University of Michigan Retirement Research Center

Working Paper

WP 2002-026

The Retirement Behavior of Federal Civil Service Workers

Beth Asch, Steven Haider, and Julie Zissimopoulos

MRRC

Project #: UM01-09

“The Retirement Behavior of Federal Civil Service Workers”

Beth Asch
RAND

Steve Haider
RAND

Julie Zissimopoulos
RAND

June 2002

Michigan Retirement Research Center
University of Michigan
P.O. Box 1248
Ann Arbor, MI 48104
www.mrrc.isr.umich.edu
(734) 615-0422

Acknowledgements

This work was supported by a grant from the Social Security Administration through the Michigan Retirement Research Center (Grant # 10-P-98358-5). The opinions and conclusions are solely those of the authors and should not be considered as representing the opinions or policy of the Social Security Administration or any agency of the Federal Government.

Regents of the University of Michigan
David A. Brandon, Ann Arbor; Laurence B. Deitch, Bingham Farms; Daniel D. Horning, Grand Haven; Olivia P. Maynard, Goodrich; Rebecca McGowan, Ann Arbor; Andrea Fischer Newman, Ann Arbor; S. Martin Taylor, Gross Pointe Farms; Katherine E. White, Ann Arbor; Mary Sue Coleman, ex officio

The Retirement Behavior of Federal Civil Service Workers

Beth Asch, Steven Haider, and Julie Zissimopoulos

Abstract

We examine the retirement behavior of federal civil service workers. This research contributes to the literature that more generally examines how retirement behavior responds to financial incentives. The civil service workers in our study provide an interesting case study because they do not participate in the Social Security system, they are only covered by a defined benefit pension plan, and this pension plan is significantly different from the Social Security system in the structure of its incentives. Moreover, there is widespread concern among policy makers of a pending retirement crisis in the federal civil service. Relying on an option value framework, our main results suggest that federal civil service workers respond to their retirement incentives in a manner that is quite similar to the responses that others have found looking at much different retirement systems. Such a result provides important additional evidence regarding the generality of previous results. On the other hand, unlike previous studies, we find little evidence of a spike in the retirement rate at age 65, nor do we find much evidence of "excess retirements" or a large fraction of retirements at age 65 that are unexplained by our financial incentive model. While past studies have attributed this age 65 effect to "social norms," those norms do not seem important to the federal civil service workers we study.

Author's Acknowledgement

We gratefully acknowledge the financial support of the Michigan Retirement Research Center. We wish thank Craig Martin for his programming assistance, Michael Dove at the Defense Manpower Data Center for his help obtaining our data, and Michael Hurd, David Loughran, and Stan Panis for many helpful discussions. All opinions and errors remain our own.

The Retirement Behavior of Federal Civil Service Workers

1. Introduction

Because of the aging of the United States population and the trend towards early retirement, the relative size of the retired population is changing dramatically. For example, 9 persons were paying into Social Security for every 1 person receiving benefits in 1955; this ratio declined to 3.4 to 1 in 2001 and is expected to decline to 2 to 1 by 2030 (U.S. General Accounting Office, 2000). This trend has important implications for the solvency of federal entitlement programs such as Social Security and Medicare, and potential legislative changes for these programs have been the subject of ongoing debate. A critical piece of information in assessing many changes to these entitlement programs and to pension systems more generally is the impact of financial incentives on retirement behavior.

Numerous studies have examined the impact of financial incentives on retirement behavior over the last three decades. These studies have focused on various aspects of the financial incentive and specified numerous behavior and empirical models to analyze their impact. Although part of the differences across studies stems from the variation in incentives that people actually face, a large part also stems from the data that researchers have available. Quite simply, the data requirements for computing social security wealth (e.g., long earnings histories and complicated benefit formulas) and pension wealth (e.g., varying and complicated formulas that often interact with social security wealth) are large, and simplifying assumptions are often necessary.

In this paper, we contribute to this literature by examining the retirement behavior of a unique set of workers, the Department of Defense civil service workers who are covered by the Civil Service Retirement System (CSRS). One important reason to study the retirement behavior of these workers is that their financial incentives are relatively simple. Civil service workers covered by CSRS do not participate in the Social Security system, and they are only covered by a defined benefit pension plan in which benefits are a function of years of service and highest salary. This information is available in our data so that we can compute financial incentives much more precisely. In addition, increasing evidence suggests that the individuals often know very little about their retirement incentives (e.g., Gustman and Steinmeier, 2001). The plan simplicity would suggest that individuals are much more likely to know their incentives, and thus provides a "best case" in assessing retirement behavior.

Another reason to study the retirement behavior of civil service workers is that the ages that are important in their retirement scheme are not coincident with the ages that are important to Social Security and Medicare. Previous studies have concluded that their exists "excess retirement" at ages 62 and 65 because, after the incentives for Social Security and Medicare are accounted for, more people retire at these ages then the models would otherwise suggest. This excess retirement is often interpreted as the impact of social norms. We can examine our data for the existence of excess retirement, as well as at what ages it occurs, to shed further light on the findings from previous studies.

A final reason for studying the retirement behavior of the CSRS-covered population is the widespread concern among policy makers of a pending retirement crisis in the federal civil service. According to the President's Management Agenda, approximately 70 percent of the federal government's current permanent employees will be eligible for early retirement by 2010 and they estimate that 40 percent of them are expected to retire

(Office of Management and Budget, 2002). The General Accounting Office placed this issue on the "high-risk list" of federal activities in 2001. Moreover, numerous conferences and commissions have been convened to study the issue, including a high-visibility commission chaired by former Federal Reserve Chairman Paul Volcker. Despite this importance, only one prior study has explicitly analyzed federal civil service retirements (e.g., Smith and Sylwester, 1988); however, this study uses a retirement model that is not forward-looking, whereas significant subsequent research has pointed to the importance of forward looking models (e.g., Lumsdaine, Stock and Wise, 1992). Direct estimates of the retirement behavior of civil service workers can be used to inform policy decisions regarding these workers.

Relying on an option value framework, our main result suggests that federal civil service workers respond to their retirement incentives in a manner that is similar to the response that other studies have found using data on much different retirement systems. Such a result provides important additional evidence regarding the generality of previous estimates. On the other hand, we find limited evidence of "excess retirement" at the key retirement ages of the civil service workers (55, 60, and 62). Although past studies find spikes at key retirement ages, especially age 65, that are unexplained by financial incentives embedded in their models our models perform quite well and seem to capture well the financial incentives for individuals to retire. Finally, we preliminarily find that our data do not support the estimation of a structural option value model. Although this finding implies that we cannot estimate the parameters of an underlying utility function, it does provide further evidence that the reduced-form option value results are robust to a wide array of assumptions.

The paper is organized as follows. We first describe different forward-looking retirement models and review past empirical implementations of them. Next, we describe the data we use and how we selected our sample. We then provide some descriptive statistics about the federal civil service workers in our sample and present the results from estimating the different retirement models. Lastly, we present our conclusions and directions for future research.

2. Financial Incentives and Retirement Behavior

The literature on retirement behavior is very large.[1] Researchers have relied on numerous types of empirical and behavioral models that vary tremendously with respect to complexity and data requirements. In this paper, we focus on several models that can usefully be viewed as special cases of the so-called "option value model" of retirement. Option value models are forward-looking models in which individuals assess the value of retiring today versus remaining at work so that they have the option of retiring in the future. These models stop short of specifying a full dynamic program that more completely specifies the uncertainty in the retirement decision (see Gotz and McCall, 1987; Berkovec and Stern, 1991; Rust, 1989; Daula and Moffitt, 1995; French, 2001; Rust and Phelan, 1997).[2] However, dynamic programs often require other simplifications due to their complexity, and previous research suggests that these simplified models still perform quite well relative to more complicated models (see Lumdsaine, Stock, and Wise, 1992).

---

[1] See Hurd (1990) and Leonesio (1996) for useful reviews. See Samwick (1998), Coile (1999), Coile and Gruber (2000), Gustman and Steinmeier (2002) for more recent examples.
[2] Asch and Warner (1994) calibrate the parameter values for a dynamic program.

2.1. The basic option value model

We first describe the basic option value model, following the notation of Stock and Wise (1990). Consider an individual who is currently working in year *t*. Let Y_s be earnings in year *s* if the individual is still working and $B_s(r)$ be retirement benefits in year *s* if the individual retires in year *r*. Denote utility while working as $U_W(Y_s)$ and utility while retired as $U_R(B_s(r))$. Then, an individual's value at time *t* of retiring in year *r* can be represented as,

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

where β is the subjective discount rate and *S* is the year the individual dies. Furthermore, the individual's expected gain from retiring in year *r* versus retiring today (year *t*) can be expressed as

(2) $$G_t(r) = E_t V_t(r) - E_t V_t(t),$$

where E_t is the expectation at time *t*.

The basic retirement decision is then characterized by whether there is a future year of retirement that returns an expected net gain to the individual. If such a future year exists, the person is assumed to continue to work. Formally, define *r** to be the future year that maximizes the expected value of retiring ($r^* = \arg\max_r G_t(r),\ \forall r > t$). Then, letting *R* be a random variable that equals one if the individual retires, the decision rule for whether a person retires is assumed to be

(3) $$R = \begin{cases} 0 & \text{if } G_t(r^*) \geq 0 \\ 1 & \text{if } G_t(r^*) < 0 \end{cases}.$$

Various sets of assumptions can then be made to empirically implement this decision rule.

2.2. Empirical implementations of this approach

We group the various sets of assumptions into three broad classes: reduced-form option value models, financial models, and structural option value models. We first present the reduced form option value model in order to make clear the interpretation of the financial models as restrictions on this model. We then discuss the financial models and the structural option value models.

2.2.1. Reduced-form option value (RFOV) models

The RFOV empirical implementation assumes the values for the underlying structural parameters of the utility functions in equation 1. The form of the utility function that is often used is the constant relative risk aversion, with a risk aversion parameter γ and an additional parameter *k* to allow for differences in the value of income flows while working and retired. Specifically, the utilities are parameterized as follows:

(4)
$$\begin{aligned} U_W(Y_s) &= Y_s^{\gamma} \\ U_R(B_s(r)) &= (kB_s(R))^{\gamma} \end{aligned}.$$

The RFOV assumes the values of these two parameters, as well as the subjective discount factor β. Additionally, it is assumed that workers face age/gender-specific mortality risk that is independent of income; denote the probability of living to year *s*, conditional on being alive in year *t*, is $\pi(s \mid t)$. The maximum expected utility gain for remaining at work can then be written as,

(5) $$G_t(r^*) = \sum_{s=t}^{r^*-1} \beta^{s-t}\pi(s \mid t) E_t y_s^{\gamma} + \sum_{s=r^*}^{S} \beta^{s-t}\pi(s \mid t) E_t (kB_s(r^*))^{\gamma} - \sum_{s=t}^{S} \beta^{s-t}\pi(s \mid t) E_t (kB_s(t))^{\gamma}.$$

Finally, it is usually assumed that the worker and researcher know the value of future earnings and benefits, allowing the expectations in equation 5 to be ignored.

With these assumptions, equation 5 can simply be computed for each individual. We can then specify a standard dichotomous outcome model that is based on the decision rule in equation 3,

(6) $$\Pr[R = 1] = \Pr[a + bG_t(r^*)].$$

Thus, the RFOV can be estimated as a probit or logit, and various other regressors can be entered.

Such models were estimated in Stock and Wise (1990), Chan and Stevens (2001), Coile (1999), Coile and Gruber (2000), Samwick (1998), and Samwick and Wise (2001). These studies tend to assume utility function parameters that were estimated in structural option value studies such as Stock and Wise (1990), one of the few that actually estimated the utility function parameters. They also tend to include a variable for social security and/or retirement wealth, intended to capture the wealth effect of retirement benefits and their positive effect on the demand for all goods, including leisure.

2.2.2. Financial models

A series of studies have implemented this basic approach by adopting a set of assumptions that are sufficient to reduce the retirement incentive to a simple financial incentive. The first assumption is that the utility function is simply a revenue function by setting γ and k equal to one (see equation 4). Setting γ equal to one is akin to assuming that the marginal utility of consumption is not decreasing, and setting *k* equal to one is akin to assuming that there is no disutility associated with work. Second, these models assume that the discount factor is known, with a typical rate being 0.90. Once again, these assumptions (in addition to the assumption that future earnings are known) are sufficient so that the value of continuing work can be computed directly, and the model reduces to a standard dichotomous outcome model (equation 6).

Such financial models have been commonly used since the late 1970s to examine the retirement and retention behavior of military personnel (Warner, 1978; Warner and Goldberg, 1984; Smith, Sylwester, and Villa, 1991). More recently, they have been used to model the retirement and separation behavior of federal civil service workers (Black, Moffitt, and Warner, 1990; Asch and Warner, 1999). In such implementations, these models are referred to as Cost of Leaving (COL) models or, with an additional adjustment for the length of time over which the costs and benefits are realized, Annualized Cost of Leaving (ACOL) models.

More recently, Coile (1999) and Coile and Gruber (2001) use a similar financial model to examine the financial incentives of the social security system, referring to their model as a "peak value" (PV) model. Coile and Gruber make several further simplifying assumptions in calculating their incentive variable. Specifically, they ignore the incentive to work that stems from wages in their financial calculation. Rather, they define a peak value incentive measure as the difference in expected pension wealth if someone retires at time r relative to retiring today, time t, appropriately discounted,

$$PV_t(r) = \sum_{s=r}^{S} \beta^{s-t} \pi(s \mid t) E_t(B_s(r)) - \sum_{s=t}^{S} \beta^{s-t} \pi(s \mid t) E_t(B_s(t)) . \tag{7}$$

Like the option value model, the peak value model assumes that the individual retires at the *r** that maximizes the value of retiring. Coile and Gruber estimate a model of retirements as a function of the Social Security and pension peak value variables, and a variable representing the level of Social Security wealth. The former variables capture the incentive effects of Social Security and pensions on retirement timing while the latter variable captures its wealth effect.

2.2.3. Structural option value models

Stock and Wise (1990) not only provided the first exposition of the structural option value (OV) model represented by equations 1 to 3, but they actually estimated the utility parameters as well. To implement the model, the CRRA utility function is used once again, but now utility shocks are included,

$$\text{(8)} \qquad \begin{aligned} &U_W(Y_s) = Y_s^{\gamma} + \omega_s \\ &U_R(B_s(r)) = (kB_s(R))^{\gamma} + \xi_s \end{aligned} .$$

The error terms for the utility function are assumed to be Gausian Markov

$$\text{(9)} \qquad \begin{aligned} &\omega_s = \rho\omega_{s-1} + \varepsilon_{Ws}, \quad E_{s-1}\varepsilon_{Ws} = 0 \\ &\xi_s = \rho\xi_{s-1} + \varepsilon_{Rs} \quad E_{s-1}\varepsilon_{Rs} = 0 \end{aligned} .$$

Just as in the previous models, it is assumed that the econometrician knows precisely how individuals forecast future earnings. Importantly, it is assumed that the utility shocks in a particularly period are observed before a person chooses to work or retire.

To derive an estimable form of the structural OV model, equation 3 is rewritten as a standard probability statement,

$$\text{(10)} \qquad \Pr[R = 1] = \Pr[G_t(r^*) < 0].$$

The assumptions structurally build an error term into the model that can be interpreted as an unobserved utility shock. If an individual is found to retire at an age that differs from her peak, it is assumed that this deviation was caused by a particular draw from the utility function.

Structural OV models are estimated infrequently because they are more computationally difficult to implement.[3] First, the probability statement in equation 10 is a highly non-linear function of the underlying utility parameters. Second, the optimal retirement date depends on the utility function parameters. Thus, as one searches for new utility parameters, the optimal retirement needs to be re-evaluated. Third, as more time periods are added, the model becomes dramatically more difficult to estimate because the structural error terms are correlated. Stock and Wise (1990) estimate models based on a cross-section of individuals and on panel data for three periods, and Lumsdaine, Stock and Wise (1992) use cross-sectional data. We discuss our estimation method and the assumptions we made to empirically implement the full OV model in Appendix A.

2.3. Other relevant literature

Many studies only focus on the financial incentives inherent in the Social Security system (see discussion in Stock and Wise, 1990). However, Stock and Wise (1990) and Samwick (1998) find that pensions are empirically much more important than Social Security when computing the accrual of financial wealth over age. This focus on social security is not surprising, given the immense amount of information necessary to calculate the incentives inherent in a pension scheme. Pension schemes vary enormously across individuals and often have very complicated rules (see Kotlikoff and Wise, 1987).

Moreover, approximately half of defined benefit pension holders are in plans where the benefits depend on social security (Slusher, 1998).

In the process of developing option value models of retirement, Stock and Wise (1990) and Lumsdaine, Stock, and Wise (1996) have noted that these models predict well the spikes in retirement rates at key ages, such as 55, 60, and 62. However, they tend to significantly underpredict the spike in the retirement rate at age 65. Only by including age dummies, an essentially ad-hoc approach, are they able to predict the spike at age 65. These papers find the high age-65 retirement is not explained by Medicare eligibility and attribute the age-65 excess retirement effect to social norms. Because CSRS does not embed any financial incentive to specifically retire at age 65, and CSRS-covered employees are not covered by social security, the study of CSRS-covered federal employees offers a good opportunity to identify the presence of an age-65 effect.

Another complication in the study of retirement incentives is that there is growing evidence that retirees do not have complete knowledge about their retirement plans. Using the Health and Retirement Study data, Gustman and Steinmeier (1999, 2001) find that many individuals are misinformed or lack information about their expected pension and Social Security benefits and about the features of the plans that cover them. It is possible that individuals do not respond to the incentives inherent in a retirement system because they simply do not know or do not understand them. However, the lack of information may itself reflect optimizing behavior. Information is costly to obtain and individuals may be making appropriate investments in their level of information; similarly, individuals may have sufficient information to know not to retire, but not know

[3] To our knowledge, the only examples of structural OV models is Stock and Wise (1990); Lumsdaine Stock and Wise (1992); and Ausink and Wise (1996).

the characteristics that make this the right decision. Regardless, CSRS personnel are likely to be among the better informed about retirement benefits because the plans are simple and because the benefits likely to provide the majority of retirement wealth.

3. The Data

In this section, we describe the structure of the CSRS retirement plan and the data that we use to analyze retirement behavior.

3.1. The Civil Service Retirement System (CSRS)

Until 1987, the Civil Service Retirement System (CSRS) was the primary retirement system covering federal civil service personnel. Because CSRS was legislated in the 1920s before the Social Security system was created, civil service employees participating in CSRS are not covered by Social Security. In 1987, the Federal Employees Retirement System (FERS) was created; importantly, FERS includes Social Security coverage. FERS covers federal civil service personnel hired for the first time after December 31, 1983 and those rehired with less than 5 years of service (YOS). Those re-hired after December 31, 1983 and who have more the 5 YOS are given the option to switch to FERS within 6 months of employment. Those who do not switch to FERS are covered by a system called CSRS-interim (later called CSRS-Offset), which included both CSRS and Social Security coverage. In 1987 and 1988, those under CSRS with more than 5 YOS were permitted to switch coverage and participate in FERS. Another opportunity to switch was given to them in 1998. Participants of CSRS, CSRS-offset, and FERS are also covered by Medicare.

CSRS is a typical defined benefit retirement plan. Benefits are vested after 5 years of service (YOS), and the benefit level is determined by the individual's highest 3 years of earnings and his or her YOS. The normal retirement age is determined by ones' years of service. Individuals who reach age 55 with 30 YOS are entitled to receive full benefits, individuals who reach age 60 with 20 YOS are entitled to receive full benefits, and individuals who reach age 62 with 5 YOS are entitled to receive full benefits.

Those who separate before they have become eligible to retire can claim benefits at age 62 if they have at least 5 years of service. Their annuity is based on the highest three years of earnings at the time of separation. Consequently, their pension annuity is eroded by inflation between the date of separation and age 62. However, those who are eligible to retire get a pension annuity that is adjusted annually by the full CPI amount. Thus, the benefit is essentially inflation protected for those who are retirement eligible at the time of separation. This protection creates a strong incentive to stay in the civil service until retirement eligibility is reached as will be seen below when we show pension accruals and peak values under CSRS.

The benefit formula under CSRS equals 1.5 percent of an individual's highest three-year average earnings times his or her years of service (YOS) for the first five YOS, plus 1.75 percent of the highest-three average earnings times YOS for the next five YOS, plus 2 percent of the highest-three average earnings times all YOS over 10. The maximum annuity an individual can receive is 80 percent of the highest-three average earnings. Normally, this is acquired after 41 years of credible civilian and military service.

3.2. The Department of Defense Civil Service (DoDCS) personnel data files

We limit our analysis to permanent federal civil service personnel in the Department of Defense (DoD), covering fiscal years 1980 to 1996. The data were provided by the Defense Manpower Data Center (DMDC) and they represent administrative personnel records for the entire population of permanent workers in DoD during this time frame. DoD is the largest employer of federal civil service workers outside of the Post Office, employing an annual average of approximately 900,000 permanent workers over the 17 years covered by our data. The personnel record includes information on age, years of federal service, retirement system coverage, demographic characteristics (e.g., gender, race/ethnicity, reported handicaps), geographic location, educational level, pay plan (e.g., General or Work Schedule), pay grade and step, annual federal earnings, and other job characteristics including occupation and functional work area. These data files contain individuals covered by CSRS, CSRS-Offset, and FERS.

Although past studies such as Stock and Wise (1990) have also used administrative personnel records to analyze retirement behavior, such data have several limitations. First, the data exclude information on marital status, non-earned income, assets, health status, and other factors that may affect the retirement decision. While we would prefer to have such information, we note that Samwick (1998) finds financial assets, marital status, and health status either to have a statistically insignificant effect or a small effect on retirement. Second, we only observe individuals exiting from the civil service, not from the labor force. Samwick (1998) compares his estimates in which exits are defined with respect to firms versus the labor force, and he finds that the latter definition provides an underestimate of the effect of financial incentives on retirement from the labor force. Thus, our estimates will likely understate the effect of CSRS on labor force withdrawal. However, for some purposes, the impact on firm exits is exactly the right concept. For example, from the perspective of a firm influencing its workforce, they are exactly interested in the impact of incentives on leaving the firm.

Offsetting these potential limitations are the advantages of using such administrative data. First, the data include millions of records on individuals covered by a single retirement plan. As discussed below, we take a one-in-ten random sample to reduce computational time, but we are still left with hundreds of thousands of observations. The large number of records contrasts with the typical situation found in survey data, such as the HRS. In those data, the sample sizes are often quite small, numbering in the hundreds, and the respondents are covered by a heterogeneous set of retirement plans, some of the features of which may be unknown to the researcher. Second, information on years of service, a key component defining retirement eligibility and benefits, is provided by the employer and not based on the possibly inaccurate recollection of respondents. Finally, we have extremely good information about retirement plan and changes over our data period.

3.3. Selection of our sample and development of our analysis file

Given our focus on retirement behavior, we restrict our data to civilian personnel in DoD who are between ages 50 and 70. Consequently, we do not model the decision to separate before or stay until age 50. We also limit our analysis to those covered by CSRS.

Participation in CSRS, particularly in the later years of our data, i.e., in the 1990s, is not entirely exogenous. As mentioned earlier, those who entered the civil service before 1984 are automatically covered by CSRS. However, individuals covered by CSRS and who had more than 5 years of service had the opportunity to switch to FERS during an "open season" that spanned from July 1987 to July 1988. Thus, those covered by CSRS in the post-1988 part of our data include those who opted to stay under CSRS. If the choice to to stay under CSRS or to switch to FERS is associated with characteristics that are also associated with the effect of retirement incentives on retirement behavior, our estimated effects of CSRS on retirement may be biased. However, from a practical standpoint, such selection bias is unlikely to be a problem in our analysis because we limit our analysis to those with 15 or more years of service, as discussed below. Earlier work suggests that the incentive to switch is very small for those with many years of service or who are older in age (Asch and Warner 1999). Consequently, it is highly unlikely that there are many individuals in our data who had an incentive to switch to FERS. We therefore ignore any selection problem that may arise by our focus on CSRS-covered personnel.

A key advantage of studying CSRS-covered is that personnel are not covered by Social Security by virtue of their federal service. However, it is possible that these personnel held jobs in the past (or plan to hold jobs in the future) that were or will be covered by Social Security. Insofar as these individuals accumulate sufficient number of quarters of social security coverage, their retirement behavior might be influenced by Social Security incentives.

We made two data restrictions (in addition to limiting our analysis to those covered by CSRS) in an attempt to ensure that no one in our data has Social Security coverage. First, we exclude those with less than 15 years of service. Individuals with fewer years of service are likely to have employment spells in covered jobs. Thus, we excluded them. Second, we eliminated all personnel whose record suggests that they had served in one of the armed services prior to becoming a federal employee. Since 1956, military personnel in the active and reserve components have been covered by social security. About 20 percent of new entrants to the DoD federal civil service have prior military service (Asch, 2001). We deleted these individuals because their retirement behavior might be influenced by social security coverage.

A final selection criterion addresses a problem that has been documented earlier about the DoD civil service personnel data. When the annual record for each individual is strung together over time, the years of service variable does not always increment in a sensible manner. In some cases, years of services may actually decrease, or jump by more the one from year to year (Asch and Warner, 1999). About 10 percent of the records had this problem in each year. These records were deleted.

In constructing our analysis file, we took into consideration several policy changes that occurred from 1982 to 1996. The most important of them was the dramatic downsizing that occurred in DoD following the end of the cold war (beginning around 1991) that resulted in a significant drop in the size of its civilian (and military) workforce (Asch and Warner, 1999). Our reduced form regression models include fiscal year dummy variables to account for changes in the size of the federal workforce over our data period and we exclude individuals who retired under an "early out window".[4] Another policy change of note is the Federal Employees Pay and Compensation Act (FEPCA) of 1990. FEPCA changed how the federal government adjusts the federal pay tables each year to reflect cost-of-living increases. All federal employees get the same baseline annual pay change. In addition, they get a change that depends on their specific geographic location. The location specific pay change is intended to account for differences in the cost-of-living change across geographic areas. These locality-specific pay adjustments began in 1994. Therefore, beginning in 1994, our annual pay variable also includes both the base adjustment and the locality adjustment.

4. Who Are the Department of Defense Civil Service Workers?

Although the Social Security Program is extremely large, coverage is not universal. The large number of federal employees under CSRS, together with many state and local government workers, and railroad workers, represent a large segment of workers who are not covered by Social Security. Table 1 shows the extent of Social Security coverage in terms of wages and salaries. In 1996, Social Security covered about 92 percent of the non-self-employed wages and salaries. Although only about 8 percent of wages and salaries were uncovered, it represents about $300 Billion in 1996. Federal employment wages and salaries represent a significant fraction of this uncovered amount.

Table 2 reports means and standard deviations of the sample. For a description of the means and standard deviations of all the explanatory variables used in the reduced form regression, see the Appendix. We compare our sample of civil service workers to a sample of respondents from the first wave of the Health and Retirement Study. From the HRS we select individuals 51-61 years old in 1992 and who are working fulltime. Our CSRS sample has 33 percent higher earnings than the HRS sample. They are also more likely to be a high school graduate and about equally likely to be a college graduate. The CSRS sample is younger by about 2 years and is less likely to be male and white.

CSRS embeds strong incentives to retire at its normal retirement ages. This point can be observed in Table 3 where we show the mean, expected present discounted value

---

[4] Our tabulations suggest only about 500 individuals were offered an early out widow, so we ignore them for now. These individuals will provide an interesting test of the parameter estimates for future work (see Lumsdaine, Stock and Wise, 1992).

(PDV) of CSRS pension wealth for our sample. The calculation assumes an annual earnings growth rate of 0.25 percent, a real interest rate of 3 percent, an inflation rate of 4 percent. The calculations are made for the first year in which we observe the workers in our sample, and we calculate the means separately for those who have a normal retirement age (NRA) of 55 and 60.

For those with an NRA of 55, we find that the PDV of pension wealth more than doubles at age 55 as compared to age 54. For those with an NRA of 60, pension wealth increases with age and then rises by over 35 percent at between ages 59 and 60. Beyond age 55 in the first case and age 60 in the second, pension wealth declines with age as the effects of fewer years of pension receipt (given an assumed death age of 99) offsets the growth in earnings and the increase in years of service.

Table 4 presents descriptive statistics for the CSRS pension peak value (see equation 7) at each age. The peak value captures the financial option value embedded in the CSRS pension system. However, as noted above, it ignores earnings in the civil service, and the value of leisure and external opportunities. The variation at each age reflects differences in years of service and earnings, the two other factors defining pension wealth at each age under CSRS. As shown in Table 4, the mean value becomes negative at age 60.

The peak values under CSRS in Table 4 are much higher at each age than the Social Security peak values reported by Coile and Gruber (2000, Table 4). For example, at age 55, the median pension peak value under CSRS is $81,109, about 4 times larger than 21,260 reported by Coile and Gruber for Social Security. At age 59, the figures are $48,344 and $13,714, respectively. The mean peak value becomes negative at age 60 under CSRS, but remains positive at $12,381 under Social Security. As shown by Coile and Gruber, the pension peak value for Social Security becomes negative at age 65.

Before presenting estimates of the effect of CSRS on retirement, it is interesting to consider the aggregate retirement hazards, shown in Figure 1. Two points are worth noting. First, the retirement hazards spike up at ages 55, 60, and 62, the three normal retirement ages embedded in the CSRS pension formula. Second, we see a small spike at ages 64, an age that have no particular significance under CSRS, and no spike at age 65..

5. Empirical Results

In this section, we first present results from the financial and the reduced form option value models. We include variables intended to capture features of the worker's budget constraint including pension wealth, the financial incentive variables and earnings. The pension wealth variable also captures the wealth effect of retirement benefits. The financial incentive variables also capture the incentive to retire while the earnings variable captures the incentive to continue to work. We also control for other factors related to retirement including disability, sex, education, occupation and age. We then turn to preliminary estimates from the structural option value model.

5.1. Financial and reduced-form option value model results

Table 5 presents logistic regression model results. The first specification follows Coile and Gruber (2000) and includes the CSRS pension peak value as the measure of the retirement incentives. The peak value is denominated in dollars thus allowing for ease in interpretation. The second specification presents the reduced-form option value model results in which utility function parameters are assumed. Based on the original Stock and

Wise (1990) study and the other papers that estimate this model, we set k to 1.5, γ to 0.75, and β to 0.95.

Our results are consistent with earlier studies. Using logistic regression we estimate a negative and statistically significant effect of the peak value and of the reduced form option value on the logit index function. Our coefficient estimate for the peak value (measure in $10,000) is -0.023. The estimate is the correct sign—an increase in the peak value associated with staying in the civil service reduces the probability of retirement—and is statistically significant at the 1 percent level. The coefficient implies that a $10,000 increase in peak values decreases the retirement rate by 0.0075 or 10% of the average retirement rate. This translates into an elasticity of 0.10 --a one percent increase in peak value decreases retirement by 0.1 percent (see Table 3).

Similarly, our estimate on the option value of retirement is negative and statistically different than zero at the 1 percent significance level (0.04). The coefficient estimate in this model is not directly comparable to the one in the peak value model because it is expressed in utility rather than in dollar units. To examine the economic content of the option value model, we use the parameters estimated from the option value model to simulate the effect on retirement of a decrease of 20 percent of retirement wealth. The model predicts an average retirement rate that is 2 percent lower than what the retirement rate would be without the 20 percent decrease in wealth. This represents a 35 percent decrease from the mean retirement rate.

In a model incorporating Social Security and pension wealth, Coile and Gruber (2000) find that $1,000 in peak value lowers retirement by 0.5 percent of the sample average retirement rate. Our estimate effect of peak value likewise leads to a slightly larger, 1 percent decrease of the average retirement rate for a $1,000 increase in peak value. Samwick and Wise (2001) estimate that a $1,000 increase in their accrual measure (accrual to age 65) reduces baseline retirement probabilities by 1.8 percent. Thus, our estimates are remarkable similar to estimates for workers covered by Social Security and pensions in the HRS as analyzed by other studies.

The results for earnings are mixed. As discussed earlier, the peak value measure does not include earnings, current or future, in its calculation. In the model that includes the peak value incentive measure, earnings have the expected negative effect on retirement. In the option value model of retirement, however, compensation also has an independent effect on retirement although the effect is now positive. We also include the log of pension wealth in the regression to account for a wealth effect of retirement independent of the incentive effect as measured by the peak value or option value. The coefficient on log wealth in the peak value regressions implies a $10,000 increase in pension wealth increases retirement by 0.043 percentage points. This is an increase in relative risk of retirement of 57 percent. In the option value model, the effect is smaller and implies an increase in relative risk of retirement of 19 percent.

Another way to assess model performance is to examine how well each model in Table 5 can predict actual retirement behavior but without relying on the underlying age patterns. In other words, we re-estimate the models presented in Table 5 (results not shown) not including the age dummy variables, and then predict retirement behavior. As can be observed in Figure 2, the models without age do good job of tracking the increase in retirement at age 55 and age 60. Moreover, the peak value and reduced-form option value models perform remarkably similarly. Returning to the main models, we also include a full set of age dummies in both reduced form specifications (see Table 5). These dummies capture the effects of age on the retirement probit, over and above their effects through the retirement incentive variables. Generally, we find statistically significant age dummy effects in both specifications and interpret this as 'excess retirements' -- retirements associated with age that are not explained by the incentive variables. On the other hand, the magnitude of the age dummy effects are not large and while statistically different than age 55, the age dummies are generally not statistically different than the prior or following age. For example, the largest age effect is at age 60, but even at that age, the change in the retirement rate implied by the coefficient estimate is only 9 percent. Thus, we do not find the age 65 spike that previous studies have found (Lumsdaine, Stock, Wise 1996; Phelan and Rust, 1991, 1993; Stock and Wise 1990), nor do we have much evidence supporting the presence of excess retirements at 65 or at any other retirement age. Civil service workers in our sample, who have at least 5 years of service have access to retiree health insurance upon separation from the job thus the availability of Medicare at age 65 is unlikely to cause an age 65 spike in retirement. The lack of an effect at age 65 also suggests that there is no 'social norm' associated with retirement at age 65.

We include several other covariates to capture characteristics of the job and the individual. The regressions also include education, rating, pay scale, years of service, grade, occupation, race, disability, agency in the civil service, and fiscal year dummies. Many of the effects are statistically different than zero suggesting, as expected, that retirement is affected by characteristics other than financial incentives.

5.2. Structural option value model results (PRELIMINARY)

The primary difference between the RFOV and the structural OV model is that the RFOV assumes particular utility function parameter values, whereas the structural OV model estimates these parameters directly. There are at least two motivations for estimating structural OV models. First and foremost, the underlying utility function parameters are interesting in their own right. Second, it is of interest to determine whether there might be better utility function parameters for the RFOV model than those currently used.

There are fewer examples of structural OV models estimated in the literature, however. One reason, as discussed previously, is that it is much more difficult to implement empirically. A second reason seems to be that data do not always support its estimation. For example, Samwick (1998) reports that his attempts at estimating a structural OV model were unsuccessful, and Ausink et al. estimate parameter values that theoretically are hard to believe.

We also estimate a structural OV model. We adopt assumptions that are very similar to Stock and Wise (1990) and Lumsdaine, Stock, and Wise (1992), with specific details provided in the Appendix. The main difference between our specification and these previous implementations is that we do not assume that the difference in error terms is mean zero, but rather has mean α. The motivation for this change is that it structurally builds an intercept into the model. When an intercept is included, the estimated parameters make a little more sense and the restriction that α equals zero is rejected by the data.[5]

---

[5] We have estimated models with and without setting the intercept equal to zero. We find that the estimation is much more numerically stable when the intercept is included and that we tend to get

Rather than providing point estimates for the structural OV model, we instead map out much of the likelihood surface. The motivation for this presentation of results is three-fold. First, our data does not seem to support the model in that parameter estimates tend to converge to non-sensical values; examining the likelihood surface is a convenient method to demonstrate these results. Second, because the model only has three structural parameters of interest (γ, β, and k), it is possible to graphically examine the likelihood surface fairly easily; the other two structural parameters (σ and α) are primarily scaling parameters. Third, the computational burden of mapping the likelihood surface is not appreciably more difficult than estimating the actual parameter values.[6] However, we still draw a simple one percent random sample of the individuals in our data to further reduce the computational burden; the resulting sample size of 6649 observations is still much larger than that used in many previous studies.

To examine the likelihood surface of the structural OV model, we choose values for the three substantive parameters (γ, β, and k) and then estimate the two scaling parameters (σ and α). Fixing the three substantive parameters causes the structural OV model to reduce to a simple probit model in which the structural scaling parameters can be recovered from the simple probit parameters. We repeat this estimation process for numerous values of the substantive parameters. To examine the likelihood surface, we graph the quantity of negative 2 multiplied by the log-likelihood of these restricted

---

parameter values that are more sensible. For example, we estimate rather reasonable estimates of the disutility of work (k) in the results presented below. When the intercept is excluded, the estimated parameter value is less than 1.

[6] To understand this claim, consider equation A4. The basic option value is much more complicated than a regular probit because the optimal retirement date r^* must be re-evaluated as new parameter values are chosen to maximize the likelihood function and the likelihood function likely has kinks because the optimal retirement age is discrete and because many people face similar and pronounced retirement incentives. Thus, probit routines in standard statistical packages cannot directly be applied. However, once the three

models; because of this transformation, maximizing the likelihood of the structural OV model is equivalent to looking for the minimum of the surface that we graph. We present two such graphs of the log-likelihood surface in Figures 3 and 4.

Figure 3 presents the surface plotted with respect to β and k, setting γ equal to 0.6. As can be observed, the graph slopes towards a higher k and a lower β. The likelihood surface becomes fairly flat in the direction of k after 2.0, but the surface is still strongly sloping towards a lower β even at a value of 0.50. Figure 4 examines the surface in the dimensions of γ and k, setting β equal to 0.50. After a value for k of 2.3, the likelihood surface is very flat in the dimension of k and γ; there is slight slope towards a smaller γ. These basic conclusions regarding the shape of the surface are quite robust to focusing on other values of γ and β. Moreover, these graphs also are indicative of the results when we estimate the five structural parameters jointly: joint estimation tends to produce a value for β that is very low compared to what would be expected when the estimation actually converges, and joint estimation will often not converge because it cannot find a unique value for γ and k.

These structural OV results are still preliminary, so we refrain from speculating about their implications regarding the behavior of the DoD civil service workers and about the performance of the structural OV model. However, it is interesting to note a couple other results from the literature. First, Samwick (1998, p. 222) reports, "In my attempts to estimate the parameters of the option value model on the SCF sample, the parameters for the value of leisure in retirement…and the discount rate…could not be simultaneously

---

structural parameters are specified, then A4 reduces to a standard probit model (but with no intercept) with the coefficient interpreted as the inverse of σ.

identified with any precision." Second, Samwick (2000) discusses several studies that have come to varying conclusions regarding the magnitude of the discount rate.

More interestingly, these structural OV results suggest important directions for sensitivity analysis of the RFOV model. Namely, by varying the utility function parameter values in the direction suggested by the structural OV results, we can examine how robust the RFOV model is to assumptions that fit the data better.[7] Specifically, we re-estimate the basic RFOV model, without age dummies, to predict retirement but use alternative utility parameters. Figure 2 presents additional predictive results for two more RFOV models: RFOV-2 sets γ =0.1, k=2.0, and β =0.95 and RFOV-3 sets γ =0.6, k=2.8, and β =0.6. As can be observed, both of these additional models predict retirement very similarly to the PV model and the RFOV model with more realistic parameter values (RFOV-1). These results suggest that the retirement elasticities are quite robust to alternative assumptions.

6. Conclusions

It is of immense interest to understand how retirement behavior responds to financial incentives, and numerous papers have focused on many different aspects of these incentives with many different models. We contribute to this literature by examining the retirement behavior of federal civil service workers in the Department of Defense, the largest federal employer outside of the U.S. Postal Service. These individuals provide an interesting case study because they do not participate in the Social Security system, they

[7] In performing this sensitivity analysis, it is important to note that the RFOV model differs from the structural OV model in respect to a scaling factor that arises from the idiosyncratic part of utility (see equation A4). The models would be more similar if the structural OV model assumed that the idiosyncratic utility shocks were *iid* rather than Markovian. Under such an assumption, the structural OV model instead

are only covered by a standard defined benefit pension plan, and this pension plan is significantly different in structure regarding retirement incentives. Moreover, there is widespread concern among policy makers of a pending retirement crisis in the federal civil service.

Relying on an option value or "forward-looking" framework, our main result suggests that federal civil service workers respond to their retirement incentives in a manner that is quite similar to the response other studies have found, using data on much different retirement systems. For example, we obtain an elasticity of 0.1 in the peak value model where Coile and Gruber (2000) estimate a .05 elasticity and Samwick and Wise (2001) estimate a .18 elasticity. Such a result is surprising given the much different nature of the pension schemes. Moreover, the similarity of elasticities provides important evidence regarding the generality of previous estimates.

In addition, we find little evidence of "excess retirement" at the key retirement ages of the civil service workers (55, 60, and 62). Thus, our results suggest that the option value model does a good job of capturing the incentives for individuals to retire. In contrast, past studies find that option value models under-predict retirement rates at key ages, particularly age 65, without the inclusion of age dummies in the model. Previous research has ruled out Medicare eligibility as an explanation and has speculated that "excess retirement" could be due to social norms. High age-65 retirement due to Medicare eligibility is unexpected in the population of civil service workers because these workers have access to retiree health insurance. Since civil service workers are

---

results in very high values of β rather than very low values. This empirical finding is consistent with results that Coile and Gruber (2000) report for their analysis of HRS data.

members of the larger U.S. labor market, it is not obvious why they would not be subject to same social norms.

Finally, we preliminarily find that our data do not support the estimation of a structural option value model. Although this finding does not allow us to estimate the structural parameters of a utility function, this finding does provide further evidence that the reduced-form option value results are robust to a wide array of assumptions.

Appendix

We provide further details about the data and the structural option value estimation methods in the appendix.

Data. In Table A1, we provide basic descriptive statistics for other variables used in our analysis. Short descriptions of the variables are provided whenever the variables are not self-explanatory.

A few variables deserve special note.

Estimating a structural option value model. Our implementation of the OV model follows Stock and Wise (1990) very closely; throughout the appendix, we will simply refer to the paper as SW. Again, the error terms are structurally built into the model by appending unobserved (to the econometrician), idiosyncratic shocks to the CRRA utility functions (see equation 8), and then the decision rule to retire can be characterized by a probability statement (see equation 10). To develop an estimable model, it is useful to first rewrite the probability statement by substituting (8) and (1) into (2),

$$\begin{aligned}
G_t(r) &= E_t\left[\sum_{s=t}^{r-1}\beta^{s-t}(Y_s^{\gamma}+\omega_s)+\sum_{s=r}^{S}\beta^{s-t}((kB_s(r))^{\gamma}+\xi_s)\right] \\
&\qquad - E_t\left[\sum_{s=t}^{S}\beta^{s-t}((kB_s(r))^{\gamma}+\xi_s)\right] \\
&= E_t\sum_{s=t}^{r-1}\beta^{s-t}Y_s^{\gamma}+E_t\sum_{s=r}^{S}\beta^{s-t}(kB_s(r))^{\gamma}-E_t\sum_{s=t}^{S}\beta^{s-t}(kB_s(r))^{\gamma} \qquad \text{(A1)} \\
&\qquad + E_t\sum_{s=t}^{r-1}\beta^{s-t}(\omega_s-\xi_s) \\
&= g_t(r)+\phi_t(r),
\end{aligned}$$

where $g_t(r)$ is the first three terms and $\phi_t(r)$ is the last term.

We now make various assumptions to empirically implement the model. The first assumption is that individuals evaluate the future until age 99 and the probability of living until year s at time t, $\pi(s \mid t)$, is independent of earnings. With such an assumption, the quantities $g_t(r)$ and $\phi_t(r)$ simplify to

$$\text{(A2)} \quad \begin{aligned} g_t(r) &= \sum_{s=t}^{r-1} \beta^{s-t}\pi(s)E_t Y_s^\gamma + \sum_{s=r}^{99} \beta^{s-t}\pi(s)E_t(kB_s(r))^\gamma - \sum_{s=t}^{99} \beta^{s-t}\pi(s)E_t(kB_s(r))^\gamma \\ \phi_t(r) &= \sum_{s=t}^{r-1} \beta^{s-t}\pi(s)E_t(\omega_s - \xi_s). \end{aligned}$$

The second assumption allows us to evaluate the expectation of future earnings and retirement benefits. SW approximates $E_t Y_s^\gamma$ using a second order Taylor expansion and approximates $E_t(kB_s(r))^\gamma$ with the approximation $(k\hat{B}_s(r))^\gamma$. We choose to evaluate both expectations using SW's latter method of approximation. Thus, we instead approximate $E_t Y_s^\gamma$ with $\hat{Y}_s^\gamma$. These approximations are akin to assuming that the workers forecast their future earnings and benefits without error. Such an assumption is more reasonable for the workers in the civil service than it would be for the general population, given the strict pay grades that exist.

The key assumption regarding how difficult the model is to estimate rests with assumptions on the error terms in the utility function (ω_s, ξ_s), which in turn gives the structure for the expression $\phi_t(r)$. Like SW, we assume that the error terms are Gaussian Markov with a zero mean. This assumption results in a composite error term ν_s that is also Gaussian Markov with a zero mean. Under this assumption, the expression $\phi_t(r)$ can be re-written as

$$\text{(A3)} \quad \begin{aligned} \phi_t(r) &= \sum_{s=t}^{r-1} \beta^{s-t}\pi(s)\rho^{s-t}\nu_t \\ &= K_t(r)\nu_t. \end{aligned}$$

The probability of retiring is then simply

$$\text{(A4)} \quad \begin{aligned} \Pr[R=1] &= \Pr[g_t(r^*) + \nu_t K_t(r^*) < 0] \\ &= \Phi[\frac{-g_t(r^*)/K_t(r^*)}{\sigma}], \end{aligned}$$

where $\Phi[.]$ is the cumulative normal distribution function.

Two additional but related assumptions remain. The first rests with whether one should use cross-sectional or panel data. As panel data is used, the evaluation of probability statement (A4) becomes much more difficult because it is a multinomial discrete choice problem in which higher-dimensional integrals must be evaluated. SW estimates both cross-sectional and panel versions. However, their panel version only includes three years of panel data, and Lumsdaine, Stock and Wise (1992) conclude that the cross-sectional version is sufficient. The second is that observed retirement in year *t* depends on a worker not retiring previously. SW concludes that this problem is intractable and ignores it. However, this implies that their sample is necessarily weighted towards individuals who chose to delay retirement.

In our structural OV estimation, we take somewhat of a middle ground. First, we rely on the cross-sectional model but we use the panel data. This method has the benefit that it corresponds more closely to the implementation of the RFOV and PV models and that it weights the sample more appropriately towards actually retirement behavior. However, the method has the drawback that it ignores that the same individuals are observed repeatedly in the data, and thus does not take adhere exactly to the Markov assumption.

Table A1: Mean and Standard Deviation of Other Variables

	Mean	Std. Deviation
Annual earnings	44538	17986
Ln(PDV pension)	12.0262	0.8184
Pay grade:	9.2311	3.3840
Performance rating: used in promotion, scale 1=outstanding, 5=unsatisfactory	1.4996	1.0660
No rating: no performance rating provided	0.1227	0.3281
Years of federal service	26.6151	5.8266
Pay Plan-General Schedule:	0.7121	0.4535
Pay Plan-WC: Corps of Engineers	0.0007	0.0271
Pay Plan-WG: Wage Grade, non-supervisory	0.1054	0.3071
Pay Plan-WS: Wage grade, supervisory	0.0235	0.1516
Pay Plan-WL: Wage grade, Leader	0.0079	0.0887
Male	0.4051	0.4909
Occupation-Blue collar	0.1497	0.3567
Occupation-Professional	0.1744	0.3794
Occupation-Administrative	0.3118	0.4632
Occupation-Technical	0.1639	0.3701
Occupation-Clerical	0.1990	0.3992
Has a disability:	0.1511	0.3581
Education-Less than High school	0.0877	0.2828
Education-Some college	0.2520	0.4342
Education-College	0.1031	0.3041
Education-Graduate degree	0.1211	0.3262
Race-White	0.7441	0.4362
Race-Black	0.1529	0.3598
Race-Hispanic	0.0502	0.2183
Race-Other	0.0529	0.2239
Agency-Army	0.3491	0.4770
Agency-Navy	0.2686	0.4432
Agency-Marine	0.0188	0.1357
Agency-Airforce	0.2524	0.4344
Agency-Other	0.1111	0.3143

Note: This table provides descriptive statistics for the other variables we use in our analysis. Number of observations is 636,331.

References

Asch, Beth, and John Warner, *A Policy Analysis of Alternative Retirement Systems*, RAND, MR-465-OSD, 1994.

Asch, Beth J. and John Warner*, Separation and Retirement Incentives in the Federal Civil Service: A Comparison of the Federal Employees Retirement System and the Civil Service Retirement System*, RAND, MR-986-OSD, 1999.

Asch, Beth J., *The Pay, Promotion, and Retention of High-Quality Civil Service Workers in the Department of Defense*, RAND, MR-1193-OSD, 2001.

Ausink, John, and David Wise, "The Military Pension, Compensation, and Retirement of U.S. Air Force Pilots," in David A. Wise (ed.), *Advances in the Economics of Aging*. Chicago: University of Chicago Press, 1996.

Berkovec, J. and S. Stern, "Job Exit Behavior of Older Men," Econometrica, Vol. 59, 1991, pg 189-210

Black, Matthew, Robert Moffitt, and John Warner, "The Dynamics of Job Separation: The Case of Federal Employees," *Journal of Applied Econometr*ics, Vol. 5, 1990, pg 245-262.

Chan, Sewin, and Ann Huff Stevens, "Retirement incentives and Expectations," National Bureau of Economic Research Working Paper 8082, 2001.

Coile, Courtney, "Retirement Incentives and Couples' Retirement Decisions," Ph.D Dissertation, MIT, 1999.

Coile, Courtney and Jonathan Gruber, "Social Security and Retirement" NBER Working Paper #7830, 2000.

Coile, Courtney, and Jonathan Gruber, "Social Security Incentives for Retirement," in David Wise (ed.) *Themes in the Economics of Aging*, Chicago: University of Chicago Press, 2002.

Daula, Thomas and Robert Moffitt, "Estimating Dynamic Models of Quit Behavior: The Case of Military Reenlistment*," Journal of Labor Economics*, Vol. 13, No. 3, 1995, pg 499-523.

French, E. "The Effects of Health, Wealth, and Wages on Labor Supply and Retirement Behavior," unpublished manuscript, Federal Reserve Bank of Chicago, 2001.

Gotz, Glenn and John McCall*, A Dynamic Retention Model for Air Force Officers: Theory and Estimates*, Santa Monica, CA: RAND, 1984, R-3028-AF.

Gustman, Alan, and Thomas Steinmeier, "Imperfect Knowledge, Retirement, and Savings, National Bureau of Economic Research Working Paper 8406, Cambridge, MA: August 2001.

Gustman, Alan and Thomas Steinmeier, "Social Security, Pensions, and Retirement Behavior in the Family," National Bureau of Economic Research Working Paper 8772, 2002.

Gustman, Alan and Thomas Steinmeier, "What People Don't Know About Their Pensions and Social Security: An Analysis Using Linked Data From the Health and Retirement Study," National Bureau of Economic Research Working Paper 7368, 1999.

Hogan, Paul, and Matthew Black, "Reenlistment Models: A Methodological Review" In Curtis Gilroy, David Horne, and D. Alton Smith (eds.), Military Compensation and Personnel Retirement: Models and Evidence. Alexandria, VA: U.S. Army Research Institute for the Behavioral and Social Sciences, 1991.

Hurd, Michael, "Research on the Elderly: Economic Status, Retirement, and Consumption and Savings," Journal of Economic Literature, Vol. 28, No. 2, 1990, pg 565-637.

Kotlikoff, L. and D. Wise, "The Incentive Effects of Private Pension Plans," in Issues in Pension Economics, Z. Bodie, J. Shoven, and D. Wise (Eds.), University of Chicago Press: Chicago, 1987.

Lumsdaine, Robin, James Stock, and David Wise, "Three Models of Retirement: Computational Complexity versus Predictive Ability," in David A. Wise (ed.), *Topics in the Economics of Aging*. Chicago: University of Chicago Press, 1992.

Lumsdaine, Robin, James Stock, and David Wise, "Why are Retirement Rates So High at Age 65," in David A. Wise (ed.), *Advances in the Economics of Aging*. Chicago: University of Chicago Press, 1996.

Office of Management and Budget, "The President's Management Agenda, Fiscal Year 2002," Washington DC: Government Printing Office, 2002.

Rust, J., and C. Phelan, "How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Information," Econometrica, Vol. 65, 1997, pg. 781-831.

Samwick, Andrew, "New Evidence on Pensions, Social Security, and the Timing of Retirements," *Journal of Public Economics*, Vol. 70, 1998, pg 207-236.

Samwick, Andrew, and David Wise, "Option Value Estimation with HRS Data," Working Paper, March 2001.

Slusher, Chuck, "Pension Integration and Social Security Reform," Social Security Bulletin, Vol. 61, No. 3, 1998, pg 20-27.

Smith, D. Alton, and Stephen Sylwester, "The Federal Retirement Decision: An Analysis of DoD Civilians," Systems Research and Applications Corporation (SRA), Arlington Virginia, 1988.

Smith, D. Alton, Stephen Sylwester, Christine Villa, "Army Reenlistment Models," In Curtis Gilroy, David Horne, and D. Alton Smith (eds.), Military Compensation and Personnel Retirement: Models and Evidence. Alexandria, VA: U.S. Army Research Institute for the Behavioral and Social Sciences, 1991.

Stock, James, and David Wise, "Pensions, the Option Value of Work, and Retirement," Econometrica, Vol. 58, No. 5, 1990, pg 1151-1180.

Stock, James, and David Wise, "The Pension Inducement to Retire: An Option Value Analysis," in David Wise (ed.) *Issues in the Economics of Aging*, Chicago: University of Chicago Press, 1990, pg 205-224.

United States Government Accounting Office, "High-Risk Series: An Update," Washington DC: GAO-01-263, January 2001.

United States Government Accounting Office, "Managing in the New Millennium: Shaping a More Efficient and Effective Government for the 21st Century," Washington DC: GAO-/T-OCG-00-9, March 29, 2000.

Warner, John and Matthew Goldberg, "The Influence of Non-Pecuniary Factors on Labor Supply: The Case of Navy Enlisted Personnel*," Review of Economics and Statistics*, Vol. 66, 1984, pg 26-35.

Warner, John, "Analysis of the Retention Effects of the Proposed Retirement System," in *Supplementary Papers of the President's Commission of Military Compensation*. Washington, DC: U.S. Government Printing Office, April 1978.

Warner, John, Military Compensation and Retention*: An Analysis of Alternative Models and Simulation of a New Model*, Alexandria, VA: Center for Naval Analyses Research Contribution 436.

Table 1: Wage and Salary Coverage by Major Public Retirement Programs [Billions]

		Wage and salary covered by major retirement programs			
Year	Total wage and salary	OASDHI	Railroad	Fed. civil servants	State/local gov'ts
1981	1,510	1,445	13	56	135
1986	2,095	1,896	12	72	190
1991	2,828	2,565	12	92	271
1996	3,632	3,328	13	107	365

Note: These data were taken from the Annual Statistical Supplement to the Social Security Bulletin, 1999, Table 3.B2, p. 141. The categories in the last four columns are not mutually exclusive. Starting in 1984, for example, some federal civil servants could elect to be part of the Social Security system. Wages that are earned in the civil service and are covered by Social Security would be counted in both columns.

Table 2: Descriptive Statistics for Primary CSRS Sample and HRS Sample

	CSRS (N=87,867)		HRS (N=)	
	Mean	Std. dev.	Mean	Std. dev.
Annual wages and salaries (2000$)	42,252	16,849	33,612	35,583
Pension coverage	1	--	0.759	0.438
Birth year	1935.3	3.8		
Less than high school degree	0.092	0.290	0.234	0.439
High school	0.420	0.494	0.322	0.466
Some college	0.249	0.433	0.214	0.407
College degree	0.096	0.294	0.230	0.408
Male	0.424	0.494	0.597	0.498
White	0.750	0.433	0.875	0.392
Black	0.150	0.357		
Hispanic	0.051	0.219		

Note: The primary sample is selected from the CSRS-covered civil service workers in the Department of Defense, 1982-1996, with age and other restrictions as described in the text. The HRS sample is full-time workers from the 1992 survey wave between the ages of 51 and 61. .

Table 3: Mean Prospective Pension Wealth by Normal Retirement Age

Retirement Age	Normal Age of 55 (N=38,546)	Normal Age of 60 (N=31,248)
50	140,993	66,297
51	151,241	71,017
52	162,721	76,546
53	175,622	84,562
54	189,326	93,121
55	427,838	102,257
56	416,143	112,002
57	403,758	122,391
58	390,750	133,463
59	377,148	145,255
60	362,817	195,833
61	347,390	190,865
62	330,783	185,385
63	313,212	179,444
64	294,808	173,093
65	256,856	155,392
66	243,315	149,508

Note: This table is based on the first year we observe everyone in our primary sample. We compute their expected PDV pension wealth for retiring at each age. We make the following assumptions: earningsgrowth=0.25%, real interest=3%, and inflation=4%. All dollar values are in 2000$. See text for additional details.

Table 4: Descriptive Statistics for the Pension Peak Value by Age

Age	N	Mean	Std. Dev	10th percentile	Median	90th percentile
50	68,727	226,094	105,520	108,757	202,518	379,531
51	75,080	219,927	106,210	103,632	195,394	374,473
52	81,340	213,187	106,943	98,263	187,846	369,210
53	76,864	212,042	107,592	96,531	186,160	369,543
54	72,612	208,646	107,763	93,089	182,406	365,878
55	67,291	75,281	92,192	-19,274	83,343	196,194
56	53,860	68,505	82,342	-20,492	78,820	174,168
57	40,768	59,596	72,097	-21,625	71,422	150,518
58	31,120	48,287	61,929	-22,552	60,970	126,772
59	23,244	36,407	52,049	-23,494	48,886	102,196
60	16,960	-12,021	8,905	-24,740	-9,093	-3,925
61	10,897	-13,351	9,513	-26,984	-10,244	-4,778
62	6,991	-14,559	10,027	-28,890	-11,297	-5,590
63	4,224	-15,901	10,720	-30,915	-12,449	-6,326
64	2,573	-17,619	11,828	-33,845	-13,742	-7,070
65	1,436	-19,542	12,699	-36,951	-15,215	-7,940

Note: This table presents the descriptive statistics for every person-year in our primary sample. See notes for Table 3 and the text for further details.

Table 5a: Retirement Logits for Financial and Reduced Form Models

Parameter	Peak Estimate	Std. Error	Option Value Estimate	Std. Error
Intercept	-13.3066	0.4338	-6.9169	0.3567
Financial Incentive				
Peak value (x1000)	-0.0023	0.0001		
Option value (x1000)			-0.0400	0.0012
Annual earnings (x1000)	-0.0300	0.0032	0.0130	0.0033
earnings squared (x10-8)	0.0133	0.0021	-0.0073	0.0021
Ln(PDV Pension wealth)	0.9446	0.0442	0.3160	0.0339
Age 50	-1.0186	0.0448	-0.8985	0.0436
Age 51	-0.8836	0.0391	-0.7792	0.0382
Age 52	-0.6771	0.0328	-0.5916	0.0321
Age 53	-0.6081	0.0297	-0.5433	0.0292
Age 54	-0.4325	0.0266	-0.3956	0.0262
Age 55 (omitted)	--	--	--	--
Age 56	-0.2001	0.0181	-0.2253	0.0181
Age 57	-0.3249	0.0199	-0.3758	0.0200
Age 58	-0.3544	0.0215	-0.4319	0.0217
Age 59	-0.4201	0.0238	-0.5274	0.0240
Age 60	0.0967	0.0238	-0.0315	0.0242
Age 61	-0.0758	0.0287	-0.2339	0.0292
Age 62	-0.0099	0.0340	-0.1977	0.0344
Age 63	-0.1904	0.0435	-0.4072	0.0438
Age 64	-0.2100	0.0522	-0.4533	0.0523
Age 65	-0.1901	0.0728	-0.4624	0.0727
Age 66	-0.1551	0.1037	-0.4462	0.1034
Age 67	-0.2154	0.1920	-0.5283	0.1913
Age 68	-0.8507	558.20	-1.2041	571.50
Age 69	-0.1024	6739.60	-0.3617	6739.60

Note: These regressions are based on the primary sample. The reduced form option value model assumes the following parameters: γ =0.75, k=1.5, and β =0.95. See the notes for Table 3 and the text for additional details. Regressions also include fiscal year dummies. Number of observations is 636,331

Table 5b: (Continued) Retirement Logits for Financial and Reduced Form Models

Parameter	Estimate	Std. Error	Estimate	Std. Error
Grade	0.0251	0.0100	0.1260	0.0101
Performance rating	0.1245	0.0069	0.1255	0.0069
No performance rating	0.3834	0.0228	0.3823	0.0228
Years of service	0.0449	0.0043	0.0684	0.0035
Grade*Years of service	-0.0009	0.0003	-0.0043	0.0003
Pay plans				
WC	0.3537	0.1961	0.3477	0.2017
WG	-0.0344	0.0472	-0.0219	0.0472
WS	0.0387	0.0531	0.0580	0.0533
WL	-0.1149	0.0693	-0.0927	0.0693
Male	-0.0266	0.0170	-0.2047	0.0171
Occupation				
Blue collar	-0.3999	0.1310	-0.4318	0.1314
Professional	-1.0290	0.1262	-1.0695	0.1265
Administrative	-0.8516	0.1248	-0.8791	0.1252
Technical	-0.8702	0.1244	-0.8877	0.1247
Clerical	-0.9908	0.1245	-1.0155	0.1248
Has a disability	0.1745	0.0138	0.1756	0.0138
Education				
Less than High school	0.0150	0.0195	0.0187	0.0194
High school (omitted)	--	--	--	--
Some college	-0.1463	0.0131	-0.1489	0.0131
College	-0.3260	0.0237	-0.3399	0.0239
Graduate degree	-0.6018	0.0272	-0.6086	0.0273
Black	-0.1459	0.0150	-0.1445	0.0150
Hispanic	-0.1608	0.0238	-0.1584	0.0237
Other	-0.3937	0.0249	-0.3895	0.0249
White (omitted)	--	--	--	--
Agency				
Army (omitted)	--	--	--	--
Navy	-0.0200	0.0139	-0.0221	0.0139
Marine	-0.2984	0.0408	-0.2928	0.0408
Airforce	0.0013	0.0146	0.0031	0.0146
Other	0.0667	0.0172	0.0702	0.0172

Note: Coefficients are a continuation from Table 5a. Number of observations is 636,331

Figure 1: CSRS Retirement Hazard

Note: This figure presents the probability of retiring by age in the primary CSRS sample.

Figure 2: Predicted Retirement Hazards

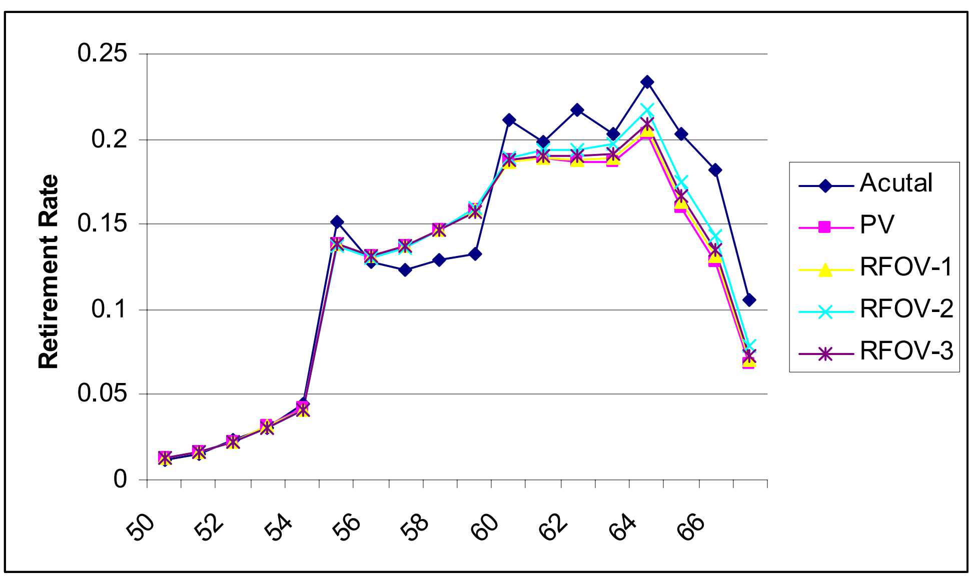

Note: This figure presents the probability of retiring by age in the primary CSRS sample (Actual) and the predicted probability of retiring in four different models. The four models used for prediction do not include age dummies in the regression. RFOV-1 sets γ =0.75, k=1.5, and β =0.95. RFOV-2 sets γ =0.1, k=2.0, and β =0.95. RFOV-3 sets γ =0.6, k=2.8, and β =0.6.

Figure 3: The Log-Likelihood Surface for Fitting Structural OV Models: Kappa vs. Beta

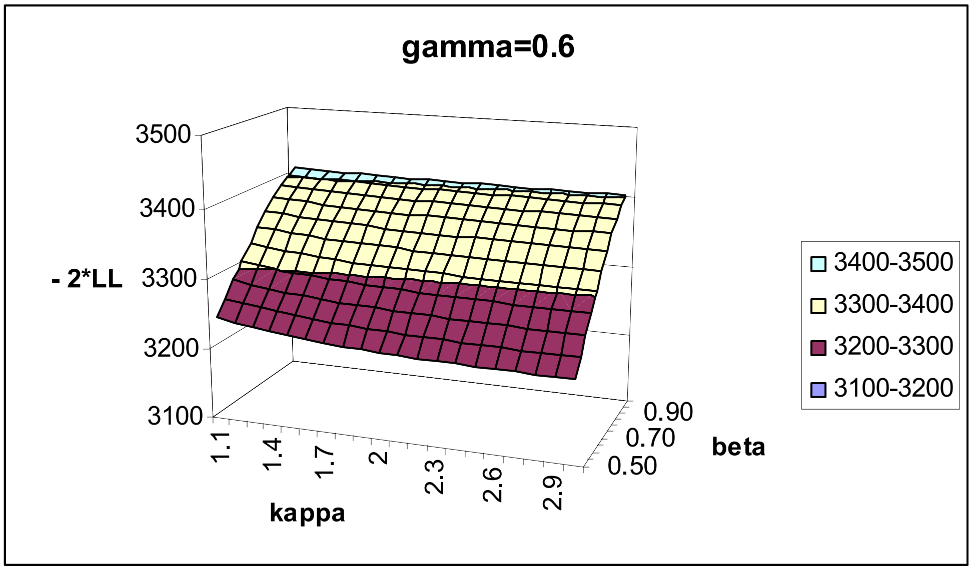

Note: This figure presents the log-likelihood surface (multiplied by –2) for a structural OV model, holding gamma, kappa, and beta fixed and estimating the two scaling parameters (the intercept and the variance). See the text for further details.

Figure 4: The Log-Likelihood Surface for Fitting Structural OV Models: Kappa vs. Gamma

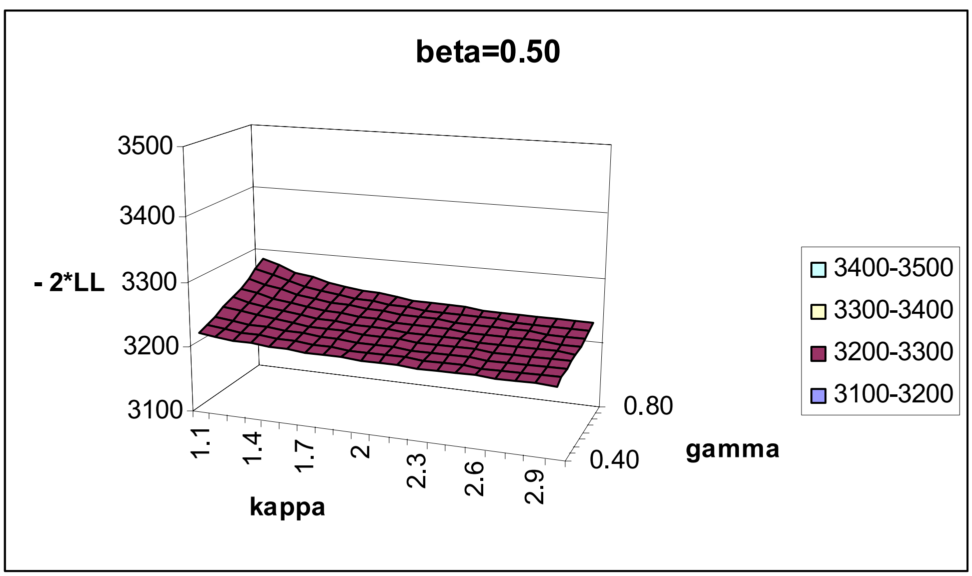

Note: This figure presents the log-likelihood surface (multiplied by –2) for a structural OV model, holding gamma, kappa, and beta fixed and estimating the two scaling parameters (the intercept and the variance). See the text for further details.