	0.985	0.955	0.961

[a] The entry for age 50, for example, is the probability that a person who was age 50 on January 1, 1980 left the firm by December 31, 1982.

is particularly noticeable at the younger ages. Third, the model seems to underpredict the retirement rates of the few persons that remain in the firm at older ages. In particular, both models underpredict retirement rates at 65. Even with ρ estimated, the model implies substantial persistence in individual preferences for continued employment versus retirement, consistent with the behavior of the vast majority of the sample. For example, if a person chooses not to retire at 62 when there was a reasonable ex ante probability that he would, the model uses this information to adjust downward the probability that he will retire in the next year. The results seem to suggest that this assumed persistence of tastes may not carry through age 65. Rather, there may be an age-65 "customary retirement age" effect. The sample size of older persons is so small, however—only 2.3 percent of the sample is 64 or older in 1980—that verification of this possibility will have to await estimation with larger samples of older employees; the current evidence can only be taken as suggestive.[18]

---

[18] Another way to compare predicted versus actual departure rates is to consider the probability that a person will retire during the three year period. The conclusions are similar to those discussed above. In general, both specifications fit the data well, except for employees who were 64 or older in 1980. For example, of persons who were age 58 on January 1, 1980, 43.2 percent left the firm between January 1, 1980 and December 31, 1982. The predicted retirement rate based on the specification with $\rho \equiv 1$ is 40.5 percent; it is 43.6 percent based on the specification with ρ estimated. "Pred $-\rho = 1$" and "Pred $-\rho$ est" are the predictions based on Table IV, models 3 and 4, respectively.

4. ILLUSTRATIVE SIMULATIONS

To demonstrate the importance of firm pension plan provisions on departure rates, the effects of two alternative plans are simulated. The first is a simple variant of the existing plan, increasing the early retirement age from 55 to 60. The second represents a more fundamental change, replacing the existing defined benefit plan with a defined contribution plan. In both cases the effects are quite dramatic. The simulations are based on the single-year estimates (model 6 in Table II, with $\rho \equiv 1$). Additional simulations, that compare the effects on retirement of changes in pension versus Social Security provisions, are reported in Stock and Wise (1988b).

A. Increasing the Early Retirement Age

Although retirement rates beginning at age 62 are very high, by that age most of those employed at age 50 have already left the firm. It is evident that this is due in large part to the plan's early retirement provisions. To quantify the importance of early retirement, we have simulated retirement behavior under an alternative provision. Early retirement under the alternative is at 60 instead of 55. Otherwise the alternative is like the existing plan. Persons who are employed at 60 or older face the same options under the alternative as under the existing plan.

The results are reported in Table VI. The base retirement rates are the single-year model predictions (specification 6 in Table II) under the existing plan. Under the existing plan, the model predicts that almost 65 percent of

TABLE VI

SIMULATION: EARLY RETIREMENT AGE IS 60 INSTEAD OF 55[a]

Age	Cumulative Retirement Rates			Retirement Rates		
	Base	Simulation	Difference	Base	Simulation	Difference
50	0.057	0.065	0.008	0.057	0.065	0.008
51	0.105	0.126	0.021	0.052	0.065	0.013
52	0.146	0.185	0.039	0.046	0.067	0.021
53	0.173	0.235	0.062	0.031	0.062	0.031
54	0.190	0.286	0.096	0.020	0.067	0.047
55	0.286	0.326	0.040	0.119	0.056	−0.063
56	0.378	0.359	−0.019	0.129	0.049	−0.080
57	0.478	0.391	−0.087	0.160	0.050	−0.110
58	0.560	0.413	−0.147	0.156	0.035	−0.121
59	0.645	0.422	−0.223	0.194	0.016	−0.178
60	0.719	0.542	−0.177	0.207	0.207	[b]
61	0.788	0.655	−0.133	0.247	0.247	[b]
62	0.860	0.772	−0.088	0.339	0.339	[b]
63	0.911	0.855	−0.056	0.365	0.365	[b]
64	0.945	0.911	−0.034	0.385	0.385	[b]
65	0.961	0.936	−0.025	0.286	0.286	[b]
66	0.973	0.956	−0.017	0.306	0.306	[b]

[a] Based on model 6 parameter estimates, reported in Table II. The simulation is described in the text.
[b] For persons employed at age 60 and older, the simulated alternative is the same as the base case.

those in the firm at age 50 will leave before age 60. Only 42 percent would have left if early retirement had been at age 60 instead of 55, according to the simulation results. With the existing plan, 45.5 percent of those employed at 50 leave the firm between 55 and 59. With early retirement at 60, only 13.6 percent would leave at these ages. Almost no one leaves just before age 60. On the other hand, departure rates before 55 are larger under the alternative, with a cumulative rate at 54 of .286 versus .190 under the existing plan. This reflects the longer wait before the early retirement bonus can be claimed. Still, the net reduction in departure rates before age 60 is substantial.

B. *A Defined Contribution versus the Defined Benefit Plan*

The incentive effects inherent in the firm's age-compensation profile are largely the result of the provisions of the pension plan. An alternative to a defined benefit plan is a defined contribution plan. Under a typical defined contribution plan an amount equivalent to a certain percentage of an employee's annual wage earnings is put in a pension fund. Once vested, the amount that the employee has in the fund depends only on the contributions on his behalf and on the return on these contributions. Retirement benefits are then based on the employee's accumulated assets in the fund at the time that he retires.

The effect of a change from the existing defined benefit to a defined contribution plan is illustrated under two alternatives. The first alternative is that the defined benefit contribution rate is such that for a person who has 30 years of service at age 60, the fair annuity value of the amount in the defined contribution fund is the same as the annuity value of the retirement benefits that the person would receive from the defined benefit plan were he to retire at 60. With the plan in the firm, this requires that the contribution to the defined contribution plan be approximately equal to 7.5 percent of earnings. The second alternative is that the contribution is equal to 5 percent instead of 7.5 percent of earnings. The results are shown in Table VII.

Consider first the annual departure rates based on the 7.5 percent contribution level. Again, the comparison is with the predicted departure rates based on the model 6 single-year estimates shown in Table II. There are two important features of the results. First, the jump in the departure rates at 55 is eliminated. Departure rates increase smoothly between ages 50 and 61. The jump at 62, due to Social Security provisions, remains, however. Second, although retirement rates after age 60 are lower under the defined contribution plan, departure rates at earlier ages are much higher. There is now no need to stay in the firm to receive the "retirement bonus" at 55, or to receive full benefits at 60 with 30 years of service. The net result is that more employees have left the firm by age 60 under the defined contribution than under the defined benefit plan. These results are consistent with the view that the defined benefit plan keeps employees in the firm until certain ages and then provides an incentive to leave. The defined contribution plan does not encourage them to stay, but if they do, neither does it encourage them to leave. Like departures under the current firm plan, it is likely that a large proportion of persons who leave the firm at the

TABLE VII

SIMULATION: DEFINED CONTRIBUTION VERSUS THE DEFINED BENEFIT PLAN[a]

| | | Annual Retirement Rate | | | Cumulative Retirement | | |
| | Number of | | Simulation | | | Simulation | |
Age	Observations	Base	7.5	5.0	Base	7.5	5.0
50	108	0.057	0.116	0.072	0.057	0.116	0.072
51	132	0.052	0.124	0.074	0.105	0.226	0.140
52	121	0.046	0.139	0.080	0.146	0.334	0.209
53	107	0.031	0.132	0.072	0.173	0.422	0.267
54	107	0.020	0.133	0.079	0.190	0.499	0.325
55	126	0.119	0.138	0.073	0.286	0.568	0.374
56	129	0.129	0.145	0.073	0.378	0.631	0.420
57	114	0.160	0.167	0.096	0.478	0.692	0.476
58	111	0.156	0.159	0.087	0.560	0.741	0.521
59	118	0.194	0.187	0.109	0.645	0.789	0.574
60	102	0.207	0.181	0.100	0.719	0.828	0.616
61	71	0.247	0.211	0.118	0.788	0.864	0.662
62	70	0.339	0.294	0.216	0.860	0.904	0.735
63	49	0.365	0.312	0.232	0.911	0.934	0.796
64	19	0.385	0.330	0.249	0.945	0.956	0.847
65	12	0.286	0.325	0.282	0.961	0.970	0.890
66	4	0.306	0.325	0.279	0.973	0.980	0.921

[a] Based on model 6 parameter estimates, reported in Table II. The simulation is described in the text.

younger ages under the simulated plan do so for another job, whereas at older ages most are leaving the labor force. The results with the 5 percent contribution rate are similar to those with the 7.5 percent level, except that the departure rates are lower.

5. SUMMARY AND CONCLUSIONS

We have presented a model of retirement based on the option value of continued work. A person continues to work if the expected value of retirement in the future is worth more than the value of retiring now. The model is forward looking at a point in time and allows expectations about future events to be updated as individuals age. It thus incorporates the advantages of nonlinear budget constraint formulations of the retirement decision and the advantages of continuous time hazard model formulations. The Markovian specification of the individual random effects, or the random walk special case of it, is an important component of the model. Single- and multiple-year versions of the model yield very similar results.

Predicted departure rates based on the model correspond closely to actual departure rates. In particular, discontinuous jumps in retirement rates at specific ages are captured by the model predictions. Out-of-sample predictions lend support to the predictive validity of the model.

Simulations of the effects of alternative pension plans show that plan provisions have very dramatic effects on retirement rates. For example, increasing the early retirement age from 55 to 60 would reduce by almost 35 percent the proportion of those employed at age 50 that has left the firm before age 60. At

the same time, it would increase departure rates between 50 and 55, reflecting the longer wait until the early retirement 'bonus" can be claimed.

Switching from the defined benefit to a defined contribution plan would have even greater effects on firm departure rates. The defined contribution formulation has no age-specific incentive effects. Annual departure rates of persons 60 and over would be reduced substantially. But the annual departure rates between 50 and 54 would be increased from around 3 to about 10 percent, close to the current departure rates between 55 and 59, ages that are after the early retirement age under the existing plan. The simulated departure rates under the defined contribution plan are also close to the departure rate of employees who are under 50 and have just become vested in the existing firm plan.[19] The net effect is to increase significantly the proportion of those employed at age 50 who have left the firm before age 60. These results support the view that defined benefit plans provide a strong and effective incentive for employees to stay in the firm until some age and then to retire at some later age. The defined contribution plan does neither.

These findings suggest several directions for future research. Although the results in this paper are based on the retirement decisions of employees in only one large firm, the incentive effects inherent in this firm's pension plan are typical of defined benefit plans. In future work we intend to determine whether the results are supported in similar analysis based on data from other firms. We will also explore the sensitivity of these findings to the decision rule that we have used. In particular, it would be informative to compare our results to results based on a more complex decision rule—that is closer to the dynamic programming rule, for example—as well as to results based on a less complex rule—one that does not involve estimation of utility function parameters, for example. Finally, although the use of firm data presented overriding advantages in this analysis, the data do not include information on wealth other than pension and SS assets, nor do the data indicate whether an employee left this firm for another job. Although we suspect that explicit consideration of this information would not have an important effect on our conclusions, we intend to address these issues in future work.

J.F.K. School of Government, Harvard University, Cambridge, MA 02138, U.S.A.
and
National Bureau of Economic Research, Cambridge, MA 02138, U.S.A.

APPENDIX: WAGE FORECASTING EQUATION

The estimation procedure uses earnings forecasts to compute the expected value of the utility of future income, both when employed and after retirement. Pension benefits depend on the entire earnings history of the individual at the firm up to the date of retirement. Thus pension benefits for future dates of retirement in general are based on both earnings history, known to the individual at

---

[19] See Kotlikoff and Wise (1989).

TABLE A-1

ESTIMATED INCOME FORECASTING EQUATION
DEPENDENT VARIABLE: $\Delta \ln Y_t$

Regressor	Coefficient	Standard Error	t statistic
$\ln(A_t)$	-0.205	0.150	-1.368
$\ln(S_t)$	0.075	0.039	1.908
$\ln(A_t)^2$	0.036	0.022	1.604
$\ln(S_t)^2$	0.021	0.003	6.891
$\ln(A_t)\ln(S_t)$	-0.050	0.012	-4.069
$\Delta \ln Y_{t-1}$	0.360	0.071	5.081
$\ln(A_t)\Delta \ln Y_{t-1}$	-0.160	0.021	-7.622
$\ln(S_t)\Delta \ln Y_{t-1}$	0.024	0.008	2.978
$\Delta \ln Y_{t-2}$	-0.152	0.070	-2.169
$\ln(A_t)\Delta \ln Y_{t-2}$	0.016	0.021	0.755
$\ln(S_t)\Delta \ln Y_{t-2}$	0.007	0.008	0.948
$D72$	0.447	0.255	1.750
$D73$	0.418	0.255	1.639
$D74$	0.407	0.255	1.595
$D75$	0.356	0.255	1.395
$D76$	0.417	0.255	1.636
$D77$	0.416	0.255	1.631
$D78$	0.450	0.255	1.762
$D79$	0.369	0.255	1.446
$D80$	0.289	0.255	1.134
$D81$	0.364	0.255	1.426
$D82$	0.347	0.255	1.361
$D83$	0.376	0.255	1.475
$D84$	0.355	0.255	1.393
SEE $= 0.173$			

the current date, and forecasts of future earnings. For example, in 1981, estimates of the pension benefits that would be received were retirement in 1986 involve known earnings through 1981 and forecasts for the remaining years.

The income forecasting equation, shown in Table A-1, was estimated using 98,465 observations, including multiple observations for the same person, taken from a panel of individuals in the same job category in the same firm as the 1500 individuals that were used in the estimation results reported in the text. The earnings data cover 1969–1984. Nominal earnings were converted to 1980 dollars using the consumer price index. S_t and A_t respectively denote years of service at the firm and age; D72, D73, etc. are dummy variables for the indicated years. Income forecasts were computed using the average of the coefficients on the dummy variables for 1978–1980.

The data set contains earnings from 1969 on. Thus earnings before 1969 (for those who joined the firm before 1969) were back-cast using a specification similar to the forecasting equation in Table A-1. The estimates (not reported here) were obtained using the same specification, except that time was reversed in the sense that all lags were replaced by leads.

REFERENCES

BERKOVEC, JAMES, AND STEVEN STERN (1988): "Job Exit Behavior of Older Men," *Econometrica*, forthcoming.

BERNHEIM, B. DOUGLAS (1987): "Dissaving After Retirement: Testing the Pure Life Cycle Hypothesis," in *Issues in Pension Economics*, ed. by Z. Bodie, J. Shoven, and D. Wise. Chicago: University of Chicago Press, 237–279.

BLINDER, ALAN, ROGER GORDON, AND DONALD WISE (1980): "Reconsidering the Work Disincentive Effects of Social Security," *National Tax Journal*, 33, 431–442.

BULOW, JEREMY (1981): "Early Retirement Pension Benefits," NBER Working Paper No. 654.

BURKHAUSER, RICHARD V. (1979): "The Pension Acceptance Decision of Older Workers," *Journal of Human Resources*, 14, 63–75.

——— (1980): "The Early Acceptance of Social Security: An Asset Maximization Approach," *Industrial and Labor Relations Review*, 33, 484–492.

BURTLESS, GARY (1986): "Social Security, Unanticipated Benefit Increases, and the Timing of Retirement," *Review of Economic Studies*, 53, 781–805.

BURTLESS, GARY, AND JERRY A. HAUSMAN (1978): "The Effect of Taxation on Labor Supply: Evaluating the Gary Negative Income Tax Experiment," *Journal of Political Economy*, 86, 139–160.

BURTLESS, GARY, AND ROBERT A. MOFFITT (1984): "The Effects of Social Security on the Labor Supply of the Aged," in *Retirement and Economic Behavior*, ed. by H. Aaron and C. Burtless. Washington, DC: Brookings Institution, 135–174.

CLARK, ROBERT L., AND ANN A. McDERMED (1984): "Life Cycle Patterns of Earnings, Pension Wealth, and Total Compensation," Mimeo.

DIAMOND, PETER, AND JERRY A. HAUSMAN (1984): "Individual Retirement and Savings Behavior," *Journal of Public Economics*, 23, 81–114.

FEINSTEIN, JONATHAN, AND DANIEL McFADDEN (1989): "The Dynamics of Housing Demand by the Elderly: I. Wealth, Cash Flow, and Demographic Effects," in *The Economics of Aging*, ed. by D. Wise. Chicago: University of Chicago Press, 55–86.

FIELDS, GARY S., AND OLIVIA MITCHELL (1982): "The Effects of Pensions and Earnings on Retirement: A Review Essay," in *Research in Labor Economics*, Vol. 5, ed. by R. Ehrenberg. Greenwich, CT: JAL Press, 115–156.

——— (1985): *Retirement, Pensions, and Social Security*. Cambridge, MA: MIT Press.

FRANT, HOWARD L., AND HERMAN B. LEONARD (1987): "Promise Them Anything: The Incentive Structures of Local Public Pension Plans," in *Public Sector Payrolls*, ed. by D. Wise. Chicago: University of Chicago Press, 215–242.

GUSTMAN, ALAN, AND THOMAS STEINMEIER (1986): "A Structural Retirement Model," *Econometrica*, 54, 555–584.

HAUSMAN, JERRY A., AND DAVID A. WISE (1980): "Discontinuous Budget Constraints and Estimation: The Demand for Housing," *Review of Economic Studies*, 47, 75–96.

——— (1985): "Social Security, Health Status, and Retirement," in *Pensions, Labor, and Individual Choice*, ed. by D. Wise. Chicago: University of Chicago Press, 159–191.

HOGARTH, JEANNE M. (1988): "Accepting an Early Retirement Bonus: An Empirical Study," *Journal of Human Resources*, 23, 21–33.

HURD, MICHAEL, AND MICHAEL BOSKIN (1981): "The Effect of Social Security on Retirement in the Early 1970s," *Quarterly Journal of Economics*, 46, 767–790.

HURD, MICHAEL, AND JOHN SHOVEN (1983): "The Economic Status of the Elderly," in *Financial Aspects of the United States Pension System*, ed. by Z. Bodie and J. Shoven. Chicago: University of Chicago Press, 359–397.

HURD, MICHAEL, AND DAVID A. WISE (1989): "The Wealth and Poverty of Widows: Assets Before the Husband's Death," in *The Economics of Aging*, ed. by D. Wise. Chicago: University of Chicago Press, 177–222.

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. by D. Wise. Chicago: University of Chicago Press, 55–85.

——— (1987): "The Incentive Effects of Private Pension Plans," in *Issues in Pension Economics*, ed. by Z. Bodie, J. Shoven, and D. Wise. Chicago: University of Chicago Press, 283–339.

——— (1989): "Employee Retirement and a Firm's Pension Plan," in *The Economics of Aging*, ed. by D. Wise. Chicago: University of Chicago Press, 279–330.

LAZEAR, EDWARD P. (1983): "Pensions as Severance Pay," in *Financial Aspects of the United States Pension System*, ed. by Z. Bodie and J. Shoven. Chicago: University of Chicago Press, 57–89.

LAZEAR, EDWARD P., AND ROBERT L. MOORE (1988): "Pensions and Turnover," in *Pensions in the U.S. Economy*, ed. by Z. Bodie, J. Shoven, and D. Wise. Chicago: University of Chicago Press, 163–188.

LUMSDAINE, ROBIN L., JAMES H. STOCK, AND DAVID A. WISE (1989): "Windows and Retirement," *Annales d'Economie et de Statistique*, forthcoming.

MERRILL, SALLY (1984): "Home Equity and the Elderly," in *Retirement and Economic Behavior*, ed. by H. Aaron and G. Burtless. Washington, DC: Brookings Institution.

PHILLIPS, DOUGLAS, AND DAVID A. WISE (1987): "Military versus Civilian Pay: A Descriptive Discussion," in *Public Sector Payrolls*, ed. by D. Wise. Chicago: University of Chicago Press, 19–46.

RUST, JOHN (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55, 999–1034.

——— (1988): "Behavior of Male Workers at the End of the Life Cycle: An Empirical Analysis of States and Controls," Mimeo.

——— (1989): "A Dynamic Programming Model of Retirement Behavior," in *The Economics of Aging*, ed. by D. Wise. Chicago: University of Chicago Press, 359–398.

SOCIAL SECURITY ADMINISTRATION (1982): *Social Security Bulletin, Annual Statistical Supplement*. Washington, D.C.: U.S. Department of Health and Human Services.

STOCK, JAMES H., AND DAVID A. WISE (1988a): "Pensions, the Option Value of Work, and Retirement," NBER Working Paper No. 2686.

——— (1988b): "The Pension Inducement to Retire: An Option Value Analysis," NBER Working Paper No. 2660.

VENTI, STEVEN F., AND DAVID A. WISE (1984): "Moving and Housing Expenditure: Transaction Costs and Disequilibrium," *Journal of Public Economics*, 23, 207–243.

——— (1988): "But They Don't Want to Reduce Housing Equity," NBER Working Paper No. 2859.

——— (1989): "Aging, Moving, and Housing Wealth," in *The Economics of Aging*, ed. by D. Wise. Chicago: University of Chicago Press, 9–48.

——— (1989b): "Aging and the Income Value of Housing Wealth," Mimeo.