Pension Benefits & Retirement Decisions: Income vs. Price Effects[*]

Dayanand Manoli[†]
University of California - Berkeley
(Job Market Paper)

Mathis Wagner
University of Chicago

November 7, 2007

Abstract

How do retirement benefits affect retirement decisions? We separately identify the income and price effects from retirement benefits on retirement decisions. This distinction is important for understanding how social security reforms may affect retirement behavior and economic efficiency. We use administrative data on the universe of employees in Austria and exploit variation in pension benefits created by multiple pension reforms between 1984 and 2003. Based on a proportional hazards specification, we estimate the elasticity of the retirement hazard with respect to pension benefits to be roughly 1.9, with the income effect accounting for 15% of this overall elasticity. We then estimate a dynamic programming model of retirement decisions that allows for wealth accumulation. Using a method of simulated moments (MSM) estimation strategy based on moments related to the hazard models, we estimate the coefficient of relative risk aversion to be roughly .8. This estimate implies that the income effect accounts for 30% of the overall response to benefits. Thus, changes in the schedule of benefits across retirement ages, rather than changes in the levels of benefits across all ages, are critical to avoiding a social security crisis.

---

[*] We are especially grateful to Alan Auerbach, Raj Chetty, Emmanuel Saez and Yuliy Sannikov for their advising. We also thank David Autor, John Friedman, Botond Koszegi, Kathleen Mullen, Sam Schulhofer-Wohl, Jesse Shapiro, Robert Topel, Andrea Weber and participants at the Berkeley Public Finance seminar and lunch and the Chicago Applied Microeconomics Working Group for helpful comments.

[†] Corresponding author. E-mail: dsmanoli@econ.berkeley.edu. Phone: (510) 684-8978. Web: https://webfiles.berkeley.edu/~dsmanoli/.

1 Introduction

Social insurance programs have been designed by governments to provide benefits to individuals at times of large declines in income. Of these programs, social security, which provides benefits during retirement, typically accounts for the largest fraction of government expenditures. Given the vast amount of resources involved in social security programs, retirement has been a widely studied topic in economics with a particular emphasis on understanding how retirement benefits affect retirement decisions. However, the precise mechanisms through which retirement benefits affect retirement decisions remain unclear. In their chapter on social security in the *Handbook of Public Economics*, Feldstein and Liebman (2002) summarize the current state of the literature on retirement benefits and retirement decisions:

> "On balance, it appears to us that when appropriate specifications are used, Social Security systems do appear to have important impacts on retirement behavior. However, significant uncertainty remains about the particular channels provoking these behavioral responses..." (p. 2285)

In this paper, we identify the distinct channels through which retirement benefits affect retirement decisions. Retirement benefits are traditionally thought to affect individuals' behavior through two channels: an income effect and a price effect. The income effect refers to changes in behavior due to changes in lifetime income from the benefits. The price effect refers to changes in behavior due to changes in marginal incentives for continued work from the benefits.

Distinguishing between these channels is important for understanding and designing effective social security reforms. Because of demographic transitions, the ratio of those receiving social security benefits relative to those paying contributions to the systems is increasing dramatically. As a result, there is growing potential for social security crises in multiple countries and hence there is rising pressure for social security reforms. Understanding the income and price effects from retirement benefits is relevant for designing reforms as the income and price effects reveal how much benefit levels and benefit schedules respectively drive retirement decisions. In particular, the magnitude of the price effect relative to the income effect provides insight into how much the schedule of benefits across potential retirement ages drives retirement behavior rather than the increase in individuals' wealth from the receipt of benefits. Additionally, this distinction is also helpful to understand the efficiency consequences of social security programs because the deadweight costs of these programs are related to the size of the price effect. This insight therefore is valuable for understanding how potential social security reforms that change the benefit schedule or the levels of benefits are likely to change retirement behavior.

Our analysis contributes to the literature in three ways. First, we are able to disentangle the income and price effects from retirement benefits on retirement decisions. Second, we build a micro-founded, dynamic model of retirement behavior. Finally, we use administrative data on the universe of employees in Austria coupled with a series of sharp reforms and discontinuities in the Austrian pension system to obtain the most precise estimates to date of the effects of retirement benefits on retirement decisions.

Using data from the Austrian Social Security Database,[1] we present empirical results based on three methodologies: (1) non-parametric graphical evidence; (2) hazard model estimates; and (3) structural estimation. Our results indicate that price effects account for roughly 70% to 85% of the overall labor supply response to changes in retirement benefits. We focus first on the severance payment system in Austria. This system mandates that individuals with qualifying years of tenure at retirement receive a lump-sum payment from their employers at retirement. Each payment is 33.3%, 50%, 75% and 100% of the individual's annual income in the year before retirement if the individual has 10 to 14, 15 to 19, 20 to 24, or more than 25 years of tenure respectively. We exploit the discontinuities created by the tenure thresholds to assess the impacts of the severance payments on retirement behavior. Graphical evidence and hazard model estimates indicate that, while receipt of additional income from the severance payments has no significant impact on the probability of retirement, the price effects from these payments lead to significant delays in retirement.

Next, we use hazard model specifications to relate retirement to pension benefits. We measure the income and price changes from retirement pensions using respectively the expected, present-discounted value of social security wealth and the one-year accrual in this wealth from delaying retirement one year. While previous retirement studies have

---

[1] Our core sample consists of male employees in Austria since we do not observe the number of children for women and maternity leave affects women's pension benefits. We discuss this and other sample restrictions in more detail below. Additionally, our analysis focuses on individual retirement decisions as opposed to household retirement decisions. In Austria, pension benefits are determined at the individual level, *i.e.* independent of marital status. For studies of couples' retirement decisions, see Gustman and Steinmeier (2000) and Coile (2004).

relied on variation in pension benefits created by nonlinearities in social security systems,[2] we exploit variation in pension benefits created by multiple reforms to the Austrian pension system between 1984 and 2003. These reforms create variation in both the level of benefits and the slope of the benefit schedule across ages, thereby allowing for identification of the income and price effects respectively. In our baseline specification, we estimate an overall elasticity of the retirement hazard with respect to pension benefits of roughly 1.9, with price effects accounting for 85% of this overall response to changes in pension benefits.

Finally, we pool the variation in pension benefits and the severance payments to recover the parameters of a dynamic model of retirement decisions. This dynamic model takes into account the pension benefit schedule beyond a one-year horizon and recursive uncertainty relating to job separations. Additionally, while previous dynamic retirement models either lack thorough theoretical foundations (see Coile and Gruber (2000b)) or do not accommodate savings behavior (see Stock and Wise (1990), Berkovec and Stern (1991), and Lumsdaine, Stock and Wise (1992)), our dynamic model is fully micro-founded and allows for wealth accumulation. Accounting for wealth accumulation is important for identifying the wealth effect from retirement benefits in the dynamic setting. We present the theoretical framework allowing for endogenous savings but, due to computation feasibility, assume an exogenous savings rate for estimation. In this model, the coefficient of relative risk aversion, denoted by γ, is identified based on the ratio of wealth and price effects with a lower γ implying smaller wealth effects and larger price effects. Using a Method of Simulated Moments (MSM) estimation strategy based on moments relating to the non-parametric graphical evidence and hazard specifications, we estimate $\gamma = 0.8$. We demonstrate, using hypothetical pension reforms, that this estimate implies that the price effects accounts for roughly 70% of the overall response to changes in retirement benefits.

Our study of income and price effects from retirement benefits relates to several previous studies. Recent research has emphasized the separate identification of income and price effects of benefits from other social insurance programs.[3] We address the relation be-

---

[2] See Coile and Gruber (2000a) and the collection of studies in Gruber and Wise (2004). These studies show that the variation in benefits created by nonlinearities in social security systems is insufficient to separately identify the income and price effects from retirement benefits.

[3] Recent research on unemployment insurance (UI), disability insurance (DI) and health insurance (HI) has emphasized the separate identification of income and price effects from social insurance benefits. Chetty (2007) shows that distinguishing between these two channels in the behavioral response of unemployment durations to unemployment insurance is important for determining the optimal level of UI benefits. Card, Chetty and Weber (2007) use variation in UI benefits to examine individuals' ability to smooth consumption at times of unemployment. Autor and Duggan (2007) and Nyman (2003) examine income and price effects from DI and HI benefits respectively. These studies find that income effects account for the majority of

tween our results and these earlier results based on other social insurance programs in the concluding section as this contrast highlights some directions for future research. Focusing more on the retirement literature, Friedberg (2000) examines income and price effects from retirement benefits using changes in the U.S. social security earnings test. While we focus on the labor force participation decision, this study focuses on hours of work as the outcome variable. Using nonlinear budget sets, Friedberg presents structural estimates of significant uncompensated elasticities with relatively small income elasticities, consistent with our results.

Other retirement studies have concentrated specifically on income effects from retirement benefits. Costa (1995) examines income effects using the introduction of Civil War pensions for Union Army veterans. This study reports an elasticity of nonparticipation with respect to pension income of .73, consistent with our estimate of .53 in our baseline hazard specification. Using reductions in social security wealth following the 1977 Social Security Act in the United States, Krueger and Pischke (1992) find little evidence for statistically significant wealth effects. Additionally, Holtz-Eakin, Joulfaian and Rosen (1999) and Brown, Coile and Weisbenner (2006) estimate income effects on retirement by exploiting variation in inheritance receipt. Our results are generally consistent with estimates from these studies.[4] Samwick (1998), Friedberg (1999), Asch, Haider and Zissimopoulos (2005) and Goda, Shoven, and Slavov (2007) provide additional evidence on the importance of retirement benefits for retirement decisions, though these studies do not focus explicitly on identifying income and price effects.[5]

We present our analysis as follows. The next section presents a Lifetime Budget Constraint Model of retirement decisions. This model allows for a clear illustration of the income and price effects from retirement benefits on retirement decisions. We also use the model to illustrate the effects of severance payments on participation decisions. Following

---

the overall response to changes in the respective benefits.

[4] Holtz-Eakin, Joulfaian and Rosen (1999) find little evidence statistically significant effects of inheritance wealth on retirement. This study reports results from an ordered logit model, so the estimates are not directly comparable to our estimates. Brown, Coile and Weisbenner (2006) present OLS estimates comparable to our OLS estimates in Table 6. Similar to our estimates, they find that a $100,000 increase in inheritance wealth increases the baseline probability of retirement by roughly 3%. The OLS estimates cannot be mapped into elasticities comparable to our hazard model estimates.

[5] There is also a literature examining health benefits and retirement decisions. Rust and Phelan (1997) estimate a dynamic programming model of retirement decisions taking into account incentives from social security and Medicare benefits. See Madrian (2006) for a review of the literature on health insurance and retirement decisions. In contrast to the United States, Austria has a universal health care system for its citizens.

this theoretical section, we describe our data and institutional features of the Austrian pension system in Section 3. Next, we present an empirical analysis of the effects of severance pay on retirement decisions in Section 4. We continue the empirical analysis with an examination of the effects of pension benefits on retirement decisions in Section 5. In section 6, we develop a dynamic programming model of retirement. We present our strategy to estimate this model and discuss the results in Section 7. We conclude our analysis and discuss directions for future research in Section 8.

2 Theoretical Foundations

In this section, we present a lifetime budget constraint (LBC) model of retirement decisions. Burtless (1986), Gustman and Steinmeier (1986) and Brown (2006) provide additional exposition and empirical work within this framework. We use this basic model to illustrate income and price effects in retirement decisions and also to highlight the effects of severance payments on these decisions. Furthermore, this static, non-stochastic framework highlights several intuitions that generalize to more complicated dynamic, stochastic settings. Thus, the LBC model provides a useful foundation for studying retirement.

2.1 A Lifetime Budget Constraint Model

The model focuses on an individual's decision regarding the age at which to retire. We normalize the age at which the decision is evaluated to 0. The individual lives to age T with no uncertainty. Preferences are specified as follows. First, c_t denotes consumption at age t, and $u(c_t)$ denotes the corresponding utility over consumption with the standard assumptions $u'(c_t) > 0$ and $u''(c_t) < 0$. Let v_t denote the disutility of work at age t with $v_t > 0$ and $v_{t+1} > v_t$ for all ages and let R denote the individual's retirement age. With no discounting, the individual's preferences over consumption and work are given by $\int_0^T u(c_t)dt - \int_0^R v_t dt$ which captures that the individual enjoys consumption at all ages but only experiences the disutility of work prior to retirement.

Next, consider the individual's budget constraint. While working, the individual receives an after-tax wage w. After retiring at age R, the individual receives a constant pension benefit $b(R)$ in each subsequent period. The individual's budget constraint is then $\int_0^T c_t dt = wR + (T - R)b(R)$.

The pension system we aim to capture is as follows. First, the individual can only claim his pension benefits after age $\underline{R} > 0$ ($b(R) = b'(R) = 0$ for $R < \underline{R}$). At age $\underline{R}$, the individual is entitled to a constant benefit level $\underline{b}$ in each subsequent period if he claims his pension at that age ($b(\underline{R}) = \underline{b}$). After age $\underline{R}$, the individual can increase his benefit level with continued work ($b'(R) > 0$ for $R > \underline{R}$). At a higher age $\bar{R} > \underline{R}$, benefits no longer increase with continued work, and instead they remain at a constant level $\bar{b}$ ($b(R) = \bar{b}$ and $b'(R) = 0$ for $R \geq \bar{R}$). For simplicity, we assume that benefits increase at a constant rate between ages $\underline{R}$ and $\bar{R}$.

The individual chooses his optimal retirement age according

$$\max_{\{c_t\},R} \int_0^T u(c_t)dt - \int_0^R v_t dt$$

$$s.t. \int_0^T c_t dt = wR + (T-R)b(R).$$

Using the first order conditions,[6] the optimal retirement age is then characterized by $R^* = R$ solving

$$\frac{v_R}{u'(c_R(y))} = w + (T-R)b'(R) - b(R).$$

where $c_R(y)$ refers to consumption at the retirement age R as a function of total income y. This equation illustrates that the optimal retirement age equates the marginal rate of substitution between work and consumption with the price of retirement.[7] The left-hand side of the equation captures the marginal rate of substitution (MRS) between additional work and additional consumption. The right-hand side captures the price of retirement, also referred to as the net-wage (i.e. the individual's wage net of benefits). Specifically, the cost of retiring at a given age, relative to continuing work for an additional year, is that the individual gives up additional wage earnings plus additional retirement benefits at all subsequent ages for retiring at a later age; the benefits the individual would have received are subtracted yielding the net marginal cost of retirement. Notice that using the budget constraint, consumption c_R can be written in terms of total income. Therefore, using y

---

[6] The second order conditions for the utility maximization problem are satisfied since $u''(c) < 0 \; \forall c$ and $v_{t+1} > v_t \; \forall t$.

[7] If the no-discounting assumption is relaxed, a similar condition for the optimal retirement age results with the dollars expressed in present value dollars.

and p to denote total income and the price of retirement,

$$\begin{aligned} y &= wR + (T-R)b(R) \\ p &= w + (T-R)b'(R) - b(R), \end{aligned}$$

we have the optimal retirement age, $R^* = R(y,p)$.

Figure 1 illustrates the determination of the optimal retirement age. Lifetime income (y) and the retirement age (R) are plotted on the vertical and horizontal axes respectively. We first describe the individual's budget constraint. Prior to age $\underline{R}$, the individual does not qualify for pension benefits. As a result, total income is determined by earnings and additional work increases total income through additional wage earnings. At age $\underline{R}$, the individual qualifies for benefit creating a positive jump in lifetime income. Between ages $\underline{R}$ and $\bar{R}$, an additional year of work increases total income from additional wages earnings and additional pension benefits during retirement, but the individual must give up a year of benefits. Beyond age $\bar{R}$, the individual has reached the maximum pension benefits. At these higher ages, additional work increases total income from wage earnings, but the individual gives up the maximal pension level when continuing work. Next, the individual's indifference curve is increasing and convex due to the increasing disutility of work and the concavity in utility over consumption. The optimal retirement age is given by the tangency between the indifference curve and the lifetime budget constraint.

2.2 Pension Reforms: Income & Price Effects

Using the Marshallian (uncompensated) labor supply function $R(y,p)$ from above and the Hicksian (compensated) labor supply function $R(U,p)$ which results from the cost-minimization problem with U as the individual's reservation utility,[8] changes in retirement with respect to benefits can be decomposed into income and price effects according to

$$\varepsilon_{R,b} = \varepsilon^c_{R,p} x_{p,b} + \varepsilon_{R,y} x_{y,b}.$$

---

[8] Specifically, the cost-minimization problem is

$$\min_{\{c_t\},R} \int_0^T c_t dt - wR - (T-R)b(R)$$

$$s.t. \int_0^T u(c_t)dt - \int_0^R v_t dt \geq U$$

Here, $\varepsilon_{R,b}$ represents the uncompensated elasticity of retirement with respect to benefits. The terms $x_{p,b}$ and $x_{y,b}$ denote respectively the percent changes in the net-wage and total income due to a change in benefits. The price effect is captured in the first term on the right-hand side which involves the compensated elasticity with respect to the net-wage, $\varepsilon^c_{R,p}$. Intuitively, increasing the price of retirement while holding utility constant leads the individual to substitute to later retirement ($\varepsilon^c_{R,p} > 0$). The income effect is reflected in the last term on the right-hand side. The intuition behind the income effect is that higher total income leads to decreased marginal utility from additional consumption and hence an earlier retirement ($\varepsilon_{R,y} < 0$).

Our goal is to separately identify the income and price effects in retirement decisions using exogenous changes in pension benefits. To illustrate the income and price effects and to demonstrate the theory behind our identification strategy, we examine a hypothetical reform to the pension system described above. Specifically, we consider a reform that increases the slope of the pension benefit schedule across retirement ages while also decreasing the level of benefits at all potential retirement ages. Figure 2 illustrates the impacts of this pension reform, focusing exclusively on retirement decisions between ages $\underline{R}$ and $\bar{R}$. We focus first on the price effects. Since the reform increases $b'(R)$ in the age range between $\underline{R}$ and $\bar{R}$, the price of retirement increases in this age range. Intuitively, retiring is more costly since the individual is giving up higher benefits from continued work if he retires. As a result of this price change, the individual substitutes to a later optimal retirement age. This change in retirement is a compensated change as the price of retirement increases from the initial pre-reform budget constraint, denoted by BC_0, while holding utility constant at U_0. Next, we focus on the income effects. The decrease in pension benefits at all potential retirement ages decreases the individual's total income (wealth) at retirement. This change in income is reflected by the downward shift to the post-reform budget constraint, BC_1. Since retirement (leisure) is a normal good, the individual consumes less retirement by retiring at a later age. Notice that the degree of curvature in the indifference curve plays a key role in determining the relative magnitude of the price effect in the total labor supply response. In particular, a lower degree of curvature in the indifference curve will imply that the price effects account for a relatively larger fraction of the total labor supply response. We discuss this point in more detail below. Thus, this single pension reform creates variation in both the price of retirement and income at retirement. By considering additional, independent pension reforms that also create variation in both the price of retirement and income at retirement, we have independent variation in these separate channels thereby allowing for identification of the price and income elasticities, $\varepsilon^c_{R,p}$ and $\varepsilon_{R,y}$ respectively.

2.3 Severance Payments

We now introduce severance payments into the LBC model and examine how these payments relate to retirement decisions. Suppose that pension benefits are adjusted to depend additionally on years of tenure so that the individual receives a lump-sum benefit B if he accumulates at least $\underline{\tau}$ years of tenure.

If we assume that a year of tenure is equivalent to a year of age, we can use a simple change of variables to re-cast the model above in terms of an individual choosing his retirement based on the years of tenure he would like to accumulate. More precisely, using τ to denote the individual's choice of tenure, the individual's optimal retirement choice is characterized by $\tau^* = \tau$ solving

$$\frac{v_\tau}{u'(c_\tau)} = w + (T - \tau)b'(\tau) - b(\tau).$$

Though seemingly identical to the equation determining the optimal retirement age, this condition has an important discontinuity at $\underline{\tau}$ years of tenure. Specifically, at $\bar{\tau}$ years of tenure, the slope of the benefit function becomes infinite ($b'(\underline{\tau}) \to \infty$); a small increase in tenure just below the tenure threshold leads to a significant bonus as the severance payment goes from 0 to B.

Figure 3 illustrates the change in the individual's retirement decision due to the severance payment. At $\underline{\tau}$ years of tenure, the budget constraint has a positive jump as total income increases with the severance payment. The individual can then choose to retire either prior to accumulating qualifying tenure or at the tenure threshold. If the individual's disutility from work is relatively low or if marginal utility from consumption does not fall by much with increased consumption from the severance payment (*i.e.* the degree of curvature in the utility over consumption is relatively low), then the individual will be more likely to be better off by delaying retirement to accumulate qualifying tenure for the severance payment.

With these basic theoretical foundations in place, we will now turn to an empirical analysis of pension reforms, severance payments and retirement decisions. We begin the empirical analysis in the next section by discussing features of the data used in our study and of the Austrian pension system.

3 Data & Institutional Background

3.1 Sample Restrictions

We use social-security record data from the Austrian Social Security Database, provided by *Synthesis Forschung*, for all individuals employed in Austria between the years 1972 and 2003. Observations are in the form of spells that are individual-specific, time-specific and insurer-specific. In the cases of employment, the insurer corresponds to the employer, while in the cases of non-employment such as unemployment or disability, the insurer corresponds to the government agency providing assistance. The time-specific characteristic of an observation means that an observation begins either at the beginning of a new spell (a new individual-insurer match) or on the 1st of January of a year. An observation ends either when that particular spell is terminated during a year, or on the 31st of December of a year. The sample covers nearly all Austrians in that period with the exceptions relating to tenured public sector employees and self-employed individuals.[9]

In addition to being characterized by begin dates and end dates, each spell is also characterized by type. The type of spell refers to a more specific classification within the main categories of employment, unemployment, retirement, and maternity leave. For each spell, the amount of earned income during the length of the spell is recorded. Specifically, if the spell corresponds to receiving social insurance, no income is recorded for the spell. Income data is top-coded based on the earnings cap for retirement pension computation.

The data include some variables specific to individuals and insurers. For each individual, the data include gender, birth date, and nationality. For each of the insurers that correspond to employers (these may correspond to firms or plants), the data include region and industry classifications. Using the available data on employees and employers, we construct firm-specific variables such as firm size and tenure.

Our main sample consists of men aged 55 or higher in 2003 (birth cohorts 1948 and earlier). Our sample restrictions and the reasons for these restrictions are as follows. We exclude individuals with less than one year of observed employment time between 1972 and 2003 since these individuals lack sufficient data to compute pension benefits. Next, we exclude foreign nationals as well as those who were have spent more than a year as self-employed or as tenured public servants, farmers, or in mining, construction, and railways since these individuals are covered by separate pension systems. Additionally, we exclude

---

[9] Tenured public sector employees are observed only starting in 1988 or in some cases 1995, and income is not observed for self-employed individuals.

individuals who claim non-disability or non-old-age pensions at the time of retirement since these claims may not correspond to retirement decisions.[10] Lastly, we exclude men claiming disability pensions before age 55 on the basis that these individuals are likely to be permanently disabled. These sample restrictions are summarized in Table 1.

After imposing these restrictions, we arrive at the severance pay sample consisting of 270,946 individuals. This sample, which includes data on all uncensored tenure spells between 1972 and 2003, is used to study the effects of severance pay on retirement decisions. For the pension sample, we focus on individuals between 1984 and 2003 since retirement pensions can only be computed for these years. The pension sample consists of 339,376 individuals. Among this sample, there are 259,517 claimants. The discrepancy between the number of individuals and the number of claimants is explained by the fact that roughly 21.5% of the pension sample is age 60 or younger in 2003 and hence not yet eligible for an old-age pension.

3.2 The Austrian Pension System

The Austrian pension system consists of two primary components: government provided retirement pensions and employer-provided severance payments. A potential third component, private retirement pensions, is virtually non-existent.

We start our description of the pension system with the simpler of the two primary components, the severance payments. Individuals qualify for severance pay from their employers at the time of retirement if the individuals have accumulated sufficient years of uninterrupted tenure by retirement. The amounts of the payments are based on tenure as follows: 10 to 14 years of tenure yields a payment of one third of the last year's salary, 15 to 19 years of tenure yields a payment of one half of the last year's salary, 20 to 24 years of tenure yields a payment of three quarters of the last year's salary and 25 or more years of tenure yields a payment of the full last year's salary. This schedule of payment fractions is shown in Figure 4.

We now turn to the government-provided retirement pensions. We focus on two forms of retirement pensions: disability pensions and old-age pensions. These pensions are computed based on similar rules. Specifically, an individual's pension is the product of two elements. The first element is the assessment basis, which corresponds roughly to the av-

---

[10] The types of pensions claimed are identified in the data. At the time of retirement, other pensions based on, income status, widow status or chronic unemployment may be claimed. We identify men claiming these types of pensions and exclude them from our sample.

erage indexed monthly earnings (AIME) used in social security computation in the United States. The assessment basis refers to the last 15 years of earnings. After applying the earnings caps to earnings in each year, the capped earnings in each year are re-valued based on wage adjustment factors. These revaluation factors are intended to adjust for wage inflation so that existing pensions grow in accordance to wages. After applying the revaluation factors, the capped, revalued earnings are averaged to determine the assessment basis. The second element, the pension coefficient, is then applied to the assessment basis to determine the actual pension level. The pension coefficient corresponds to the percentage of the assessment basis that the individual receives in his pension. This percentage increases to a maximum of 80% based on the number of insurance years and the retirement age. Insurance years correspond to periods of employment as well as periods of unemployment, military service and similar periods of labor market participation. Prior to 2001, disability pensions are computed identically to old-age pensions. In 2001 and after, the pension coefficient used in the disability pension is reduced relative to that of the old-age pension. The reduction in the disability pension coefficient is based on insurance years with lower insurance years receiving larger reductions.

Eligibility for the pensions is as follows. Disability pensions can be claimed at any age, provided that the claimant has been classified as disabled. Generally, an individual is classified as disabled if his working capacity is reduced by more than 50% relative to another individual of similar education. At ages 55, 57 and 65, the conditions to be classified as disabled are relaxed. By claiming a retirement pension, the individual essentially exits the labor market.[11] Men are first eligible for old-age pensions at age 60. In addition to being at least age 60, an individual who claims an old-age pension prior to the statutory retirement age, 65, must have 37.5 insurance years or 35 contribution years (years of contributions to the pension system).

Figure 5 presents average retirement hazard rates by age. In this figure, retirement hazard are based on claiming either an old-age pension or a disability pension. Separate hazard rates into disability pensions are presented in Figure 6. The hazard rates in both figures are based on individuals claiming the respective pensions at any time during the year at the specified age level. Focusing first on Figure 5 with both types of pensions, the

---

[11] In our sample, roughly 9% of individuals continue some work within the year after claiming a pension. About 3.5% of old-age-pension claimants continue work, while 12% of disability claimants continue work. After claiming a pension, there is a mandatory 6 month break required to continue work with the same employer, and additionally, there are minimum earnings restrictions. As a result, we focus on the pension claiming decision as an exit from the labor market.

spikes in the hazard rates occur at ages 60 and 65; these ages correspond to the minimum early retirement age and the statutory retirement age. The hazard rates at these ages are roughly equal at about 80%. The plot in Figure 6 presents details regarding claiming disability pensions. First, the hazard rates are more uniform than the hazard rates into both types of pensions. Since individuals can only retire prior to age 60 via disability pensions, the hazard rates for disability pensions at ages 55 through 59 are much higher than the corresponding hazards based on both types of pensions. The disability hazard rates increase sharply at age 57; at this age the conditions to be qualified as disabled are relaxed allowing for early retirement due to reduced working ability. The hazard rates into disability pensions decline after age 65; after this age, individuals have reached the statutory retirement age and are more likely to qualify for old-age pensions.

Between 1984 and 2003, there were five significant pension reforms in Austria in 1985, 1988, 1993, 1996 and 2000. Our detailed knowledge of these reforms and the computation of the pensions is based on Marek (1985, 1987-2003).[12] Table 2 presents a summary of each reform. The reforms generally reduced the generosity of the retirement pension system as government officials felt the pension system was not financially sustainable. The pension reforms in the 1980s reduced benefits through changes in the length of the assessment basis. The 1985 reform changed the assessment basis from the last 5 years of an individual's earnings to the last 10 years. Because wages are generally increasing with age in Austria, this change decreased benefits. The reform was implemented at the start of the 1985 calendar year. The 1988 reform changed the length of the assessment basis from the last 10 years to the last 15 years. This change was phased in between 1988 and 1992 based on birth cohort. Specifically, the legislation determined the length of an individual's assessment basis based on the year the individual reached age 60.

The reforms in the 1990s continued the reduction in benefits and also specifically aimed to get individuals to retire at later ages. The 1993 reform linked pension coefficients to retirement ages so that the coefficients would rise with both insurance years and retirement ages up to the statutory retirement age, 65. The 1993 reform also changed the assessment basis from the last 15 years of earnings to the highest 15 years of earnings. However, this change generally did not affect retirement pension benefits; since wages generally rise with age, the best 15 years of earnings correspond to the last 15 years of earnings for most

---

[12] Ney (2004) and Linnerooth-Bayer (2001) provide information on the historical contexts of the reforms. See also Koman, Schuh and Weber (2005) and Hofer and Koman (2006) for studies of the Austrian severance pay and pension systems respectively.

individuals. This aspect of the reform is likely to have been more relevant for other non-retirement pensions that are also based on an individual's assessment basis. These changes from the 1993 reform became effective at the start of the 1993 calendar year.

The 1996 and 2000 reforms also focused primarily on changes in pension coefficients. Specifically, the 1996 reform introduced a bonus/malus system to discourage early retirement (before the statutory age) by penalizing early retirees with reduced pension coefficients. The 2000 reforms further developed the bonus/malus system by increasing the reductions in pension coefficients for early retirements and also by offering bonus increases in pension coefficients for retirements after the statutory ages. The 2000 reform also affected eligibility by raising the minimum retirement age from 60 to 61.5. The increase was phased-in between October of 2000 and October of 2002.

4 Severance Pay & Retirement Decisions

Our primary goal is to understand the mechanisms, specifically price and income effects, through which retirement benefits affect retirement decisions. The first step to accomplishing this goal is establishing a causal relationship between benefits and retirement behavior. We use an empirical model motivated by insights from the static LBC model. We first examine the effects of severance pay on retirement decisions and then turn to the effects of pension benefits in the next section.

4.1 Empirical Framework & Identification

Our goal in this section is to determine the causal effect of the severance payments on the probability of retirement at a given age. To do this, we estimate the following hazard model,

$$R_i(a) = \bar{R}(a) \exp\{\beta_1 \ln(SEV_{i,a}) + \beta_2 \ln(\Delta SEV_{i,a}) + \delta X_{i,a}\}.$$

In this specification, $R_i(a)$ denotes the relative hazard for individual i at age a. The relative hazard is the probability that individual i retires at age a conditional on not having retired at an earlier age relative a baseline probability across all individuals at age a. The term $\bar{R}(a)$ denotes the baseline hazard rate at age a. This baseline hazard is common across individuals at each age and thus the intuition regarding the baseline hazard closely follows the intuition of age fixed effects in a linear model. In regard to the severance pay variables, $\ln(SEV_{i,a})$ reflects the impacts of the level of the severance payment and $\ln(\Delta SEV_{i,a})$

reflects the impacts of foregone severance pay. These severance variables therefore roughly capture the wealth and price effects of the severance payments on retirement decisions. The term $X_{i,a}$ refers to covariates for individual i at age a. We include a base and full set of controls. The base controls include education dummies and quartic polynomials in calendar year, log annual earnings and log total earnings from the prior 10 years. The full controls include the base controls as well as industry and region dummies and quartic polynomials in log annual earnings from each of the prior 10 years.

The severance pay variables are constructed as fractions on the individual's previous annual earnings. Specifically, let $sevfrac(t_{i,a})$ denote the fraction of previous annual earnings that is received in as a severance payment when individual i at age a has t years of tenure. These fractions are illustrated in Figure 4. Using these fractions, the severance pay variables are defined as

$$\begin{aligned} SEV_{i,a} &= 1 + sevfrac(t_{i,a}) \\ \Delta SEV_{i,a} &= 1 + [\frac{1}{4}\sum_{k=1}^{4} sevfrac(t_{i,a}+k) - sevfrac(t_{i,a})]. \end{aligned}$$

In particular, notice that the price measure, $\Delta SEV_{i,a}$, is based on foregone severance payments over a 4 year projection assuming that the individual would remain employed with the same employer for the additional time. This degree of forward-looking behavior is motivated by the length of time between the tenure thresholds.[13] The scaling of the severance pay variables as fractions allows the coefficients β_1 and β_2 to be interpreted as the impacts of one percentage point changes in the severance pay fractions. Additionally, the common scale for the wealth and price measures allows the coefficients to reflect comparable wealth and price elasticities as a once percentage point increase in the current severance pay fraction increases the wealth measure by .01 while also decreasing the price measure by .01.

To interpret the estimates of β_1 and β_2 from this specification as causal effects, it is necessary to make an identifying assumption. An assumption regarding the exogeneity of severance pay a given age is essentially an assumption regarding the exogeneity of tenure at a given age. Therefore, interpreting the effects as causal effects requires the assumption that tenure at a given age is exogenous. The intuition behind this identifying assumption can be illustrated with the following example. Suppose there are two possible retirement ages,

---

[13] Specifically, using the current year plus four years in the future covers five years, which is the length of time between the tenure thresholds.

60 and 61. Under the identifying assumption, the individual's tenure upon reaching age 60 is exogenous. Intuitively, the individual may not be able to choose whether he started his current job at age 50 or 51, so he cannot control whether he completes 9 or 10 years of tenure by age 60. Thus, a comparison of those who receive the severance pay with those who do not allows for identification of the causal effect of these payments on retirement. Continuing the example, an individual at age 60 could choose to delay his retirement to age 61 if, at age 60, he were one year away from a qualifying tenure threshold. The price measure captures this delay in retirement due to the tenure threshold, assuming that the individual cannot control the age at which his current tenure spell began.

Heterogeneity across individuals poses a threat to identifying the causal effects of the severance payments. Specifically, individuals with higher tenure may have higher probabilities of retirement due to their preferences regarding work. To address this threat, we include a quartic tenure polynomial to reflect continuous changes in retirement preferences with years of tenure. We also include flexible controls for individuals' earnings histories to capture additional heterogeneity across individuals. Specifically, we include quartic polynomials in annual earnings from each of the previous 10 years, and well as in total earnings from the prior 10 years.

4.2 Graphical Evidence

Before presenting the hazard model estimates for the specification above, we present graphical evidence demonstrating the significance of these severance payments for retirement decisions. More precisely, we present graphical results based on the probability of retirement at a given age by each year of tenure.

Using individuals aged 55 to 70 in the severance pay sample, we estimate a Cox proportional hazard model of the form

$$R_i(a) = \bar{R}(a)\exp\{\beta_1\mathbf{1}(tenure_{i.a}=1)+\beta_2\mathbf{1}(tenure_{i,a}=2)+...+\beta_{32}\mathbf{1}(tenure_{i,a}=32)+\delta X_{i,a}\}.$$

In this case, the baseline hazard corresponds to individuals with less than 1 year of tenure at each age. As above, $X_{i,a}$ corresponds to covariates for individual i at age a. The independent variables of interest are the dummy variables for each tenure level where $1(.)$ denotes the indicator function. We are interested in the coefficients on these dummies to examine changes around the tenure thresholds. These coefficients capture the percentage difference in the hazard between individuals with less than 1 year of tenure and individuals at the specified tenure level.

Figure 7 presents a plot of the coefficients from the above specification with the full controls. Table 3 summarizes the amounts of the severance payments at each tenure level. The graph illustrates discontinuities in the retirement hazard near each of the tenure thresholds. Between the thresholds, there is a downward slope in the relative hazard. The downward slopes between the tenure thresholds provides evidence of forward-looking behavior and the importance of price effects in retirement decisions. Individuals closer to the thresholds have a higher effective wage since additional work increases the likelihood of receiving a (larger) severance payment in addition to providing additional wage income. Thus the downward slopes between thresholds indicate decreases in the probability of retirement as the price of retirement increases. The precise magnitude of the increase in the price of retirement will depend importantly on the discount rate and uncertainty regarding survival and job stability.

Additionally, notice that the hazards are positive at all levels of tenure. This may seem at odds with the theoretical model above which implies that no individual will retire just before a tenure threshold. However, a key element missing from the model that can explain this feature is job uncertainty. Specifically, individuals realistically face separation shocks due to layoffs for example. As a result, individuals who face a risky prospect of qualifying for (additional) severance pay with an additional year of work may chose to retire instead.[14]

4.3 Results

Table 4 presents the results from estimating the above hazard specification based on the severance payment variables. Beginning with the base control results for the full sample, these results indicate a significant price effect and no distinguishable wealth effect from the severance payments. In particular, the coefficient on the wealth measure indicates a lightly positive effect of additional severance pay wealth on the probability of retirement, but the standard error indicates that this effect cannot be distinguished from zero. The coefficient on the price measure indicates that a one percentage point increase in the foregone gain in the severance pay fraction decreases the probability of retirement by roughly .5%. The full sample results with the full controls (*i.e.* including additional earnings history controls) are nearly identical to the base control results.

The results in the remaining columns of Table 4 separately examine individuals at age

---

[14] While a regression framework does not account for this job uncertainty, the separation shocks will be incorporated explicitly in the structural analysis below.

60 as individuals retiring at this age of first eligibility for old-age pensions may be less responsive to the severance pay incentives. Consistent with this intuition, the results for ages $\neq 60$ indicate larger wealth and price effects relative to the full sample results, though the wealth effects remain statistically insignificant. At the point estimates, the estimated coefficients with the full controls imply that the price effects from a one percentage point change in the current severance pay fraction account for roughly 75% of the total response to this change (based on $\frac{.5807}{.5807+.1934}$). Next, we focus explicitly on individuals at age 60. With the tenure and earnings history controls, it becomes difficult to identify statistically significant effects in this smaller sample. Nonetheless, these point estimates are smaller in magnitude than the results in the other columns, indicating that, on average, these individuals are less responsive to the severance pay incentives. Overall, the results in Table 4 confirm the graphical evidence presented in Figure 7 that the severance payments have relatively significant price effects on retirement decisions.

5 Pension Benefits & Retirement Decisions

5.1 Empirical Framework

Using the sample of men at ages 55 through 70 between 1984 and 2003,[15] we estimate a Cox proportional-hazard model of the form,

$$R_i(a) = \bar{R}(a)\exp\{\beta_1 \ln(SSW_{i,a}) + \beta_2 \ln(\Delta_{i,a}) + \delta X_{i,a}\}.$$

The first independent variable, $\ln(SSW_{i,a})$, captures the log of the expected present value of the individual's retirement pension if he were to retire at age a. More precisely,

$$SSW_{i,a} = \sum_{t=a}^{100} \beta^{t-a} \pi_{t|a} b_i(a)$$

where $b_i(a)$ denotes individual i's annual benefit when retiring at age a, $\pi_{t|a}$ denotes the probability of survival until age t conditional on having survived until age a and $\beta = .93$ cap-

---

[15] In estimating this model across ages 55 to 70, we are treating disability pensions as identical to old-age pensions. Specifically, this implies that becoming classified as disabled is a choice by each individual. Furthermore, this assumes that individuals' retirement decisions are based on the best pension available to them. These assumptions are essential to model the decision to claim either pension as a single retirement decision.

tures the individual's discount factor.[16] Each individual's retirement pension is calculated based on the rules of the Austrian pension system and the individual's observed earnings history. The second independent variable, $\ln(\Delta_{i,a})$, captures the individual's expected pension accrual, i.e. the change in $SSW_{i,a}$ from delaying retirement by one additional year. Specifically, $\Delta_{i,a}$ is defined in terms of present value social security wealth according to

$$\Delta_{ia} = 1 + \underbrace{\tfrac{E_a(SSW_{i,a+1}) - SSW_{i,a}}{SSW_{i,a}}}_{ACC_{i,a}}$$

where $ACC_{i,a} = \frac{E_a(SSW_{i,a+1}) - SSW_{i,a}}{SSW_{i,a}}$ captures the 1-year accrual in pension wealth net of pension contributions. We assume 1.5% real wage growth to project earnings one year ahead. Lastly, individual covariates are captured in the vector X_{ia}. We discuss these covariates in more detail below.[17]

In estimating this model across ages 55 to 70, we are treating disability pensions as identical to old-age pensions. Specifically, this implies that becoming classified as disabled is a choice by each individual. Furthermore, this assumes that individuals' retirement decisions are based on the best pension available to them. These assumptions are essential to model the decision to claim either pension as a single retirement decision. We address this assumption in more detail in our dynamic model.

This empirical model is based on previous work in the literature. Lumsdaine, Stock and Wise (1992), Coile and Gruber (2000a, b), Gruber and Wise (2004) and others have focused primarily on probit and linear probability models relating pension incentives and retirement decisions. We focus on a hazard model to adopt a more dynamic perspective on each retirement decision as a stopping-time event following a duration of a career. Furthermore, the hazard model presents results precisely in terms of the elasticities we

---

[16] The survival probabilities are taken from life tables available through Statistics Austria (www.statistik.at). The value of β corresponds to a real interest rate of roughly 7.5% which is consistent with the long-term real interest rate in Austria in the mid-1990s.

[17] The hazard specification can also be written in terms of current and foregone pension wealth

$$R_i(a) = \bar{R}(a) \exp\{b_1 \ln(SSW_{i,a}) + b_2 \ln(E_a(SSW_{i,a+1})) + \delta X_{i,a}\}.$$

where $b_1 = \beta_1 - \beta_2$ and $b_2 = \beta_2$. While the ratio $\frac{\beta_2}{\beta_1}$ illustrates the ratio of the price effect to the income effect, the ratio $\frac{b_2}{b_1}$ illustrates the ratio of the price effect to the sum of the income and price effect. This difference arises because the second specification does not separate the income and price effects of changes in pension wealth at age a whereas the first specification does. In both specifications, the coefficients reflect comparable elasticities as a 1% change in pension wealth at age a changes the wealth measure by 1% and also the price measure by 1%.

are interested in, whereas the alternative models present coefficients that cannot be easily converted into elasticities.

The coefficients β_1 and β_2 relate to the income and price effects discussed in the LBC model. First, β_1 captures the elasticity of retirement with respect to pension wealth. Second, β_2 captures the elasticity of retirement with respect to the one-year accrual. However, in this empirical model, notice that we have assumed that the price of retirement is based only on looking ahead one year as reflected in the accrual measure. To the extent that individuals are more forward-looking, the price of retirement should take into account the profile of benefits are multiple future ages. Thus β_1 and β_2 will correspond to the income and price elasticities conditional on assuming this limited forward-looking behavior.

The Cox model makes the particular assumption that the effects of benefits on retirement decisions are proportional across ages. Consider the following example to illustrate the implications of this proportionality assumption. Suppose that the vector X_{ia} includes years of experience. The proportionality assumption implies that a one year increase in years of experience has the same proportional effect at age 55 as it does at age 70. Given the rules of the Austrian pension system, additional experience at age 55 changes benefits differently than additional experience at age 70. This leads to a plausible violation of the proportionality assumption. Therefore, it will be particularly important to relax this assumption and examine effects at different ages.

5.2 Identification

We exploit exogenous variation in retirement benefits created by the 5 pension reforms in Austria between 1984 and 2003 to identify a causal relationship between benefits and retirement decisions. More precisely, we use only this specific variation in pension benefits created by the reforms to identify β_1 and β_2, corresponding to the pension wealth and accrual elasticities in the hazard model above. Without this exogenous variation from, the reforms, the identification of causal effects is threatened by unobserved heterogeneity in preferences for work. Intuitively, individuals with more willingness to work may have higher earnings and hence higher pension benefits, thereby creating a correlation between benefits and retirement decisions. We therefore include sets of polynomials in individuals' earnings histories to control for variation in pension benefits based on earnings histories. Additionally, the baseline hazard controls for changes in the pension benefit schedule that are common across ages. Thus, only the remaining variation in pension benefits due to the pension reforms is used to identify the pension wealth and accrual elasticities.

The key to identifying both the income and price effects is that the pension reforms create exogenous variation in pension wealth and the accrual that is independent across the reforms. In particular, if there are two pension reforms that change pension wealth and the accrual by the same amounts, it is not possible to separately identify the effects of pension wealth and the accrual on retirement decisions. In our case, each of the pension reforms creates separate, independent variation in pension wealth and the accrual thereby allowing us to identify the parameters of interest.

5.3 Graphical Evidence

Figure 8 illustrates the identification strategy based on the pension reforms. This figure presents three time-series for individuals at age 55. The first time-series is the mean accrual at age 55. The second time-series is the median change in social security wealth, where changes are computed relative to the previous year's legislation. An increase in the first time-series reflects an increase in the price of retirement while a negative value for the second time series reflects a decrease in pension wealth at retirement. The final time-series is the retirement hazard for individuals at age 55. This figure concentrates on individuals at age 55 since the current discussion will be based on two particular pension reforms in 1988 and 1996 that first affect individuals at age 55.

We consider first the identification of income effects from pension benefits on retirement decisions based on the changes in pension wealth. The 1988 pension reform creates variation in pension wealth since the reform phased in a five-year increase in the length of the assessment basis from the last 10 years to the last 15 years of earnings. Since earnings further back in the earnings history are generally lower (*i.e.* earnings are generally increasing with age), this increase in the length of the assessment basis lowers pension wealth. As illustrated in Figure 8, median pension wealth decreases by roughly 1% with each additional year for the assessment basis. Notice that this reform only affects the level of pension wealth as the accrual is unchanged. Focusing on the retirement responses, the retirement hazard time-series has only a slight decrease at the time of the reform, and this decrease does not persist over the entire phase-in. The lack of distinct changes in the retirement hazard indicate that the wealth effects from pension benefits are likely to be relatively small.

Next, we consider the 1996 pension reform which creates both income and price effects of pension benefits on retirement decisions. This reform increases the penalties for early retirement (retirement before the statutory age, 65). As a result of these penalties, the mean accrual increases between 1995 from roughly -.095 to -.08, reflecting a higher price of retirement. Additionally, the penalties for early retirement reduce pension wealth. Relative to the pre-reform legislation, pension wealth decreases by roughly 3.5% after the reform. While the 1988 pension reform indicates that wealth effects are likely to be relatively small, the 1996 reform indicates that the price effects are likely to be relatively large. Specifically, with this reform that includes price changes in addition to the wealth changes, the hazard falls sharply at the time of the reform from roughly .15 to .05.

As mentioned above, the key to our identification strategy is that the pension reforms create exogenous variation in pension wealth and the accrual that is independent across the reforms. In particular, notice that it is not essential that one pension reform affects only pension wealth while another reform affects both pension wealth and the accrual. This example is simply a special case of independent variation in pension wealth and the accrual across two pension reforms. In the regression analysis below, we pool the exogenous variation in pension wealth and the accrual across the five reforms and across multiple retirement ages to precisely identify the income and price effects from pension benefits on retirement decisions.

5.4 Results

The results from estimating the Cox model are presented in Table 5. The estimates from the full sample with the base controls indicate that a 1% increase in pension wealth increases the hazard by .5%, while a 1% increase in the accrual measure decreases the hazard by 1%. After including the full control set, the estimate for pension wealth decreases slightly to .33 while the estimate for the accrual increases in magnitude to roughly 1.6. Consistent with the earlier graphical evidence, these effects highlight larger price effects and smaller income effects.

In the remaining columns of Table 5, we split the sample by age to examine heterogeneity across retirement ages and also to relax the proportionality assumption of the Cox model. Specifically, we examine the impacts of the spike in the hazard rate at age 60 by omitting observations at this particular age. Focusing on the results for ages $\neq 60$, the coefficients on the pension wealth variable increase slightly from .53 and .33 to .68 and .48 in the base and full control specifications respectively. The accrual coefficients increase in magnitude more significantly. In particular, comparing the full control specifications, the accrual coefficient changes from -1.6 to roughly -2.2. This suggests that individuals retiring at age 60 may be insensitive to financial incentives relative to individuals considering retirement at all other potential retirement ages. Indeed, the estimates for the sample at age 60 are generally smaller in magnitude than the estimates for the full sample. While the estimates do differ across years of age, the stability of the estimates indicates that proportionality assumption does not seem violated. Lastly, we examine individuals retiring at ages beyond the disability ages and the minimum retirement age. The estimates for this sample indicate generally larger elasticities for these individuals. Furthermore, the pension wealth coefficients are larger than the corresponding coefficient for the ages $\neq 60$ sample. This indicates that pension wealth has larger effects on retirement decisions for individuals retiring at older than for individuals retiring at younger age. Thus, while responsiveness to pension incentives seems to increase with age, the share of the overall elasticity due to changes in pension wealth also seems to increase with age.

For comparison with results from previous studies and for robustness, we present estimates from a linear probability model with age fixed effects in Table 6.[18] Focusing on the results from the full sample with full controls, the estimated coefficients for pension wealth and the accrual are roughly .06 and -1.5. In the right columns, we estimate separate coefficients for the different age groups as in the hazard model. These estimates confirm the patterns and magnitudes presented in Table 5. First, the coefficients at age 60 are generally smaller than the corresponding coefficients for the remaining ages. Second, the share of the responsiveness that is attributed to changes in pension wealth increases with age. However, these coefficients cannot be directly interpreted as elasticities as in the hazard model. Since the linear model is additive rather than multiplicative like the Cox model, the estimates are not presented in terms of percentage changes in the dependent variable.

6 A Dynamic Framework

In this section, we develop a dynamic model of retirement decisions. While the LBC model provides a clear presentation of the intuition behind income and price effects in retirement decisions, this model is poorly suited to capture realistic uncertainty that affects retirement decisions. In particular, we are interested in modelling uncertainty relating to

---

[18] In the linear probability model, we also include a quartic polynomial in insurance months divided by 10. This additional control cannot be included in the hazard model since age is closely correlated with insurance months, and hence insurance months are accounted for in the baseline hazard function.

employment status.[19] Intuitively, an individual may become laid off so that, upon learning of his job separation, he chooses to retire earlier than previously expected. Accounting for such separations is an important component of modelling uncertainty in severance pay. Additionally, an individual's retirement decision may be driven by the schedule of benefits across multiple ages rather than only considering benefits in the next year. The dynamic model accommodates for this forward-looking behavior and allows for recursive uncertainty relating to job separations.

The dynamic model that we develop is closely related to previous work in the literature. In particular, Stock and Wise (1990) present a dynamic model of retirement decisions that emphasizes the role of the option value of work in these decisions. In addition to discussing the option value model, Lumsdaine, Stock and Wise (1992) present a dynamic programming model of retirement decisions. Berkovec and Stern (1991) present a dynamic model to examine movement between full-time work, part-time work, and retirement. While these models have assumed that consumption is equal to current income, we allow for wealth accumulation that is important for our purposes. Additionally, Coile and Gruber (2000a) study the dynamic nature of retirement decisions in a regression framework. Their analysis is based on a peak-value measure that corresponds at an intuitive level to the option value measure examined by Stock and Wise. While the peak-value measure is not based on a particular theoretical framework, we develop a dynamic model with explicit theoretical foundations.

We present the theoretical framework allowing for endogenous savings. However, when estimating this model, we assume a fixed savings rate for computational feasibility. We discuss this in more detail in the estimation section below.

6.1 The Model

The intuition behind the model is as follows. In each period, an employed individual must chose whether to retire or whether to continue working. A period in the model corresponds to an individual's age. At the beginning of the period, the individual knows his assets and he learns his labor market status. Specifically, labor market status refers to the individual's wage, his tenure, whether or not the individual was separated from his previous job, and other observables affecting his retirement benefits and labor market income. If he chooses to retire, the individual receives his retirement benefits and faces no

---

[19] Jimenez-Martin and Sanchez-Martin (2007) present a structural estimation of an LBC model with uncertainty regarding the time of death.

remaining uncertainty from the labor market. Based on his assets and retirement benefits, the retired individual chooses his optimal consumption. On the other hand, if an individual chooses to continue working, he receives his income from the labor market and chooses his optimal consumption. When optimizing, the working individual takes into account the continued uncertainty from the labor market. The individual repeats this decision problem at each age in the labor market.

We now turn to formalizing the model. Consider first the optimization problem for an individual who has chosen to retire. Let $R_a(A_a, X_a)$ denote the value of retirement at age a for an individual with assets A_a and labor market observables X_a, where the subscript a refers to the individual's age. Once an individual has chosen to retire, the individual solves the following optimization problem:

$$\begin{aligned} R_a(A_a, X_a) &= \max_{c^R} u(c^R) + \beta\pi_{a+1|a} R(A_{a+1}, X_{a+1}) \\ s.t.\ \ A_{a+1} &= (1+r)[A_a + y^R(X_a) - c^R] \end{aligned}$$

The function $u(.)$ captures utility over retirement consumption c^R with $u' > 0$ and $u'' < 0$. The term $y^R(X_a)$ denotes the individual's retirement income based on his observables X_a. This retirement income consists of pension benefits and severance pay. Pension benefits are computed according to the rules of the pension system. Thus, X_a contains the relevant variables such as age, insurance years and the individual's earnings history as well as projections these variables and pension rules to compute expected benefits at future dates. If the individual qualifies for a severance payment, the individual receives the entire, lump-sum payment at the age of his retirement. Therefore, X_a also includes the individual's separation status and years of tenure in the current and prior period. The variables r, β and $\pi_{a+1|a}$ respectively denote the interest rate, the discount factor and the probability of survival to age $a+1$ conditional on survival to age a.

Next, consider the problem facing an individual who has chosen to work. As in the case of retirement, the individual must chose his optimal consumption. The optimization problem in the case of continuing to work differs from that in the case of retirement in the following respects. First, the individual must take into account his preference for work. While the individual's work preference evolves across periods, each individual knows his work-preference profile at the initial age. Specifically, at the initial age, v_0 is drawn for each individual from the distribution $\Psi(v)$ defined over $[0, \infty)$ with $\Psi(0) = 0$ and $\lim_{v\to\infty} \Psi(v) = 1$. At subsequent ages, the individual's work preference is given by $v_a(v_0)$

such that work utility is decreasing and concave in age. Second, the individual's income is based on his wage income and any severance pay if he was separated from his previous job and had qualifying tenure. After-tax income from working is denoted by $y^W(X_a)$. Third, the individual must take into account the continued uncertainty from the labor market. In particular, $E_a[.]$ captures the individual's expectation at age a taking into account uncertainty regarding future labor market status (i.e. separation status). Let $W_a(A_a, v_a, X_a)$ denote the value of working at age a with assets A_a, work utility v_a and observables X_a. The individual's optimization problem conditional on continuing to work is given by[20]

$$\begin{aligned} W_a(A_a, v_a, X_a) &= \max_{c^W} u(c^W, v_a) + \beta\pi_{a+1|a} E_a[D_{a+1}(A_{a+1}, v_{a+1}, X_{a+1})] \\ s.t.\ \ A_{a+1} &= (1+r)[A_a + y^W(X_a) - c^W] \end{aligned}$$

The function $D_a(A_a, v_a, X_a)$ captures the value of being in the labor market at age a and having the decision between retiring or continuing to work. When deciding between retirement and work, the individual simply chooses the option that presents the highest value,

$$D_a(A_a, v_a, X_a) = \max_{retire, work} \{R_a(A_a, X_a), W_a(A_a, v_a, X_a)\}.$$

In regard to heterogeneity, retirement income $y^R(X_a)$, labor market income $y^W(X_a)$, and work utility v_a are allowed to vary across individuals. The interest rate r, discount rate β, survival probabilities $\pi_{a+1|a}$, and consumption-utility function $u(.)$ are restricted to be common across individuals. Additionally, the utility over consumption when retired is assumed to be the same as the utility over consumption when employed.

In this optimal stopping time setting, the individual's optimal strategy is to set a reservation level for his work preference such that he will retire if his utility from work is smaller than his reservation level. Let $\bar{v}_a$ denote the individual's reservation level at age a. Formally, the reservation level is defined as the preference for work that leaves the

---

[20] We model the utility from work as being multiplicative, rather than additively separable, with the utility over consumption. This is consistent with evidence presented by Aguiar and Hurst (2005). Specifically, these authors establish that while food expenditures decline at retirement, total food consumption at retirement, measured in terms of quantity and quality, remains stable. The change in expenditure patterns at retirement suggests that the marginal utility over expenditures changes with the additional leisure at retirement. This feature is captured in the multiplicative form for within-period utility.

individual indifferent between retiring and continuing to work

$$R_a(A_a, X_a) = W_a(A_a, \bar{v}_a, X_a) \Rightarrow \bar{v}_a = \bar{v}_a(A_a, X_a).$$

Given this reservation level, the individual's retirement decision rule is then

$$\text{Retire at age } a \text{ if } v_a \leq \bar{v}_a(A_a, X_a).$$

To compute the value functions used to determine the reservation work utility, we assume that there is a terminal age a^* at which all working individuals are assumed to retire. With this assumption, the value functions can be computed recursively.

Figure 9 illustrates an example with the computed value functions and also presents the determination of the individual's reservation work utility.[21] This figure plots utility on the vertical axis and work utility on the horizontal axis. Since the value of retirement is independent from work utility, the value function R_a is a horizontal line. In this example, the marginal utility of work at a given age is diminishing; the value of continued work increases more at lower values of work utility than at higher values. The reservation work utility is then determined by the intersection of these two value functions.

6.2 Wealth and Price Effects

Similar to the static LBC model, retirement benefits affect retirement decisions through income (wealth) and price effects in the dynamic model. To develop intuition for wealth and price effects in the dynamic setting, we will first consider a change in retirement benefits at a given age. Following this discussion, we will build upon this intuition to consider a change in the profile of retirement benefits across multiple ages as this more closely corresponds

---

[21] The functional forms used to compute this example are the same as the functional forms described in the next section on estimation. Specifically, we use:

$$\begin{aligned} u(c) &= \frac{c^{1-\gamma}}{1-\gamma} \\ u(c, v_a) &= \frac{[cv_a]^{1-\gamma}}{1-\gamma} \\ v_a &= \frac{(a-54)^{-\eta_2}}{\eta_2}. \end{aligned}$$

The parameter values are $\beta = .95$ and $\gamma = \eta_2 = 1.25$. Furthermore, we consider an individual at age 65 with $A = 10,000$, $y^W = 8,000$, tenure $= 8$, benefits $y_{65}^R = 7,800$ and expected benefits $Ey_{66}^R = 8,100$.

to actual pension reforms.

The income and price effects from retirement benefits can be seen by examining changes in an individual's optimal retirement strategy. Following Chetty (2007) and Card, Chetty and Weber (2007) who study wealth and price effects in the context of unemployment insurance, consider separately the effects of changes in benefits, earnings and savings at a given age a. Differentiating the equation for $\bar{v}_a$ with respect to $y^R(X_a)$, $y^W(X_a)$ and A_a yields

$$\begin{aligned}
[y^R(X_a)] &: \frac{du(c_a^R)}{dc_a^R} = \frac{du(c_a^W,\bar{v}_a)}{dv_a}\frac{d\bar{v}_a}{dy^R(X_a)} \\
[y^W(X_a)] &: 0 = \frac{du(c_a^W,\bar{v}_a)}{dc_a^W} + \frac{du(c_a^W,\bar{v}_a)}{dv_a}\frac{d\bar{v}_a}{dy^W(X_a)} \\
[A_a] &: \frac{du(c_a^R)}{dc_a^R} = \frac{du(c_a^W,\bar{v}_a)}{dc_a^W} + \frac{du(c_a^W,\bar{v}_a)}{dv_a}\frac{d\bar{v}_a}{dA_a}
\end{aligned}$$

The first condition reflects that a small increase in retirement benefits increases the value of retirement through increased current consumption during retirement. This change leads the individual to adjust his reservation level so he is more likely to retire. Notice that all of the changes in the reservation utility are scaled by the flow marginal utility from work, $\frac{du(c_a^W,\bar{v}_a)}{dv_a}$. The second condition illustrates that a small increase in labor market income increases the value of continuing to work though increased consumption while working. This effect leads the individual to adjust his strategy so he is more likely to continue working. The last condition illustrates that a change in assets affects both the value of retirement and the value of continuing work. The individual adjusts his reservation level to exactly offset the differences in marginal utility so as to remain indifferent between his options.

The third condition illustrates that the effects of a change in savings must be balanced between the two options. The first two conditions allow the effects from the change in savings at age a to be related to the effects from changes in benefits and earnings at age a. We can therefore combine these conditions to write the following condition capturing the response to a change in retirement benefits at age a

$$\frac{d\bar{v}_a}{dy^R(X_a)} = \underbrace{\frac{d\bar{v}_a}{dA_a}}_{\text{wealth effect}} - \underbrace{\frac{d\bar{v}_a}{dy^W(X_a)}}_{\text{price effect}}.$$

This condition illustrates that the response to a change in the profile of retirement benefits can be decomposed into an wealth effect and a price effect. The income effect, $\frac{d\bar{v}_a}{dA_a}$, reflects that the change in benefits creates additional wealth for the individual. The price effect, $\frac{d\bar{v}_a}{dy^W(X_a)}$, reflects that the change in retirement benefits creates a change in the effective wage from working. This price effect is associated with the deadweight costs created from distorting the individual's marginal incentives with the change in benefits.

Before moving on, we consider the effects of additional savings at a given age in more detail. As discussed by Lentz and Tranaes (2005) in the context of search models, ambiguity in the sign of the wealth effect arises with multiplicative within-period utility due to the complementarity between consumption and the preference for work. Intuitively, a decrease in wealth affects the probability of retirement through two channels. First, the decrease in wealth makes the gains from continued work seem relatively more attractive. As a result, the probability of retirement decreases ($\bar{v}_a$ decreases). Second, the marginal utility of work decreases since consumption decreases and the two are complementary given the multiplicative form of utility. As a result, the probability of retirement increases ($\bar{v}_a$ increases). The first effect dominates at higher asset levels when $\bar{v}_a$ and c_a^W are higher, while the second effect will dominate at lower asset levels when $\bar{v}_a$ and c_a^W are lower since the marginal utility of work changes more at lower asset levels. Thus, $\frac{d\bar{v}_a}{dA_a} > 0$ at higher asset levels while $\frac{d\bar{v}_a}{dA_a} \leq 0$ at lower asset levels.

Figures 10A-C provide a computed example that illustrates the decomposition of the labor supply response from a change in retirement benefits into a wealth effect and a price effect.[22] First, Figure 10A illustrates the behavioral response to a decrease in retirement benefits at the current age. This decrease in benefits reduces the value of retirement; the value of continued work is unchanged since benefits at future ages are unchanged. Given the decrease in R_a, the reservation work utility decreases so that continued labor force participation is more likely.

We begin the decomposition of this behavioral response by considering a change in wealth. The effects from a change in wealth are illustrated in Figure 10B. Unlike the change in current retirement benefits considered in Figure 10A, the change in wealth affects both the value of retirement and the value of continued work. Lower wealth reduced the value of retirement since the individual will have less to consume while retired. Additionally, the

---

[22] The functional forms used in this example are the same as those described in the footnote for Figure 10 and in the next section on estimation. We use the following parameter values: $\beta = .95$, $\gamma = 1.5$ and $\eta_2 = .75$. For the baseline, we use an individual at age 65 with $A = 75,000$, $y^W = 20,000$, tenure = 8, benefits $y_{65}^R = 16,600$ and expected benefits $Ey_{66}^R = 16,700$. The reduced benefit level for Figure 11A is $y_{65}^R = 15,000$. The reduced wealth level for Figure 11B is $A = 3750$. The increased wage for Figure 11C is $y^W = 35,000$.

reduction in wealth also lowers the value of continued work since consumption while working declines. As a result of these changes, the reservation utility decreases. Intuitively, the reduction in wealth makes the gains from continued work seem relatively more attractive. Therefore, continued labor force participation is more likely. As shown in the figures, this wealth effect partially accounts for the behavioral response to the reduction in benefits.

They key parameter governing the magnitude of the wealth effect is the curvature of the value function W_a which is based on the curvature of the utility function with respect to consumption. A higher degree of curvature will imply a larger wealth effect. When the utility function is very curved, marginal utility rises rapidly as consumption decreases. As a result, a decrease in wealth that reduces consumption leads to a sharp increase in marginal utility of consumption and hence the gains in consumption from additional work seem relatively more attractive. This leads to a significant downward revision in the reservation work utility so continued work is much more likely. This can be seen in Figure 10B by noticing that when utility is very curved, the value function W_a will be relatively flat at higher values of work utility. As a result, the shift in the reservation utility due to the downward change in R_a will lead to a relatively large decrease in $\bar{v}_a$. Thus, the decomposition of the behavioral response to a reduction in pension benefits into wealth and price effects will be determined by the degree of curvature in consumption utility.

Figure 10C presents the price effect that accounts for the remaining amount of the behavioral response to the benefit reduction. In this figure, we consider a wage increase given that assets have been reduced as in Figure 10B. The wage increase leads to an increase in the value of continued work and leaves the value of retirement unchanged. Thus, the reservation utility declines so that continued work is more likely. Combining Figures 10B & C, we see that the responses to the changes in wealth and wages account for the response to the change in benefits shown in Figure 10A.

Thus far, we have focused on a change in benefits at a given age to develop intuition for wealth and price effects from retirement benefits. However, the pension reforms discussed above and often considered by policymakers correspond to changes in the schedule of benefits across multiple retirement ages. Nonetheless, the intuition from considering changes at a given age can be generalized to understand the effects from a reform that changes the profile of retirement benefits. Such a reform affects the current and future values of retirement and continued work. As we have shown above, changes in current wealth and the current wage can account for changes in the current values of retirement and continued work. However, we need to consider changes in wealth and wages at future ages to account for the changes in the value functions at future ages.

To develop intuition for the effects of a change in the profile of retirement benefits, we consider a setting in which benefits and wages increase linearly with age. In particular, benefits and wages are given by

$$\begin{aligned} y_a^R &= \bar{b} + a\Delta^R \\ y_a^W &= \bar{y} + a\Delta^W. \end{aligned}$$

In this case, consider a reform that alters the profile of benefits by changing the intercept $\bar{b}$ and the slope Δ^R. To relate the behavioral response to this reform to a wealth effect and a price effect, it is important to account for timing in the changes in wealth. Intuitively, a change in wealth at an earlier age affects the optimal path of savings and hence alters wealth at later ages as well. Following Chetty (2007), we introduce two annuities, one that pays the individual a constant amount $\bar{A}$ at each age and another that pays the individual different amounts at each age but with a constant slope Δ^A across ages, and then consider comparative statics with respect to these annuities to decompose the behavioral response to the change in benefits. Specifically, we show in Appendix 1 that the behavioral response to the reform can be decomposed into wealth and price effects,

$$\begin{aligned} \frac{d\bar{v}_a}{d\bar{b}} &= \frac{d\bar{v}_a}{d\bar{A}} - \frac{d\bar{v}_a}{d\bar{y}} \\ \frac{d\bar{v}_a}{d\Delta^R} &= \frac{d\bar{v}_a}{d\Delta^A} - \frac{d\bar{v}_a}{d\Delta^W}. \end{aligned}$$

The central point from this exercise is that, following the intuition from considering changes at a given age, a change in the benefit profile can be related to a change in the wealth profile (wealth effect) and a change in the wage profile (price effect).

Having developed some intuition behind an individual's retirement decision within the context of the dynamic framework, we now turn to recovering the parameters of the model. Once the parameters have been identified, it is possible to determine the magnitudes of the income and price effects.

7 Estimation of the Dynamic Model

In this section, we focus on using the data on severance payments, pension benefits and retirement decisions to estimate the parameters of the dynamic model. We start by discussing the estimation strategy. Following this discussion, we present the results from implementing the strategy. Lastly, we discuss some additional estimation strategies to be considered in the future.

7.1 Estimation Strategy

Our baseline estimation strategy is based on matching key moments observed in the data. The intuition behind the strategy is as follows. First, suppose that M key moments have been identified from the data. Let m_i denote the ith moment with $i = 1...M$. Next, consider the theoretical model. Given a set of parameter values and data on assets, wages, tenure, and pension benefits, we simulate retirement outcomes for each individual. Using these outcomes, the analogues to the identified moments from the data can be constructed. Using θ to denote the vector of parameters and X to refer to the data, the ith simulated moment is given by $\hat{m}_i(\theta, X)$. The estimation strategy is then to find the parameters that minimize the distance between the identified empirical moments and the simulated moments. Specifically, the parameters are estimated according to[23]

$$\hat{\theta} = \arg\min_{\theta} \sum_{i=1}^{M} [m_i - \hat{m}_i(\theta, X)]^2.$$

For our estimation, we use the following moments: the retirement hazard rates by years of age for ages 55 through 65 and the probabilities of retirement by years of tenure for tenure 0 through 30. Therefore, the initial age for retirement is age 55.

This generalized method of moments (MSM) estimation strategy presents some advantages over an alternative, maximum-likelihood (ML) estimation strategy. General intuition of ML is sufficient for the purposes of the present discussion, but we discuss an ML estimation strategy in more detail in the appendix. A first advantage of the MSM estimator lies in transparency. In particular, the MSM estimator clearly illustrates the variation in the data the is used to estimate the parameters. By selecting the moments based on economic theory, the MSM estimator ensures that the economically relevant variation in the data will

---

[23] It is possible to introduce a weighting matrix to improve the estimation. Using Ω to denote an M by M weighting matrix, the estimator in this case is given by $\hat{\theta} = \arg\min_{\theta} m(\theta, X) W m(\theta, X)'$, where $m(\theta, X)$ denotes a 1 by M vector with the ith position corresponding to the ith moment condition $m_i - \hat{m}_i(\theta, X)$. Intuitively, the optimal weighting matrix places more weight on the moments with lower variances since these moments provide the most information for the parameter values (i.e. $W = \Omega^{-1}$ where Ω is the asymptotic covariance matrix of the GMM estimator). In our case, we weight each moment uniformly (i.e. W is an M by M identity matrix). In this case, our estimator will be consistent but inefficient.

drive the estimation. An ML estimator, on the other hand, may use additional variation in the data and additional information from less well-founded functional form (distributional) assumptions to estimate the parameters. If the estimation employs or relies on economically irrelevant variation, this may be problematic.

A second advantage of the MSM estimator lies in computation since an ML estimator relies on a full solution of the model for each individual while the MSM estimator makes use of a partial solution. In particular, an ML estimator is based on the probability that an individual retires at each age the individual is in the sample and hence it requires identifying the reservation utility for each individual at multiple ages. In contrast, the MSM estimator is based only on whether the individual retires or not. As a result, the MSM estimator does not require identifying a reservation utility at any age for any individual. Since each reservation utility must be determined via grid search (due to the intertemporal optimization, it is not possible to obtain an analytic solution), an ML estimator will require more computation time.

Implementing the MSM estimation strategy requires additional assumptions relating to functional forms and expectations. We assume that flow utility over consumption and work at age a is given by

$$u(c^W, v_a) = \frac{[c^W v_a]^{1-\gamma}}{1-\gamma} \text{ with } \gamma > 0.$$

Next, we assume that work utility at age $a \geq 55$ is given by

$$v_a = v_0 \frac{(a-54)^{-\eta_2}}{\eta_2} \text{ with } \eta_2 > 0$$

where v_0 is drawn from an exponential distribution with mean $\eta_1 > 0$. During retirement, flow consumption utility is given by

$$u(c^R) = \frac{[c^R]^{1-\gamma}}{1-\gamma}.$$

In particular, the curvature parameter γ is assumed to be the same across work and retirement. Additionally, notice that the flow consumption utility during retirement is equal to the flow utility during work with the work utility normalized to 1. Thus the parameters to be estimated are

$$\theta = \{\beta, \gamma, \eta_1, \eta_2\}.$$

In regard to expectations, we assume that individuals expect future earnings to grow at a constant rate of 2% per year. We assume that the probability of being separated from a given job varies only by years of tenure. The probabilities of job separation at each level of tenure are computed directly from the data. We assume that the separation shocks do not affect wages, so that conditional on separations, wages are still expected to grow at 2% per year. This assumption simplifies the computation of projected pension benefits. With this assumption, project pension benefits for each individual at each age can be computed based on the single path of projected earnings based on constant wage growth rather than based on multiple paths from different histories of job separations. Furthermore we assume that individuals form expectations for pension benefits based on the projected earnings with the constant growth rate and the current legislative framework for pensions.

Estimation of the model relies on asset data which we do not observe. To overcome this, we impute assets based on annual earnings, a constant savings rate, and a constant real interest rate. We assign initial values of wealth based on one quarter of mean annual earnings between ages 44 and 46. Given this initial wealth, we compute assets at subsequent ages by assuming that individuals save a constant fraction s of their annual income at each age y_a while earning interest at a rate r,

$$A_a = (1+r)[sy_a + A_{a-1}].$$

For this imputation, we use $r = .04$ and $s = .10$ where these numbers are based on the real interest rate and household savings rate in Austria during the sample period (see Dirschmid and Glatzer (2004)). Table 7 provides a summary of the imputed asset values at age 55. These asset statistics are roughly consistent with micro-data on household wealth in Austria from the Luxembourg Wealth Study.[24]

We estimate the dynamic model assuming a constant savings rate. Intuitively, the moment that is used to identify household savings is the average household savings rate over the sample period. In particular, rather than trying to fit this moment by estimating savings, this assumption restricts the model to exactly match this moment. This assumption also simplifies computation to estimate the model since the savings path for each individual can be computed mechanically rather than by using a grid search to solve for the optimal savings path. In the next section, we present our results and discuss the implications of this savings assumption, as well as other assumptions, for the estimates.

---

[24] The Luxembourg Wealth Study (LWS) provides the only micro-data on household savings in Austria of which we know. We do not use the LWS data more directly since doing so is likely to introduce significant measurement error due to the relatively small sample size (roughly 2,500 observations).

Lastly, we focus on the identification of γ in more detail. Specifically, a key identifying assumption is that utility is additively time-separable. With other, more general utility specifications, it is not possible to identify the curvature in the utility function since any monotonic transformation of the utility function can be reconciled with the data while also changing the degree of curvature. However, it is not possible to reconcile any monotonic transformation of additively time-separable utility with the data. Therefore, with the assumption of additively time-separable utility, we are able to identify the curvature parameter γ based on the ratio of income and price effects in the overall response to changes in benefits. Chetty (2006) presents a more general discussion of the relation between labor supply changes and the curvature of the utility function.

7.2 Estimation Results

Table 8 presents the results from estimating the structural model under the estimation strategy and assumptions discussed above. We present the results for the baseline specification of the model in the first column of the table. For this specification, we estimate a discount factor of .9974, a curvature with respect to consumption of .7905 and a work utility curvature with respect to age of .3947. The mean of the exponential distribution of initial work utility v_0 is .8829. Focusing first on the discount factor, β, which applies to utility at each age, this estimate is consistent with existing estimates for an annual discount factor. This estimate indicates that the model places relatively equal weight on current and future utility. Some of the discounting for future utility may be captured by the declining work utility across ages. We explore this in more detail in a separate specification of the model that we discuss below. Next, we focus on the curvature with respect to consumption, γ, the point estimate indicates a relatively low degree of curvature. This implies relatively large price effects and relatively small wealth effects. This result is exactly consistent with the earlier evidence from the hazard models.

In columns 2 and 3 of Table 8, we present estimates based on separately matching the retirement age and tenure moments respectively. The estimates do not change much indicating that the estimation fits both sets of moment conditions based on relatively large price effects and relatively small wealth effects. This is consistent with the results from separately examining the effects of pension benefits and severance payments on retirement decisions.

Figure 11A & B illustrate the goodness-of-fit of the baseline specification of the model.[25] These figures present the empirical moments and the predicted moments under the parameter estimates from column 1 of Table 8. Figure 11A presents the empirical retirement hazard rates by age and the corresponding predicted moments from the model. Generally the model predicts smother hazard rates than the empirical hazard rates. As the low degrees of curvature predicts relatively small wealth effects, we see that the model accounts from the spikes in the hazard rates at 60 and 65 based primarily on the price effects from the pension benefits. Some of the difference between the empirical moments and the predicted moments at these ages may be driven by over-emphasis of these particular ages since they correspond to the first eligibility age and the statutory retirement age for standard old-age pensions. The predicted hazards for ages 55 to 59 are higher than the corresponding empirical hazards. This may be due to the fact that individuals can only retire by claiming disability pension are these ages and eligibility for these pensions may be more restrictive than we have accounted for in the model. We consider this possibility in more detail below.

Figure 11B presents the empirical probabilities of retirement by years of tenure and the corresponding predicted moments from the model. The predicted moments fit the dips and spikes just before and after the tenure thresholds. The predicted retirement probabilities also follow a declining trend between the thresholds. These results indicate that the relatively large price effects created by the severance payments are able to account for the pattern illustrated in the graphical evidence on the severance payment effects in Figure 7.

The results in Figure 11B also show that the model over-estimates the retirement probabilities at the very low and very high tenure levels. This over-estimation may be driven by over-estimation of the separation probabilities at these tenure levels. To see this, suppose an individual becomes separated from his job either voluntarily or involuntarily. After the separation, the individual may work for a short period and then claim his retirement pension. Since we calculate the separation probabilities based on all experiences prior to retirement, this experience would count as a separation to zero years of tenure. However, this experience may correspond more realistically to a transition into retirement rather than a separation. In particular, it is plausible that such experiences should not be included as separations in the computation of the separation probabilities. Given the available data,

---

[25] In evaluating the goodness-of-fit of the model, we do not present the q statistic (where $q = \sum_{i=1}^{M}[m_i - \hat{m}_i(\hat{\theta}, X)]^2$) from the estimation since the goodness-of-fit test based on this statistic can only be implemented when using the optimal weighting matrix in the estimation.

however, it is difficult to make such distinctions regarding actual separations versus transitions into retirement.

In Table 9, we present estimates of the parameters under alternative specifications of the dynamic programming model to improve upon the previously-discussed shortcomings of the baseline specification. We first examine the assumption of a 10% savings rate. To get a sense of how much allowing for savings affects the estimated coefficient of relative risk aversion, we estimate the model with a 0% savings rate. Specifically, we assume that consumption while working is equal to current income and consumption during retirement is equal to current benefits. These results are presented in column 2 of Table 9. The estimates are nearly identical to those from the baseline specification. This is not surprising given that the change in the savings rate primarily affects individuals' wealth and the baseline model indicates relatively small wealth effects. Furthermore, the lack of change in the estimates indicates that allowing for liquidity constraints may not greatly affect the estimates or the empirical fit of the model. This may be due to the relatively high degree of generosity in the Austrian pension system.

It is also important to consider how endogenous savings may affect the estimates. First, restricting the model to only allow for adjustments to labor supply rather than allowing for adjustments to both savings and labor supply may lead the model to over-estimate the responsiveness to pension incentives. Second it is important to specifically consider how the estimated curvature may change. Since the reforms generally reduce the generosity of the pension benefits, individuals may choose to accumulate private savings. As a result, the relative wealth changes generating the observed labor supply responses would be smaller, thereby indicating larger wealth effects and a higher degree of curvature. Given the current estimates, it seems that savings would have to adjust dramatically in a short period of time to greatly change the estimates.

We next consider how the curvature in work utility across ages, η_2, affects the estimates of the coefficient of relative risk aversion, γ. Due to the complementarity between work utility and consumption utility, it may seem that the low estimate of γ in the baseline specification is driven by the specification of work utility. To address this concern, we estimate the model assuming that work utility declines linearly with age,

$$v_a = (66 - a)\, v_0.$$

The results from this specification of the model are presented in column 3 of Table 9.

Though the estimate for γ increases to .9371, the estimated degree of curvature remains relatively low. The estimated discount factor declines to .7283, while the mean of the exponential distribution of work utility increases to .9330. These results indicate that the curvature in work utility across ages relates more to the estimated discount factor than to the estimated γ. Thus, the specification of work utility across ages is not the driving factor behind the low coefficient of relative risk aversion.

In the next specification, we consider revisions to the separation probabilities. In the baseline specification, the separation probabilities may have been biased due to transitions to retirement. We attempt to account for this bias by computing revised separation probabilities based only on observations prior to age 55. As shown in column 4 of Table 9, the alternative specification does not change significantly from the baseline specification. Two features of the model may explain these results. First, we have not explicitly accounted for wage uncertainty associated with job separations. In particular, we assumed that wages do not change with a separation to simplify the computation of pension benefits. We have explored a specification allowing for wage changes at separations while not updating pension benefits based on the histories of separations. Estimates based on this specification do not change significantly. Second the model may have too much heterogeneity (variance) in work utility within each year of tenure, and this creates the more uniform pattern that is observed. Since we have used the exponential distribution for work utility, the mean and variance of the distribution are determined by a single parameter. Taking into account wage uncertainty and parameterizing the model with separate parameters to govern the mean and variance of work utility may improve the fit. However, it does not seem that this would alter the estimated coefficient of relative risk aversion as this parameter is identified off of labor supply responses to variation in pension benefits created by the pension reforms.

Next, we introduce a positive shift in work utility at ages 55 through 59. This positive shift is meant to capture the restrictions in claiming a disability pension for retirement prior to age 60. This shift increases work utility from v_a to $\tilde{v}_a = (1+k)v_a$ for $a = 55, ..., 59$ where $k > 0$. With this shift, individuals have higher work utility at these ages so that, prior to age 60, fewer individuals try to qualify as disabled for retirement. The estimates from this specification are show in column 5 of Table 9 and the corresponding predicted moments are shown in Figure 12. Focusing on the discount factor, the curvature parameter and the shift in work utility, this specification yields estimates of $\beta = .9306$, $\gamma = .8625$ and $k = 10.3750$. Thus, including the shift in work utility at ages 55-59 does not significantly change the estimated coefficient of relative risk aversion.

Turning to Figure 12, the predicted retirement hazards at ages 55 through 59 are lower and hence closer to the empirical moments relative to the baseline specification. Furthermore, the predicted spikes in the hazard rates at ages 60 and 65 are closer to the corresponding empirical moments. However, the hazards between ages 60 and 65 move further from the corresponding empirical moments. Thus, these results indicate that in addition to accounting for restrictions on claiming a disability pension, retirement may be concentrated at ages 60 and 65 due to some rule-of-thumb behavior corresponding to the age of first eligibility and the statutory retirement age.

Lastly, we present estimates based on an additively separable model. Specifically, we allow utility from consumption to be additively with the disutility of work. In this specification of the model, utility in each period is given by $u(c_a) - v_a$ where

$$\begin{aligned} u(c_a) &= \frac{c^{1-\gamma}}{1-\gamma} \\ v_a &= v_0 \frac{(a-54)^{1+\eta_2}}{1+\eta_2} \end{aligned}$$

and $v_0 \sim Exp(\eta_1)$. The estimates for this specification are presented in the last column of Table 9. We estimate a low degree of curvature, $\gamma = .4618$, consistent with the baseline, non-separable model. Additionally, the estimated discount factor declines to $\beta = .8784$. This is consistent with the decline in β in the non-separable case with the linear decline in work utility (column 3 of Table 9). Intuitively, the rate of decline in work utility across ages interacts more closely with the rate of time preference across ages rather than with the degree of curvature at each age.

Figures 13A & B present a graphical comparison of the additively separable and non-separable models. The hazard rates in Figure 13A indicate that the predictions from the additively separable model more closely correspond to the empirical moments at lower retirement ages but the non-separable model has a better fit at higher retirement ages. In Figure 13B, both models capture the patterns of declines and jumps just before and after the tenure thresholds, with the additively separable model predicting sharper declines and jumps than the non-separable model. However, the levels of the tenure moments from the non-separable model more closely correspond to the empirical moments. These results indicate that both models fit the data with a low degree of curvature.

The main point from these estimation results is that the estimated curvature, the coefficient of relative risk aversion γ, is relatively low. Thus, we see that the estimation fits the model to the data through relatively smaller wealth effects and larger price effects.

7.3 Policy Implications: Hypothetical Reforms

Given the estimates from the dynamic programming model, we now consider two hypothetical reforms to demonstrate the magnitudes of the income and price effects in the dynamic setting. We begin with a baseline that is determined using the predicted moments under the baseline parameter estimates of the structural model. The first reform we examine reduces pension benefits at all ages by 20%. Notice that this reform creates wealth effects only since the slope of the benefit schedule remains unchanged from the initial baseline. The second reform introduces penalties for retirement before age 65. Specifically, benefits are reduced by 2% per year prior to age 65. This reform therefore creates wealth and price effects. Using the results from the first reform allows for a decomposition of the wealth and price effects in the second reform.

The simulated labor supply responses to these reforms are shown in Figures 14A & B. Comparing the two figures, the second reform generates noticeably larger changes in behavior. This graphically demonstrates that the price effects are relatively larger than the wealth effects, though the wealth effects are non-negligible. To get a more precise sense of the relative magnitudes of these effects, notice that under the first reform the hazard rate at age 55 declines 13% from roughly .15 to .13. Given the 20% change in wealth, the implied elasticity with respect to pension wealth is .65. Turning to Figure 14B, the combination of wealth and price effects causes roughly a 45% decline in the age-55 hazard from .15 to .08. Since the wealth effect accounts for a decline of .02 in the hazard, this indicates that the price effect accounts for a decline of .05, roughly 70% of the total change and about 2.5 times larger than the wealth effect. Given an initial accrual of about 8%, the 2% change in the accrual implies an elasticity of roughly 1.6. These elasticities are consistent with those from the hazard model estimates (see Table 5).

Next, we consider the importance of accounting for forward-looking behavior beyond a one-year horizon. While the price measure in the hazard model is based on benefits in only the current and next year, the dynamic model takes into account benefits in the current year and all future, potential retirement years. We demonstrate the implications of this difference between the models by considering a hypothetical reform that provides a 20% bonus for retirement at age 65. Thus, for ages prior to age 64, the accrual measure is unchanged, but the overall benefit schedule is changed. Figure 15 plots the predicted retirement hazards under the baseline and reform specifications. While there is a decline in the hazard rates at all age prior to age 65, the hazards decline more noticeably at ages 61 through 64. These changes prior to age 64 would not be predicted by the hazard model. These results demonstrate that, while the accrual measure provides a rough measure of the price of retirement, it is important to use more forward-looking price measures that take the entire profile of benefits into account. This is consistent with results on the option-value and peak-value price measures developed by Stock and Wise (1990) and Coile and Gruber (2000b) respectively.

8 Conclusions & Future Research

While there is a large literature in economics examining the causal impacts of retirement benefits, the precise channels through which these benefits affect retirement decisions has not been clarified. In this study, we have separately identify the income and price effects from retirement benefits on retirement decisions. Our analysis using administrative, social-security-record data from the Austrian Social Security Database exploits variation in pension benefits created by multiple pension reforms in Austria between 1984 and 2003. Our results based on estimating hazard models and a dynamic programming model indicate that price effects account for 70% to 85% of the overall labor supply response to retirement benefits.

Our results have important implications for understanding social security reform. The relative importance of price effects indicates that the schedule of benefits across potential retirement ages greatly impacts retirement behavior. Hence, reforms based on changing the benefit schedule are likely to have larger impacts on retirement behavior compared to reforms using across-the-board changes in benefits. However, since our results are based only on such changes in the benefit schedule, it is difficult to use our results to evaluate more fundamental social security reforms.

Finally, our analysis highlights a potentially important avenue for future research relating to liquidity constraints at retirement. This avenue is motivated by the contrast between our results based on retirement benefits and earlier results based on other social insurance programs. Specifically, research on unemployment insurance (see Chetty (2007) and Card, Chetty, and Weber (2007), disability insurance (see Autor and Duggan (2007)) and health insurance (see Nyman (2003)) has identified larger income effects that account for a majority of the overall response to changes in the respective benefits. Our results may contrast from these earlier results because we do not explicitly account for liquidity constraints at retirement. However, liquidity constraints may be more relevant in the context of unanticipated changes at times of job loss or illness and less relevant in the context of anticipated changes at the time of retirement (see Aguiar and Hurst (2005)). Nevertheless, in future work, it would be useful to estimate a model taking liquidity constraints into account. Comparing the estimated income effect with liquidity constraints to the corresponding estimate from the permanent income model considered here would provide insight into individuals' abilities to smooth marginal utilities at retirement.

References

Aguiar, M. and E. Hurst (2005). Consumption versus Expenditure. *Journal of Political Economy 113*(5), 919–948.

Asch, B., S. Haider, and J. Zissimopoulos (2005). Financial Incentives and Retirement: Evidence from Federal Civil Service Workers. *Journal of Public Economics 89*, 427–440.

Autor, D. and M. Duggan (2007). Distinguishing Income from Substitution Effects in Disability Insurance. *American Economic Review Papers and Proceedings 97*(2), 119.

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

Brown, J., C. Coile, and S. Weisbenner (2006). The Effect of Inheritance Receipt on Retirement. Working Paper 12386, National Bureau of Economic Research.

Brown, K. (2006). The Link between Pensions and Retirement Timing: Lessons from California Teachers. Working paper, University of California - Berkeley.

Burtless, G. (1986). Social Security, Unanticipated Benefit Increases, and the Timing of Retirement. *Review of Economic Studies 53*(5), 781–805.

Card, D., R. Chetty, and A. Weber (2007). Cash-on-Hand and Competing Models of Intertemporal Behavior: New Evidence from the Labor Market. *Quarterly Journal of Economics*.

Chetty, R. (2006). A New Method of Estimating Risk Aversion. *American Economic Review 96*(5), 1821–1834.

Chetty, R. (2007). Moral Hazard vs. Liquidity in Unemployment Insurance. Working paper, University of California - Berkeley.

Coile, C. (2004). Retirement Incentives and Couples' Retirement Decisions. *Topics in Economic Analysis & Policy 4*(1), 1277–1277.

Coile, C. and J. Gruber (2000a). Social Security Incentives for Retirement. Working Paper 7651, National Bureau of Economic Research.

Coile, C. and J. Gruber (2000b). Social Security and Retirement. Working Paper 7830, National Bureau of Economic Research.

Costa, D. (1995). Pensions and Retirement: Evidence from Union Army Veterans. *Quarterly Journal of Economics 110*(2), 297–319.

Dirschmid, W. and E. Glatzer (2004). Determinants of the Household Saving Rate in Austria. *Monetary Policy & the Economy, December* (4), 25–38.

Feldstein, M. and J. Liebman (2002). Social Security. *The Handbook of Public Economics*, 2245–2324.

Friedberg, L. (1999). The Effect of Old Age Assistance on Retirement. *Journal of Public Economics 71*, 213–232.

Friedberg, L. (2000). The Labor Supply Effects of the Social Security Earnings Test. *Review of Economics and Statistics 82*(1), 48–63.

Goda, G., J. Shoven, and S. Slavov (2007). Removing the Disincentives in Social Security for Long Careers. Working Paper 13110, National Bureau of Economic Research.

Gruber, J. and D. Wise (2004). *Social Security Programs and Retirement around the World: Micro-Estimation*. University of Chicago Press.

Gustman, A. and T. Steinmeier (1986). A Structural Retirement Model. *Econometrica 54*(3), 555–584.

Gustman, A. and T. Steinmeier (2000). Retirement in Dual-Career Families: A Structural Model. *Journal of Labor Economics 18*(3), 503–545.

Hofer, H. and R. Koman (2006). Social Security and Retirement Incentives in Austria. *Empirica 33*(5), 285–313.

Holtz-Eakin, D., D. Joulfaian, and H. Rosen (1999). Estimating the Income Effect on Retirement. Aging Studies Program Paper 18, Maxwell Center for Demography and Economics of Aging, Syracuse University.

Jimenez-Martin, S. and A. Sanchez-Martin (2007). An Evaluation of the Life-cycle Effects of Minimum Pensions on Retirement Behavior. *Journal of Applied Econometrics* (5), 923–950.

Koman, R., U. Schuh, and A. Weber (2005). The Austrian Severance Pay Reform: Toward a Funded Pension Pillar. *Empirica 32*(3), 255–274.

Krueger, A. and J. Pischke (1992). The Effect of Social Security on Labor Supply: A Cohort Analysis of the Notch Generation. *Journal of Labor Economics 10*(4), 412–437.

Linnerooth-Bayer, J. (2001). Pension Reform in Austria. *International Institute for Applied System Analysis. Report to the PEN-REF project*.

Lumsdaine, R., J. Stock, and D. Wise (1992). Three Models of Retirement: Computational Complexity Versus Predictive Validity. Working Paper 3558, National Bureau of Economic Research.

Madrian, B. (2006). The U.S. Health Care System and Labor Markets. Working Paper 11980, National Bureau of Economic Research.

Marek, E. (1985, 1987-2003). *Pensionsberechnung Leicht Gemacht*. Weiss Verlag.

Ney, S. (2004). Pension Reform and Active Ageing in Austria. *Project Deliverable, Interdisciplinary Centre for Comparative Research in the Social Sciences*.

Nyman, J. (2003). *The Theory of Demand for Health Insurance*. Stanford University Press.

Rust, J. and C. Phelan (1997). How Social Security and Medicare Affect Retirement Behavior In a World of Incomplete Markets. *Econometrica 65*(4), 781–831.

Samwick, A. (1998). New Evidence on Pensions, Social Security, and the Timing of Retirement. *Journal of Public Economics 70*, 207–236.

Stock, J. and D. Wise (1990). Pensions, the Option Value of Work, and Retirement. *Econometrica 58*(5), 1151–1180.

Appendix 1

Decomposing the Response to a Change in the Profile of Retirement Benefits.

As mentioned in the text, we consider a setting in which benefits and wages increase linearly with age,

$$\begin{aligned} y_a^R &= \bar{b} + a\Delta^R \\ y_a^W &= \bar{y} + a\Delta^W. \end{aligned}$$

With these linear functions, note that $\frac{dy_a^R}{d\bar{b}} = \frac{dy_a^W}{d\bar{y}} = 1$ and $\frac{dy_a^R}{d\Delta^R} = \frac{dy_a^W}{d\Delta^W} = a$. We focus first on a change to the intercept, $\bar{b}$. Differentiating the equation for $\bar{v}_a$ with respect to $\bar{b}$ yields

$$\begin{aligned} [\bar{b}] \quad : \quad & \sum_{t=a}^{T} \pi_{t|a} \beta^{t-a} \frac{du(c_t^R)}{dc} \frac{dc_t^R}{dy_a^R} \frac{dy_a^R}{d\bar{b}} = \\ & \frac{du(c_a^W, \bar{v}_a)}{dv} \frac{d\bar{v}_a}{dy_a^R} \frac{dy_a^R}{d\bar{b}} + \\ & \underbrace{\sum_{t=a+1}^{a^*} \pi_{t|a} \beta^{t-a} [\Psi(\bar{v}_t) \frac{dR_t}{d\bar{b}} + \frac{d\Psi(\bar{v}_t)}{dv} \frac{d\bar{v}_t}{dy_t^R} \frac{dy_t^R}{d\bar{b}} (R_t - W_t)]}_{Term\ 1} \end{aligned}$$

where

$$\frac{dR_t}{d\bar{b}} = \sum_{s=t}^{T} \pi_{s|t} \beta^{s-t} \frac{du(c_s^R)}{dc} \frac{dc_s^R}{dy_t^R} \frac{dy_t^R}{d\bar{b}}.$$

Next, we introduce an annuity that pays the individual a constant $\bar{A}$ in each period, regardless of his work status. Note that $\frac{dA_a}{d\bar{A}} = 1$ for each age a. Differentiating the equation for $\bar{v}_a$ with respect to $\bar{A}$ yields

$$\begin{aligned} [\bar{A}] \quad : \quad & \sum_{t=a}^{T} \pi_{t|a} \beta^{t-a} \frac{du(c_t^R)}{dc} \frac{dc_t^R}{dA_a} \frac{dA_a}{d\bar{A}} = \\ & \frac{du(c_a^W, \bar{v}_a)}{dc} \frac{dc_a^W}{dA_a} \frac{dA_a}{d\bar{A}} + \frac{du(c_a^W, \bar{v}_a)}{dv} \frac{d\bar{v}_a}{dA_a} \frac{dA_a}{d\bar{A}} + \\ & \underbrace{\sum_{t=a+1}^{a^*} \pi_{t|a} \beta^{t-a} [\frac{du(c_t^W, v_t)}{dc} \frac{dc_t^W}{dA_t} \frac{dA_t}{d\bar{A}} + \Psi(\bar{v}_t) (\frac{dR_t}{d\bar{A}} - \frac{du(c_t^W, v_t)}{dc} \frac{dc_t^W}{dA_t} \frac{dA_t}{d\bar{A}}) + \frac{d\Psi(\bar{v}_t)}{dv} \frac{d\bar{v}_s}{dA_t} \frac{dA_t}{d\bar{A}} (R_t - W_t)]}_{Term\ 2} \end{aligned}$$

where

$$\frac{dR_t}{d\bar{A}} = \sum_{s=t}^{T} \pi_{s|t} \beta^{s-t} \frac{du(c_s^R)}{dc} \frac{dc_s^R}{dA_s} \frac{dA_s}{d\bar{A}}.$$

Lastly, we consider a change in the wage intercept. Differentiating the equation for $\bar{v}_a$

with respect to $\bar{y}$ yields

$$[\bar{y}] : 0 = \quad \tfrac{du(c_a^W,\bar{v}_a)}{dc}\tfrac{dc_a^W}{dy_a^W}\tfrac{dy_a^W}{d\bar{y}} + \tfrac{du(c_a^W,\bar{v}_a)}{dv}\tfrac{d\bar{v}_a}{dy_a^W}\tfrac{dy_a^W}{d\bar{y}} + \underbrace{\sum_{t=a+1}^{a^*} \pi_{t|a}\beta^{t-a}[\tfrac{du(c_t^W,v_t)}{dc}\tfrac{dc_t^W}{dy_t^W}\tfrac{dy_t^W}{d\bar{y}} + \Psi(\bar{v}_{a+1})(-\tfrac{du(c_t^W,v_t)}{dc}\tfrac{dc_t^W}{dy_t^W}\tfrac{dy_t^W}{d\bar{y}}) + \tfrac{d\Psi(\bar{v}_t)}{dv}\tfrac{d\bar{v}_t}{dy_t^W}\tfrac{dy_t^W}{d\bar{y}}(R_t - W_t)]}_{Term\ 3}.$$

Using the relations $\frac{dc_a^R}{dy_a^R} = \frac{dc_a^R}{dA_a}$ and $\frac{dc_a^W}{dA_a} = \frac{dc_a^W}{dy_a^W}$, these equations can be combined to result in

$$\tfrac{du(c_a^W,\bar{v}_a)}{dv}\tfrac{d\bar{v}_a}{dy_a^R}\tfrac{dy_a^R}{d\bar{b}} + [Term\ 1] = \tfrac{du(c_a^W,\bar{v}_a)}{dv}\tfrac{d\bar{v}_a}{dA_a}\tfrac{dA_a}{d\bar{A}} + [Term\ 2] - \tfrac{du(c_a^W,\bar{v}_a)}{dv}\tfrac{d\bar{v}_a}{dy_a^W}\tfrac{dy_a^W}{d\bar{y}} - [Term\ 3].$$

If $[Term\ 1] = [Term\ 2] - [Term\ 3]$, we can use $\frac{dy_a^R}{d\bar{b}} = \frac{dA_a}{d\bar{A}} = \frac{dy_a^W}{d\bar{y}} = 1$ to arrive at the desired result. Lastly, notice that $[Term\ 1] = [Term\ 2] - [Term\ 3]$ follows directly from the observations that $\frac{d\bar{v}_a}{dy_a^R} = \frac{d\bar{v}_a}{dA_a} - \frac{d\bar{v}_a}{dy_a^W}$, $\frac{dc_a^R}{dy_a^R} = \frac{dc_a^R}{dA_a}$ and $\frac{dc_a^W}{dA_a} = \frac{dc_a^W}{dy_a^W}$.

Next, consider a change in the slope of the benefit profile, Δ^R. In this case, we introduce and annuity that pays the individual a different amount in each period, but the slope of annuity income across ages is constant and given by Δ^A. Notice that we can use the same equations as above, except now we must replace $\frac{dy_a^R}{d\bar{b}}$, $\frac{dA_a}{d\bar{A}}$ and $\frac{dy_a^W}{d\bar{y}}$ with $\frac{dy_a^R}{d\Delta^R}$, $\frac{dA_a}{d\Delta^A}$ and $\frac{dy_a^W}{d\Delta^W}$ respectively. Thus, following the identical arguments as above and using $\frac{dy_a^R}{d\Delta^R} = \frac{dA_a}{d\Delta^A} = \frac{dy_a^W}{d\Delta^W} = a$, we can arrive at the result given in the text.

Appendix 2

Maximum Likelihood (ML) Estimation

This appendix develops an ML estimation strategy for the dynamic model presented in the text. While we have focused on an MSM estimation strategy in the current work, we plan to implement this ML estimation strategy as well. Doing so allows for a comparison of the two estimation strategies on both economic and technical grounds.

Conditional on the individual's assets and labor market observables, the probability of retirement at age a is

$$\Pr(\text{retire at age } a \mid A_a, X_a) = \Pr(v_a \leq \bar{v}_a(A_a, X_a) | A_a, X_a).$$

Integrating over the distribution of assets, denoted by $\Phi(A_a)$, yields the probability of retirement at age a conditional on the observables X_a

$$\Pr(\text{retire at age } a \mid X_a) = \int \Pr(v_a \leq \bar{v}_a(A_a, X_a) | A_a, X_a) d\Phi(A_a | X_a).$$

This probability is all that is necessary to form a likelihood function if the data is structured so that retirement status at a single age is observed for each individual. However, since the data contains information on individuals' retirement decisions over multiple ages, this single-age probability must be augmented.

Suppose the data is structured so that individuals' retirement decisions are observed over the age range $[\underline{a}, \bar{a}]$ and all individuals are not retired prior to age $\underline{a}$. The probability of observing an individual retire at age $a \in [\underline{a}, \bar{a}]$ must take into account that the individual did not retire at earlier ages. In this case, the probability that the individual retires at age a is given by

$$\begin{aligned} \Pr(\text{retire at age } a &\in [\underline{a}, \bar{a}] \,|(A^a, X^a)) = \\ \Pr(v_{\underline{a}} &> \bar{v}_{\underline{a}}(A_{\underline{a}}, X_{\underline{a}}), v_{\underline{a}+1} > \bar{v}_{\underline{a}+1}(A_{\underline{a}+1}, X_{\underline{a}+1}), \ \ldots, v_a \leq \bar{v}_a(A_a, X_a) | (A^a, X^a)). \end{aligned}$$

where (A^a, X^a) denotes the histories of assets and observables from age $\underline{a}$ to age a

$$(A^a, X^a) = ((A_{\underline{a}}, A_{\underline{a}+1}, \ \ldots, A_a), (X_{\underline{a}}, X_{\underline{a}+1}, \ \ldots, X_a)).$$

Furthermore, notice if an individual is observed to not retire over the given age range, this observation still contributes to the likelihood function through the probability Pr(retire at age $a \geq \bar{a} | (A^{\bar{a}}, X^{\bar{a}}))$ which can be computed using the probability above. Analogous to the single-age probability of retirement, we can integrate over the distribution of assets at each age conditional on observables to get the corresponding probabilities of retirement conditional on observables only

$$\begin{aligned} \Pr(\text{retire at age } a &\in [\underline{a}, \bar{a}] \; | X^a) = \\ \int \dots \int \Pr(\text{retire at age } a &\in [\underline{a}, \bar{a}] \; | (A^a, X^a)) d\Phi(A_a | X_a) \; \dots \; d\Phi(A_{\underline{a}} | X_{\underline{a}}). \end{aligned}$$

The estimation strategy is follows. Let θ denote the parameters to be estimated and let $p_i(\theta)$ denote the retirement probability for individual i, i.e. we have $p_i(\theta) = \Pr($retire at age $a_i \in [\underline{a}, \bar{a}] \; | X_i^a)$. The log-likelihood function is therefore

$$\mathcal{L}(\theta) = \sum_{i=1}^{I} p_i(\theta).$$

Once a parameter vector is specified, the model can be solved recursively for each individual. This yields a retirement probability for each individual so that the log-likelihood function can be computed. Therefore, given an initial parameter vector, the parameters are estimated by numerically maximizing the log-likelihood function.

Table 1
Sample Restrictions, Initial Sample (Males, Birth Cohorts ≥ 1948): 2403454

Sample Restriction	Sample After Restriction	# of Individuals Excluded
1. Less than 1 year of employment in 1972-2003	1512323	891131
2. Non-Austrian nationality	1417209	95114
3. Public servants, mining, rail, farmers, construction for 1 or more years	1075285	341924
4. Self-employed for 1 or more years	744597	330688
5. Claiming non-old-age or non-disability pensions	720308	24289
6. Claiming before age 55	648305	72003
7. Claiming or last observed before 1984	394934	253371
Samples Based on Complete Data, Ages 55-70:		
Severance Payment Sample, # of Individuals	270946	
Pension Sample, # of Individuals (# of Claimants)	339376 (259517)	

Notes: The Severance Payment Sample includes observations with uncensored tenure spells from all years observed, 1972 through 2003. The Pension Sample includes observations from 1984 through 2003 since pensions can only be computed for these years given the legislative information available to us. The number of claimants in the Pension Sample is less than the number of individuals in the sample since younger individuals toward the end of the sample have yet to claim old-age pensions. Further details regarding the samples are contained in the text.

Table 2

Summary of Austrian Pension Reforms - 1984 - 2003

1985 Pension Reform	1988 Pension Reform	1993 Pension Reform	1996 Pension Reform	2000 Pension Reform
change in assessment basis from last 5 years to last 10 years of earnings	change in assessment basis from last 10 years to last 15 years of earnings, phased in 1988-1992	change in assessment basis from last 15 to best 15 years of earnings	introduction of bonus / malus system (lower pension coefficient to penalize early retirement)	development of bonus / malus system (increased penalities for early retirement)
		change in revaluation factors used in assessment basis		increase in minimum retirement age from 60 to 61.5, phased in 2000 - 2002
		linking pension coefficient to retirement age		increased restrictions for claiming disability pension
		introduction of early retirement due to reduced working capacity at age 57		elimination of early retirement due to reduced working capacity at age 57

Notes: Please see text for more details regarding the pension reforms.

Table 3
Severance Pay Statistics

Years of Tenure	Mean	Median	Std. Dev.
≤ 9	0	0	0
10	5509.022	3702.813	5634.52
11	5851.349	4128.048	5306.572
12	6329.918	4589.569	5470.313
13	6807.162	5075.87	5727.344
14	7219.541	5474.673	5606.503
15	11452.46	8883.572	9167.197
16	11989.68	9430.788	9044.891
17	12643.19	9971.554	9295.396
18	13467.34	10713.31	10269.48
19	14270.69	11486.08	10393.53
20	23254.69	18782.44	16431.01
21	25112.16	20456.8	16786.69
22	26939.98	22082.51	17162.67
23	28546.99	23395.7	18780.9
24	29913.81	24836.43	18648.91
25	42192.09	35404.13	26219.39
26	43670.82	37029.73	24553.73
27	44783.03	38271.75	24743.35
28	46744.89	39739.88	26738.13
29	49003.58	41786.64	28064.31
30	50721.91	43178.67	28380.45
31	52247.66	44395.91	28987.65
32	53642.18	45913.36	29132.96

Notes: Statistics are in 2003 euros.

Table 4
Effects of Severance Payments, Hazard Model Estimates

	Full Sample		Ages ≠ 60		Ages = 60	
	Base Controls	Full Controls	Base Controls	Full Controls	Base Controls	Full Controls
ln(SEV)	0.0455	0.0479	0.1754	0.1934	0.0225	-0.0341
	(0.1260)	(0.1265)	(0.1754)	(0.1762)	(0.1820)	(0.1824)
ln(ΔSEV)	-0.4936	-0.4931	-0.5477	-0.5807	-0.2071	-0.2313
	(0.0930)	(0.0934)	(0.1297)	(0.1301)	(0.1340)	(0.1344)
Observations	969589	969589	832478	832478	137111	137111

Notes: Standard errors are shown in parentheses. All coefficient estimates should be interpreted as changes in the retirement hazard. All specifications include the following base controls: education dummies and quartic polynomials in calendar year, tenure, annual earnings, and total earnings in the prior 10 years. The full controls specifications include the base controls, industry and region dummies, and quartic polynomials in earnings from each of the prior 10 years.

Table 5
Effects of Pension Incentives, Hazard Model Estimates

	Full Sample		Ages ≠ 60		Age = 60		Ages ≥ 61	
	Base Controls	Full Controls	Base Controls	Full Controls	Base Controls	Full Controls	Base Controls	Full Controls
ln(SSW)	0.5284	0.3332	0.6817	0.4876	0.2101	0.1297	0.8175	0.6484
	(0.0116)	(0.0139)	(0.0158)	(0.0187)	(0.0179)	(0.0219)	(0.0302)	(0.0345)
ln(Δ)	-1.0192	-1.6114	-1.7407	-2.1921	-0.5680	-1.0322	-1.3691	-1.4520
	(0.0476)	(0.0611)	(0.0496)	(0.0610)	(0.0978)	(0.1470)	(0.1222)	(0.1566)
Observations	1338384	1338384	1167210	1167210	171174	171174	71721	71721

Notes: Standard errors are shown in parentheses. All coefficient estimates should be interpreted as changes in the retirement hazard. The base controls include education dummies and quartic polynomials in year, log annual earnings and log total earnings from the previous 10 years. The full controls include the base controls and additionally, quartic polynomials in log annual earnings from each of the ten previous years as well as region, industry and tenure dummies. To account for censoring, tenure dummies are interacted with a censored dummy indicating whether or not the spell started in 1972 or if it is fully observed. All specifications include dummies to capture changes in eligibility at age 57 between years 1993 and 2000, at ages 60 and 61 in 2000 and after and also for ages less than age 60 after 2000.

Table 6

Effects of Pension Incentives, Linear Probability Model Estimates

	Full Sample		Ages ≠ 60			Full Sample	
	Base Controls	Full Controls	Base Controls	Full Controls		Base Controls	Full Controls
ln(SSW)	0.0708	0.0605	0.0624	0.0582	$\ln(SSW)_{ages \leq 59}$	0.0567	0.0411
	(0.0016)	(0.0019)	(0.0017)	(0.0020)		(0.0018)	(0.0020)
					$\ln(SSW)_{age = 60}$	0.0732	0.0718
						(0.0027)	(0.0029)
					$\ln(SSW)_{age \geq 61}$	0.1464	0.1340
						(0.0026)	(0.0027)
$\ln(\Delta)$	-1.3645	-1.4640	-1.6197	-1.7302	$\ln(\Delta)_{ages \leq 59}$	-2.4645	-2.6112
	(0.0242)	(0.0245)	(0.0275)	(0.0278)		(0.0327)	(0.0331)
					$\ln(\Delta)_{ages = 60}$	-0.2891	-0.3264
						(0.0425)	(0.0427)
					$\ln(\Delta)_{ages \geq 61}$	-0.6783	-0.8463
						(0.0400)	(0.0398)
Observations	1338384	1338384	1167210	1167210		1338384	1338384

Notes: Standard errors are shown in parentheses. The base controls include dummies for age, education and wether or not a claim date is observed. Additionally, the base controls include quartic polynomials in year, insurance months divided by 10, log annual earnings and log total earnings from the previous 10 years. The full controls include the base controls and additionally, quartic polynomials in log annual earnings from each of the ten previous years as well as region, industry and tenure dummies. To account for censoring, tenure dummies are interacted with a censored dummy indicating whether or not the spell started in 1972 or if it is fully observed. All specifications include dummies to capture changes in eligibility at age 57 between years 1993 and 2000, at ages 60 and 61 in 2000 and after and also for ages less than age 60 after 2000.

Table 7
Imputed Wealth at Age 55, Summary Statistics

Percentile	Wealth
1%	0
5%	1973.74
10%	14367.42
25%	22510.38
50%	33149.06
75%	49028.09
90%	69005.14
95%	83151.05
99%	115982.7
Mean	38015.38
Observations	172993

Table 8

Structural Parameter Estimates

	10% Savings Rate	10% Savings Rate, Age Moments Only	10% Savings Rate, Tenure Moments Only
Discount Rate: β	0.9974	0.9911	1.0002
	(0.0182)	(0.0180)	(0.0179)
Curvature of Utility: γ	0.7905	0.8189	0.8243
	(0.0447)	(0.0372)	(0.0446)
Distribution of Work Utility: η_1	0.8829	0.8778	0.8592
	(0.0567)	(0.0504)	(0.0533)
Curvature of Work Utility w.r.t. Age: η_2	0.3947	0.3840	0.3970
	(0.0276)	(0.0252)	(0.0320)

Notes: Estimates are based on a sample of 244,759 individuals with 853,111 observations. To estimate the structural parameters, we use the same Pension Sample used to estimate the hazard models but with the additional restriction that full tenure is observed. Thus, observations with censored tenure are excluded in the estimation. Standard errors are shown in parentheses. The standard errors are estimated using numerical derivatives of the moments conditions with respect to the parameters. Please refer to the text for additional details on the structural estimation.

Table 9

Structural Parameter Estimates, Alternative Specifications

	Baseline	0% Savings Rate	Linear Decline in Work Utility Across Ages	Revised Separation Probabilities	Shift in Work Utility Ages 55-59	Additively Separable Model
Discount Rate: β	0.9974	1.0022	0.7283	0.9936	0.9306	0.8784
	(0.0182)	(0.0192)	(0.0305)	(0.0190)	(0.0196)	(0.0148)
Curvature of Utility: γ	0.7905	0.8050	0.9371	0.8053	0.8625	0.4618
	(0.0447)	(0.0454)	(0.0333)	(0.0415)	(0.0436)	(0.0100)
Distribution of Work Utility: η_1	0.8829	0.8309	0.9330	0.8758	0.9039	2.0974
	(0.0567)	(0.0616)	(0.0437)	(0.0512)	(0.0344)	(0.0772)
Curvature of Work Utility w.r.t. Age: η_2	0.3947	0.4093		0.3940	0.3701	1.1200
	(0.0276)	(0.0341)		(0.0345)	(0.0204)	(0.0197)
Shift in Work Utility, Ages 55-59: k					10.3750	
					(0.0294)	

Notes: Estimates are based on a sample of 244,759 individuals with 853,111 observations. To estimate the structural parameters, we use the same Pension Sample used to estimate the hazard models but with the additional restriction that full tenure is observed. Thus, observations with censored tenure are excluded in the estimation. Standard errors are shown in parentheses. The standard errors are estimated using numerical derivatives of the moments conditions with respect to the parameters. Please refer to the text for additional details on the structural estimation.

Figure 1: Optimal Retirement Age

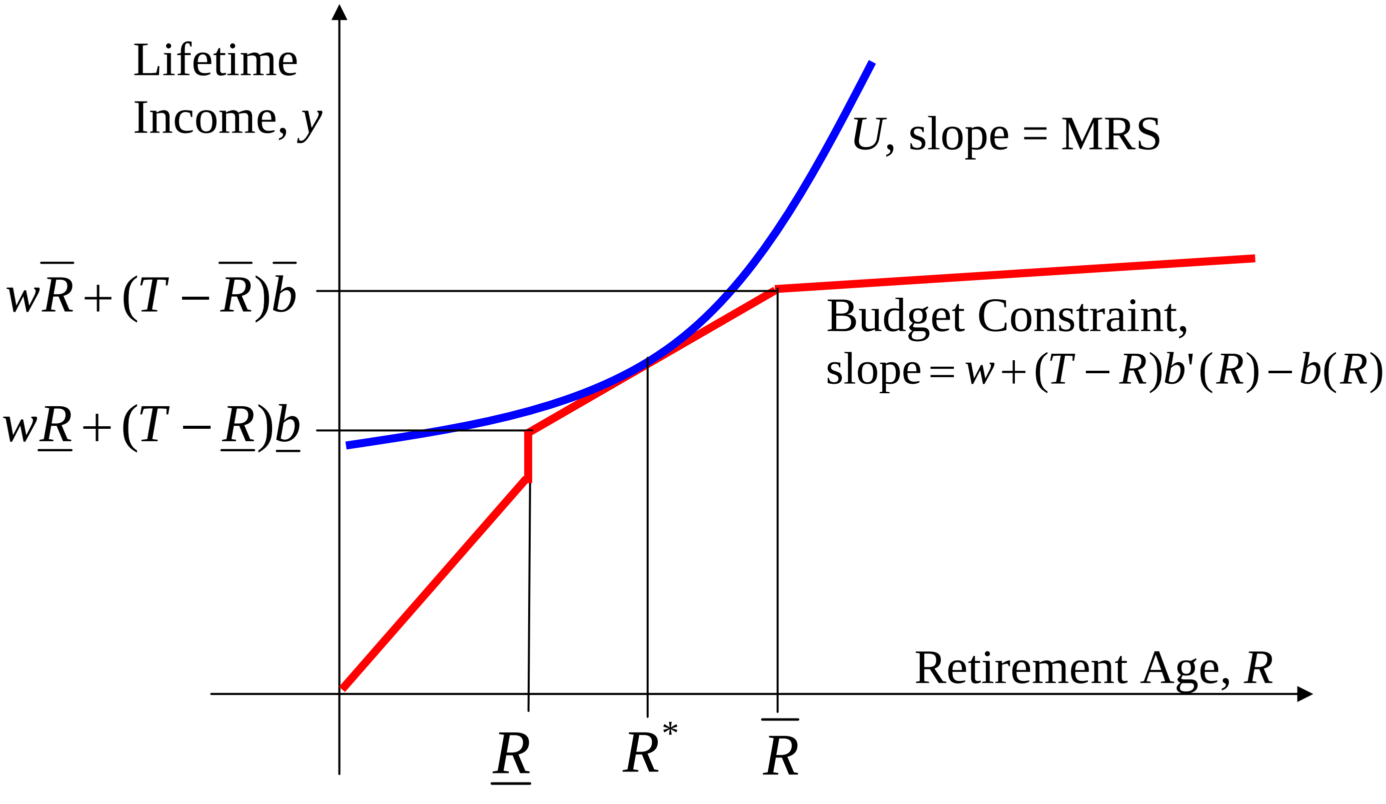

Figure 2: Income & Price Effects:

Identification from Changes in the Level & Slope of the Benefit Schedule

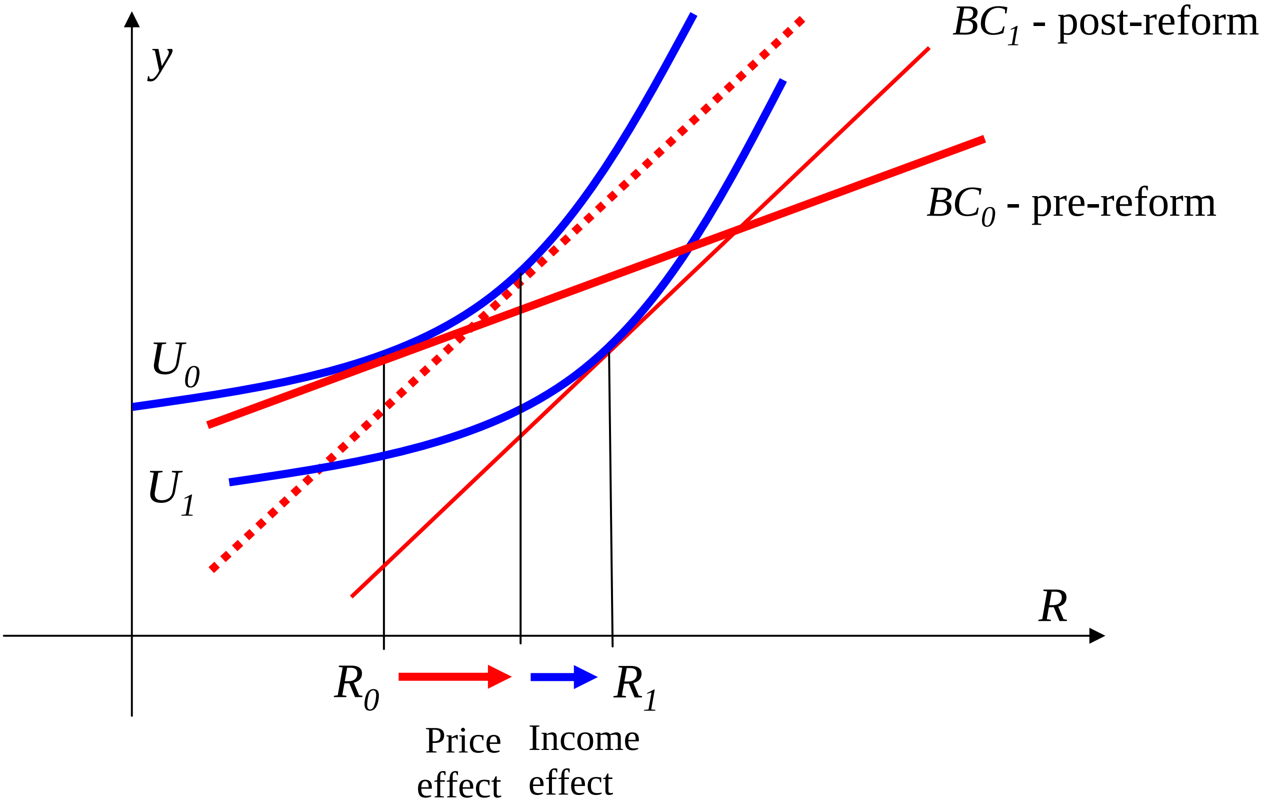

An increase in $b'(R)$ increases the price of retirement.

A decrease in $b(R)$ decreases income (wealth) at retirement.

Both cause later retirement.

Figure 3: Effects of Severance Payments

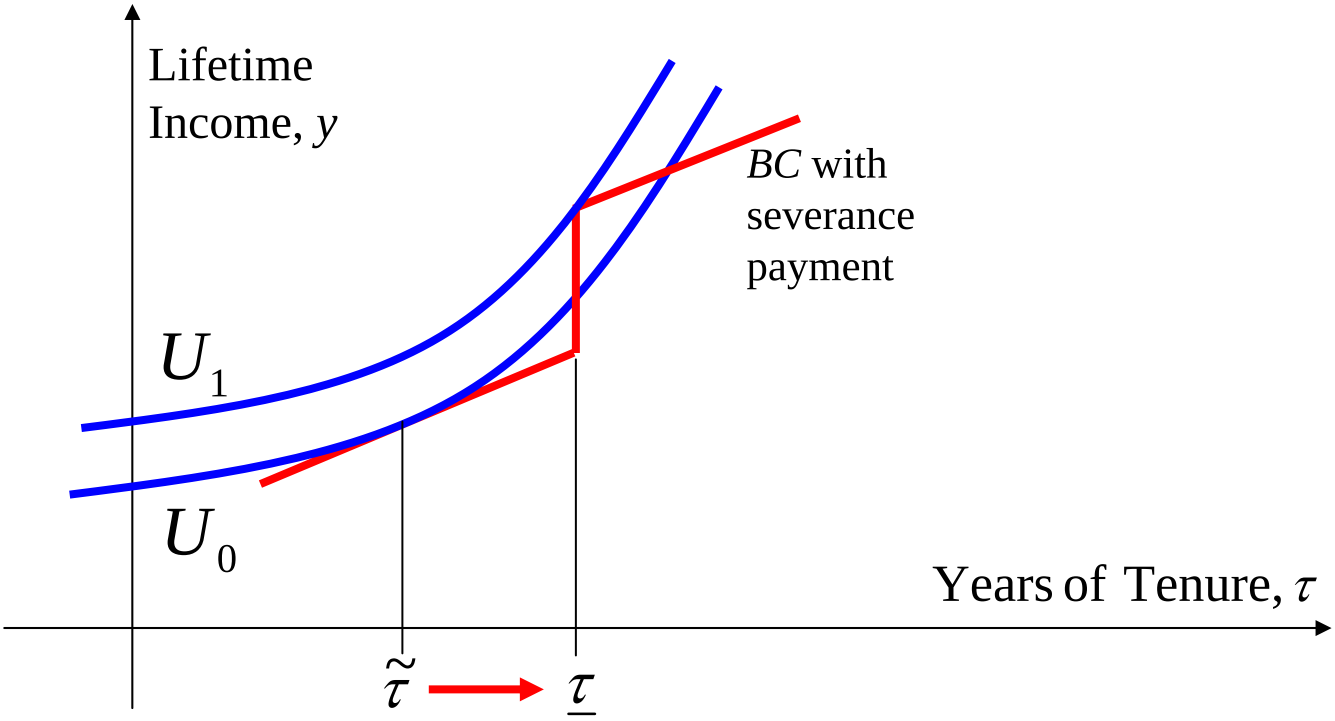

Severance payments create incentives for individuals to delay retirement and accumulate qualifying years of tenure.

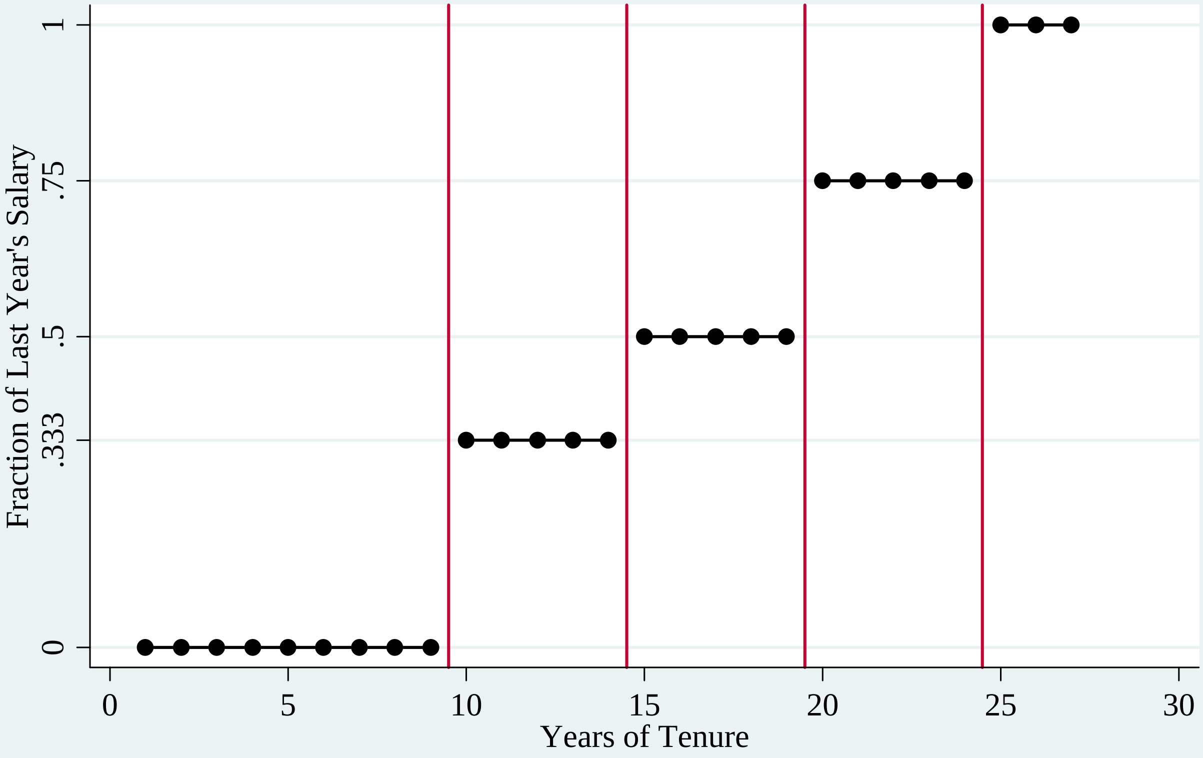

Figure 4
Severance Pay Schedule

Figure 5

Hazard Rates Into Retirement Pensions

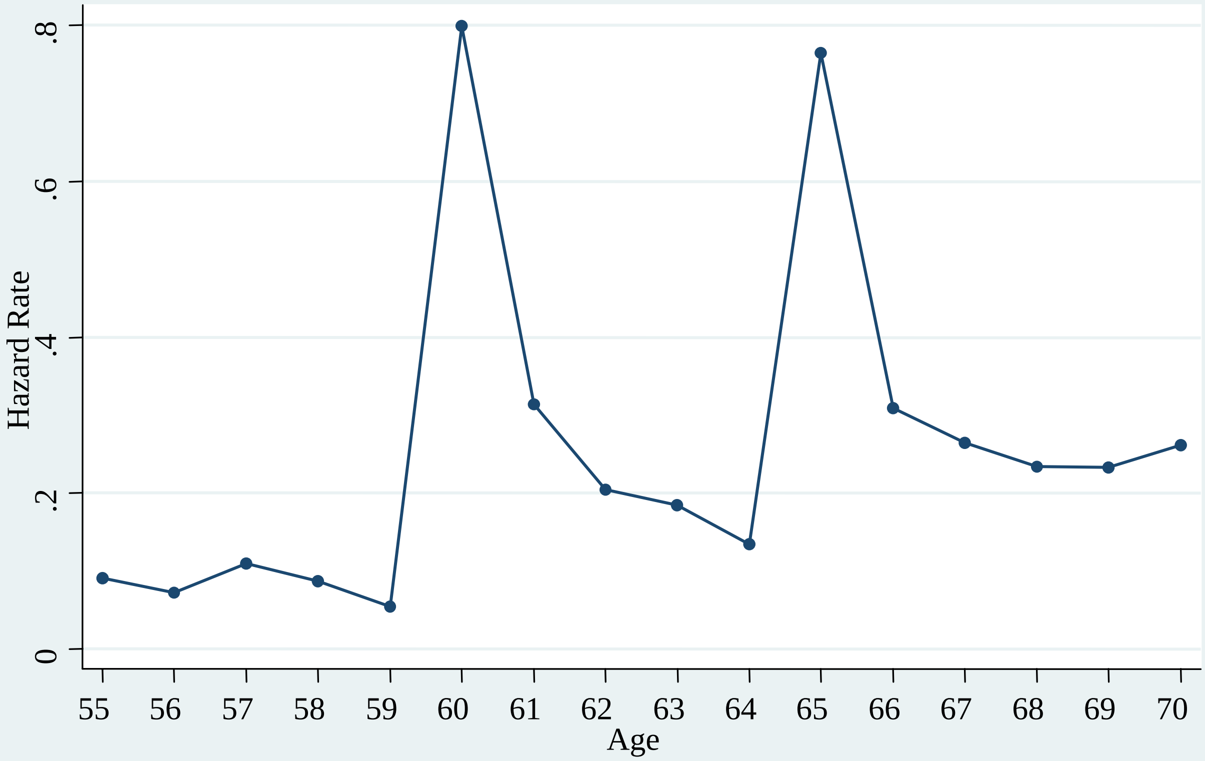

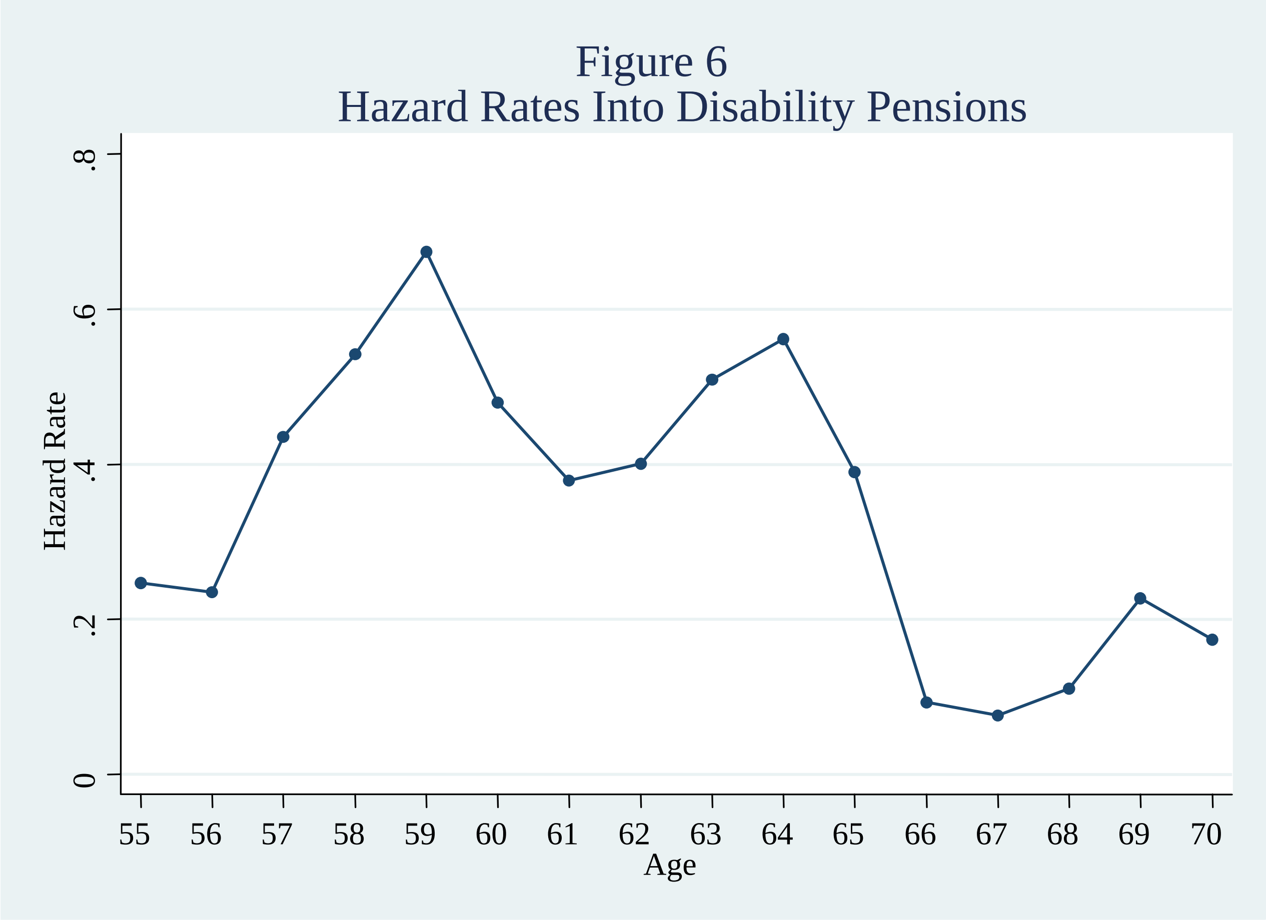

Figure 6
Hazard Rates Into Disability Pensions

Figure 7: Effects of Severance Pay on Retirement Decisions

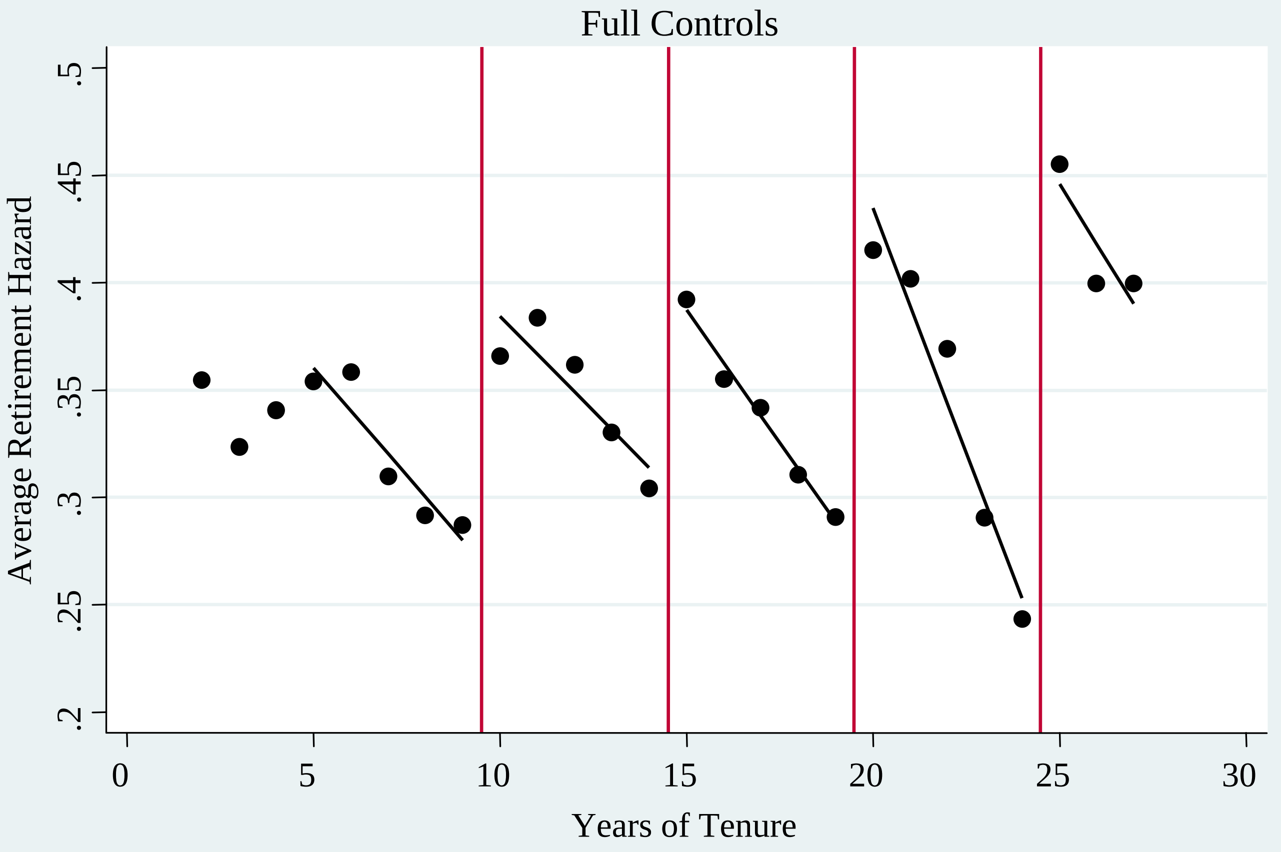

Figure 8: Identification of Income & Price Effects

Legislated Changes & Retirement Responses
1988 & 1996 Pension Reforms, Age 55

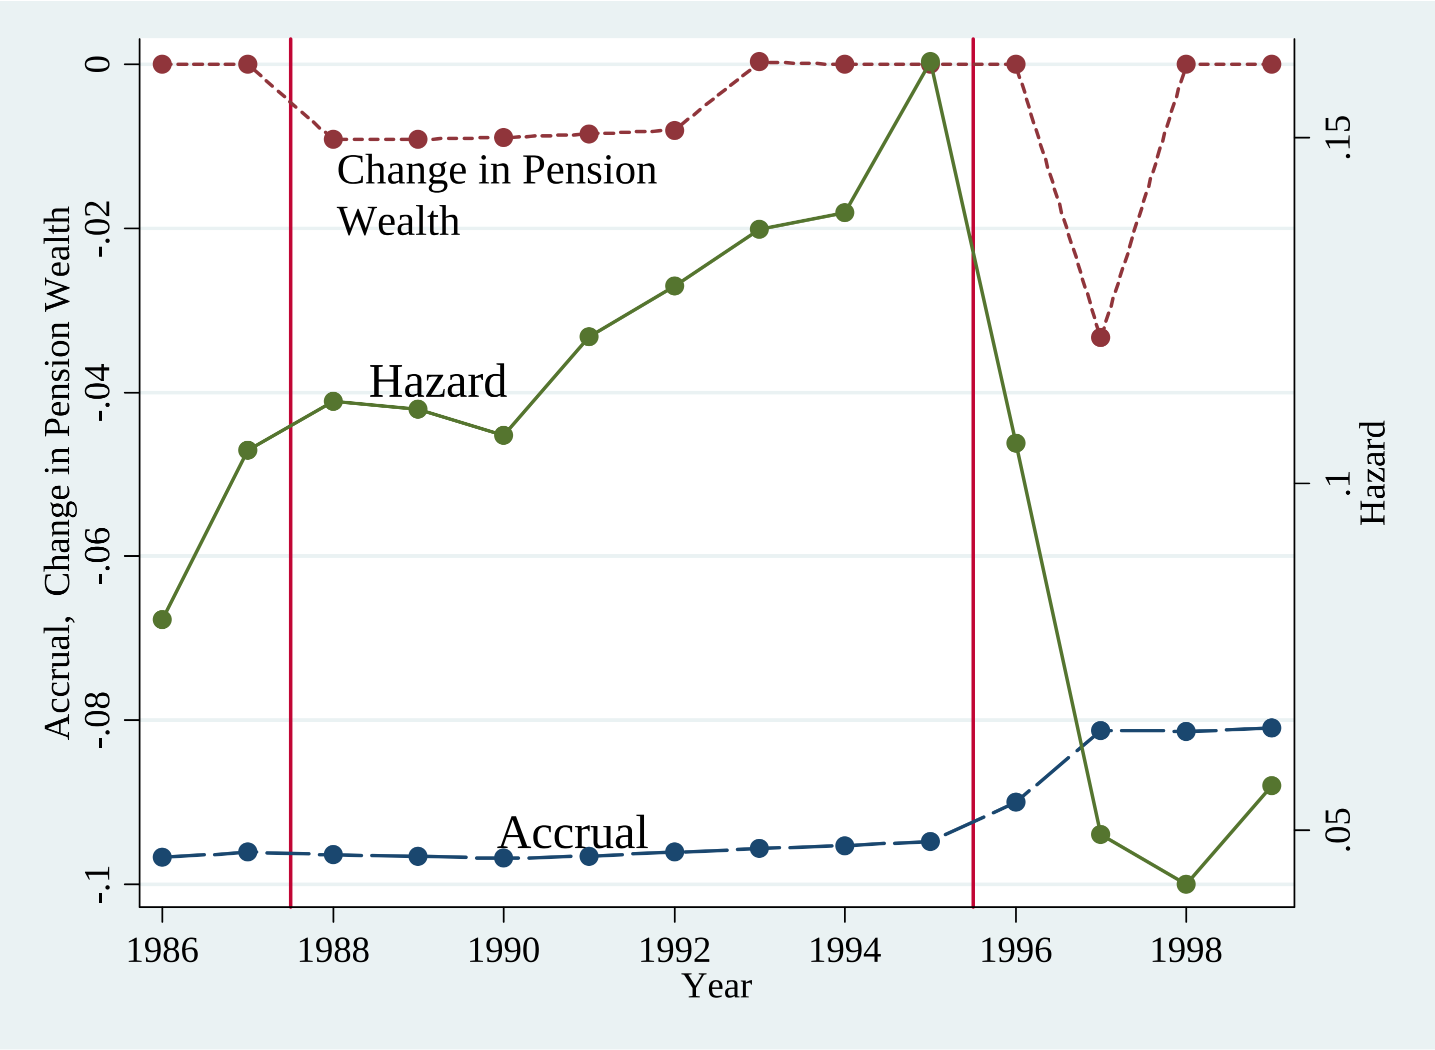

Note: Legislated effects are for individuals at age 55. The change in pension wealth is computed relative to the previous year’s legislation.

Figure 9: The Retirement Decision

Determination of the Reservation Work Utility

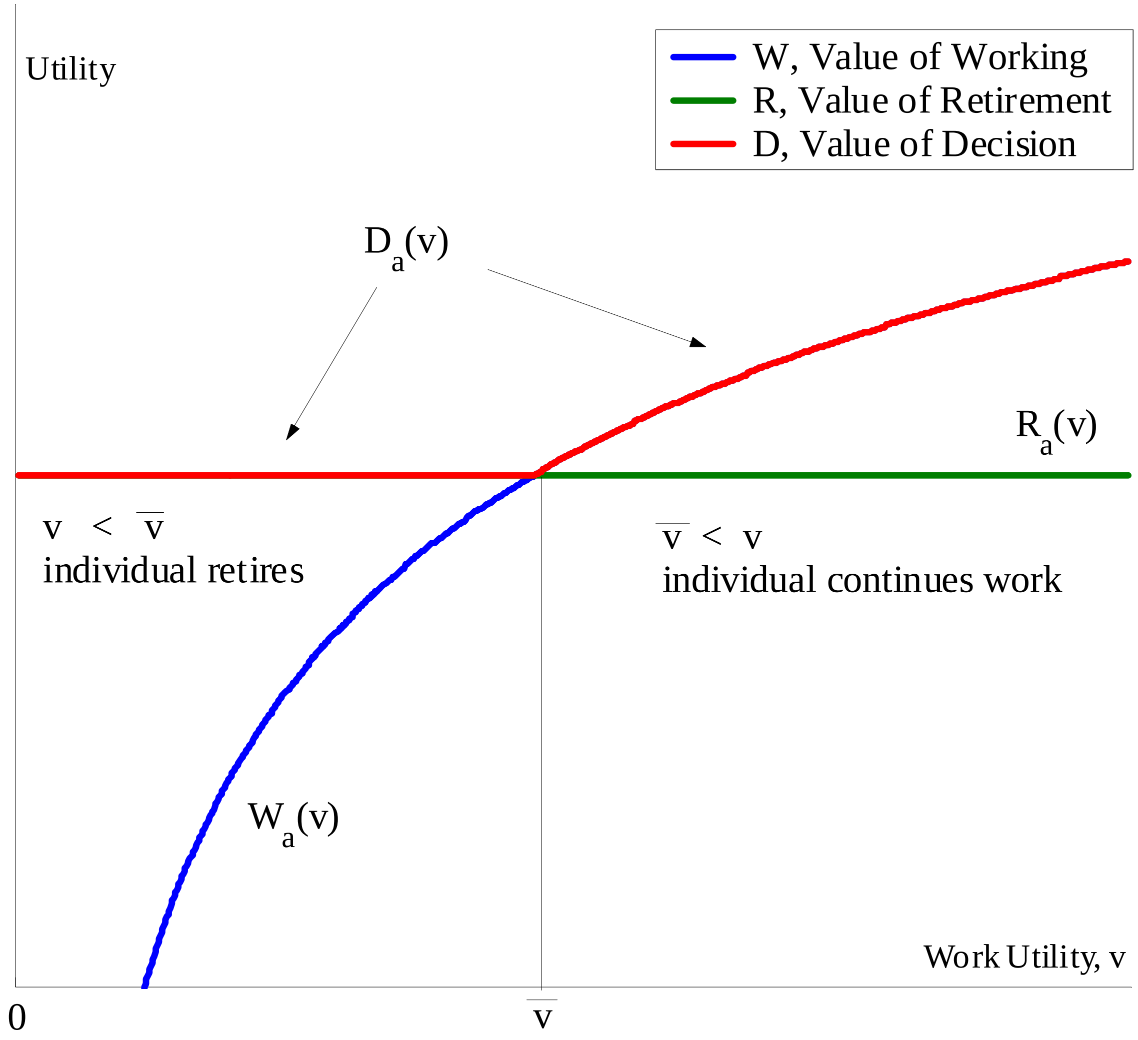

Figure 10:
Decomposing the Labor Supply Response to Benefits

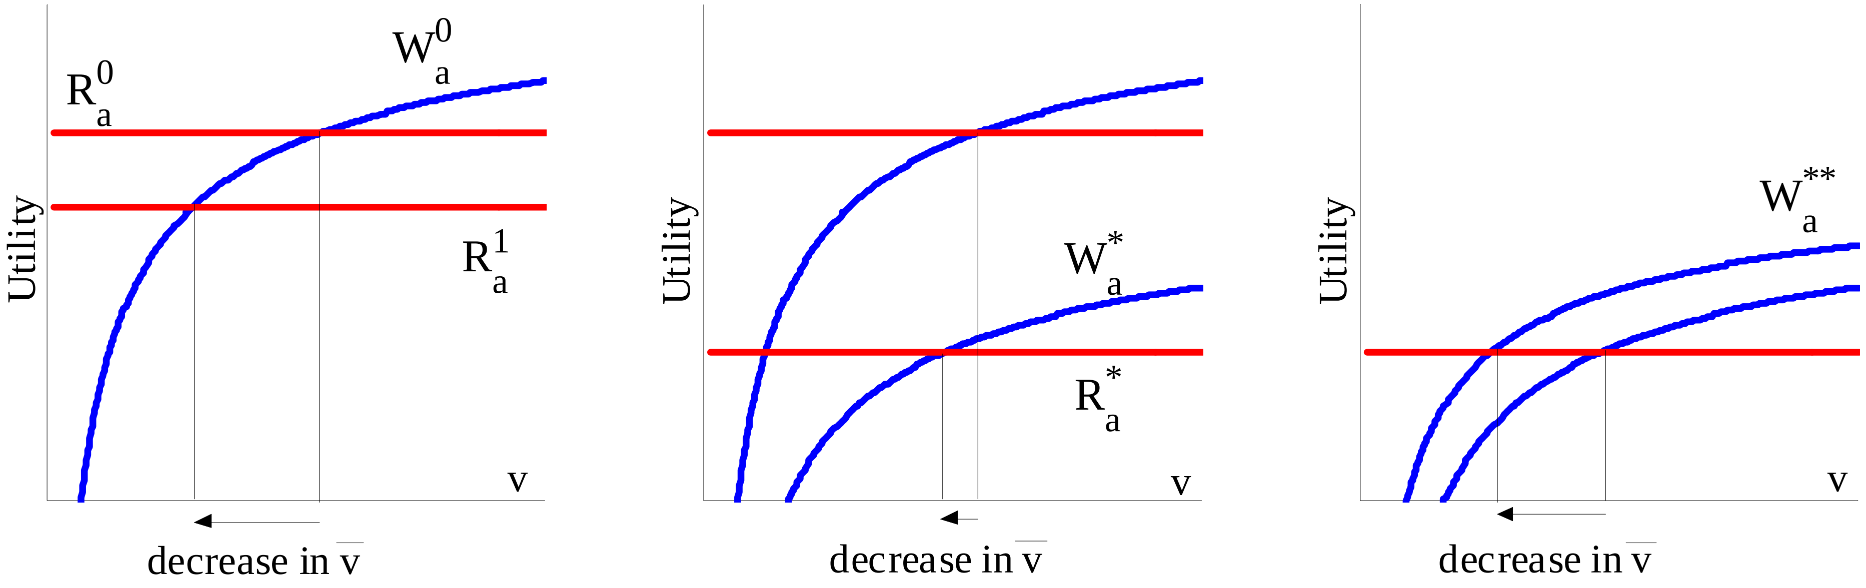

A. Overall Response to Decrease in Benefits

B. Wealth Effect from Decrease in Wealth

C. Price Effect from Increase in Wage

Figure 11A:

Retirement Hazard Rates by Age

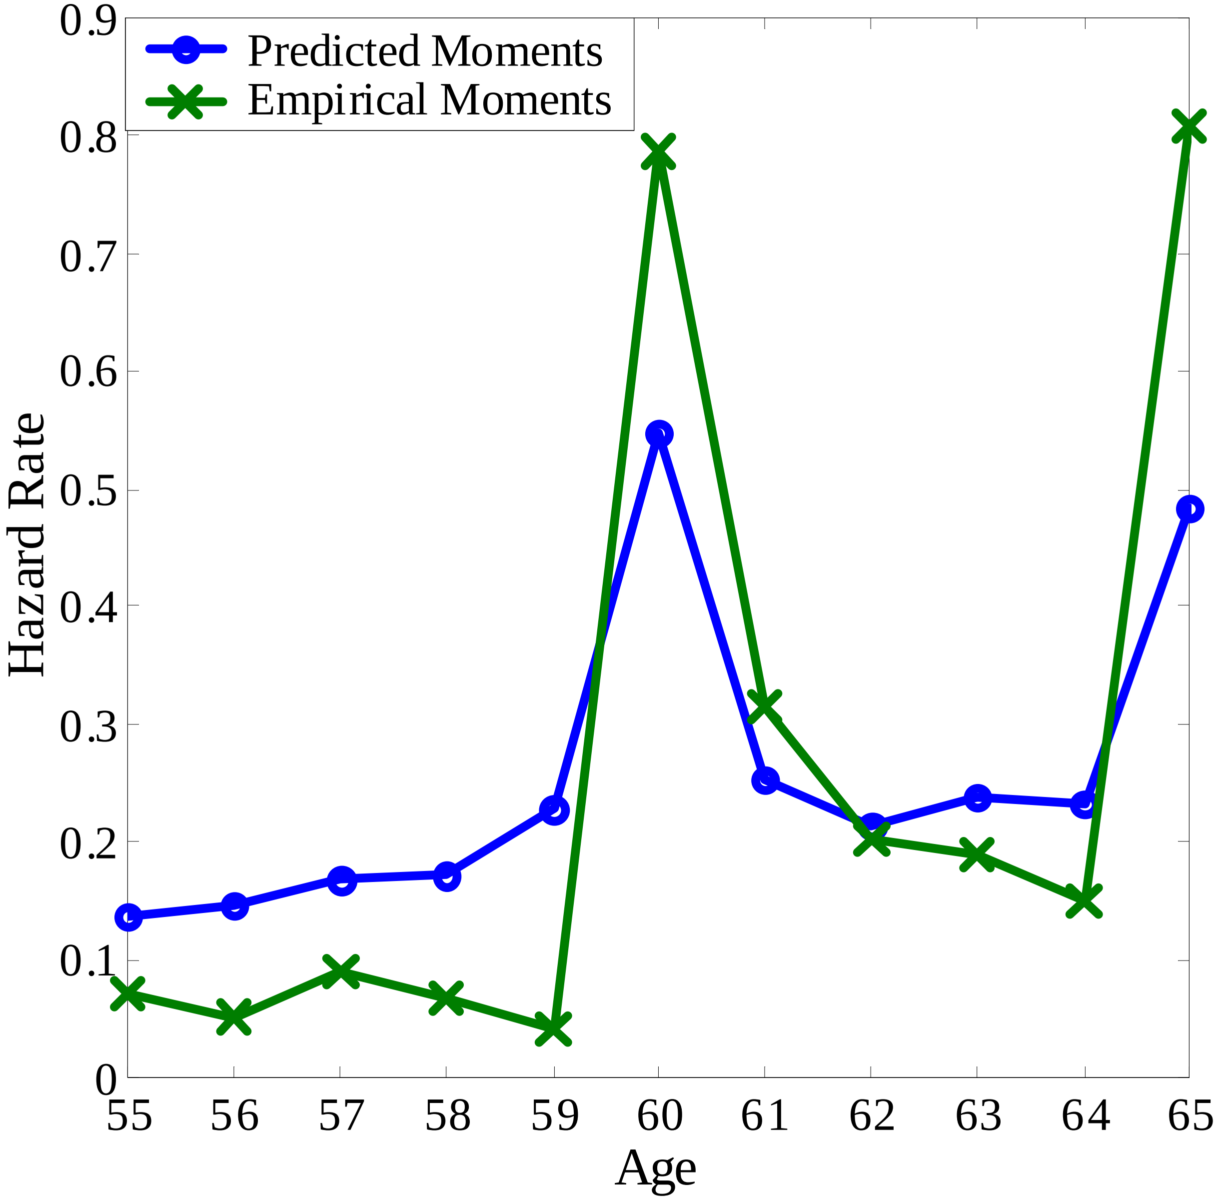

Figure 11B:

Retirement Probabilities by Tenure

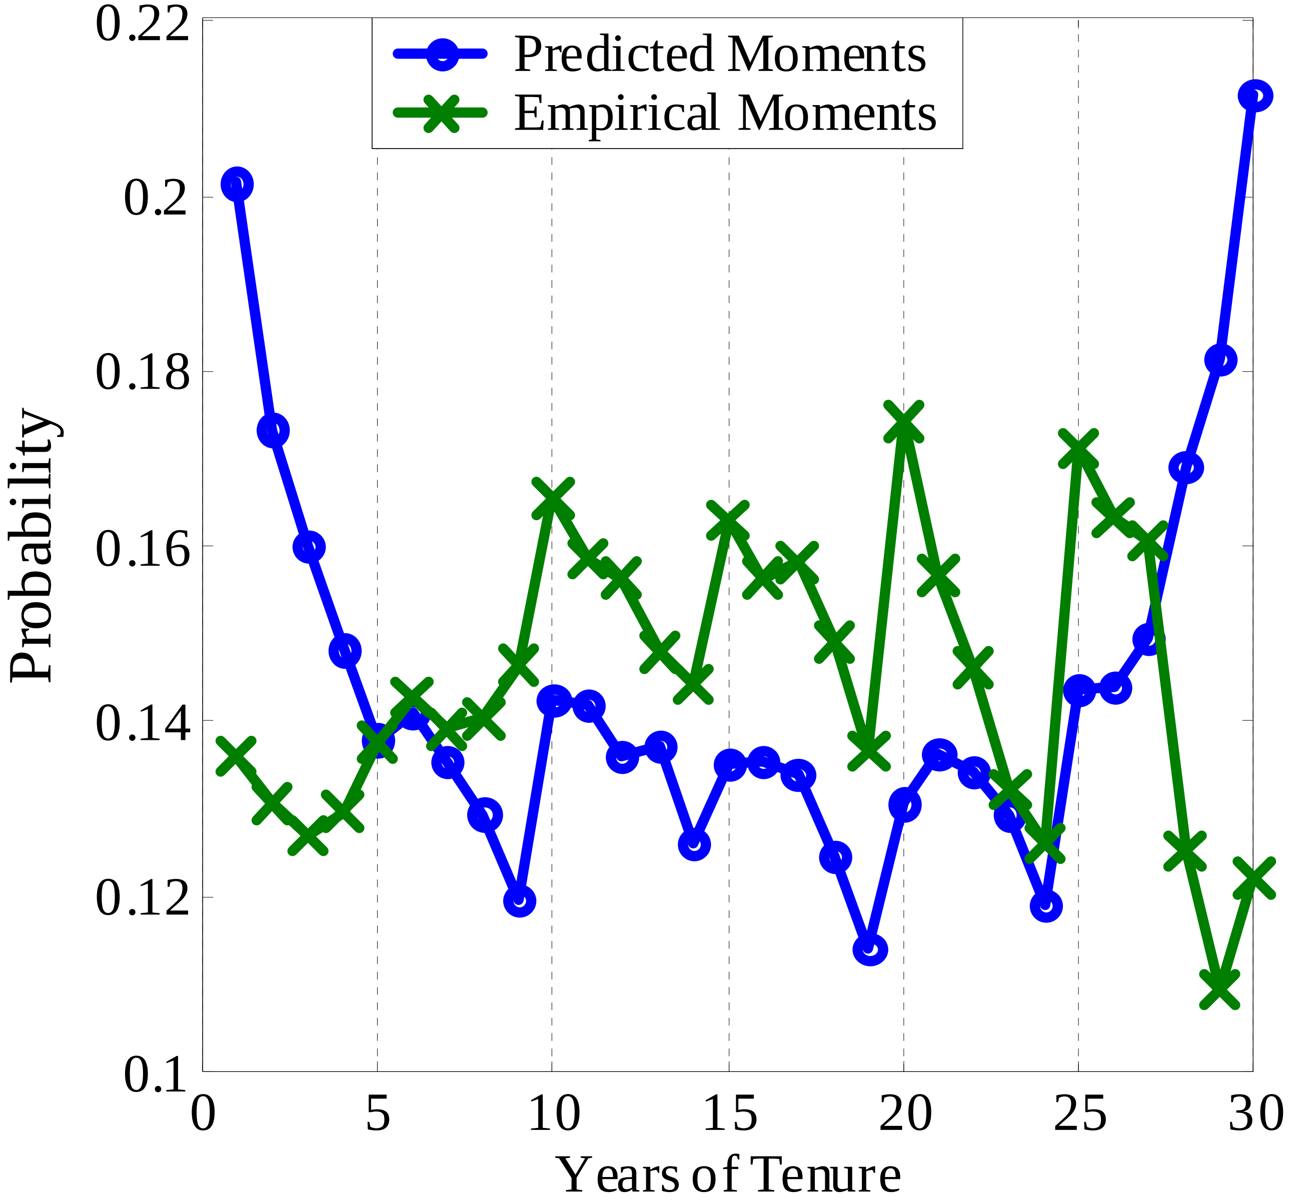

Figure 12:

Shift in Work Utility at Ages 55-59, Retirement Hazard Rates by Age

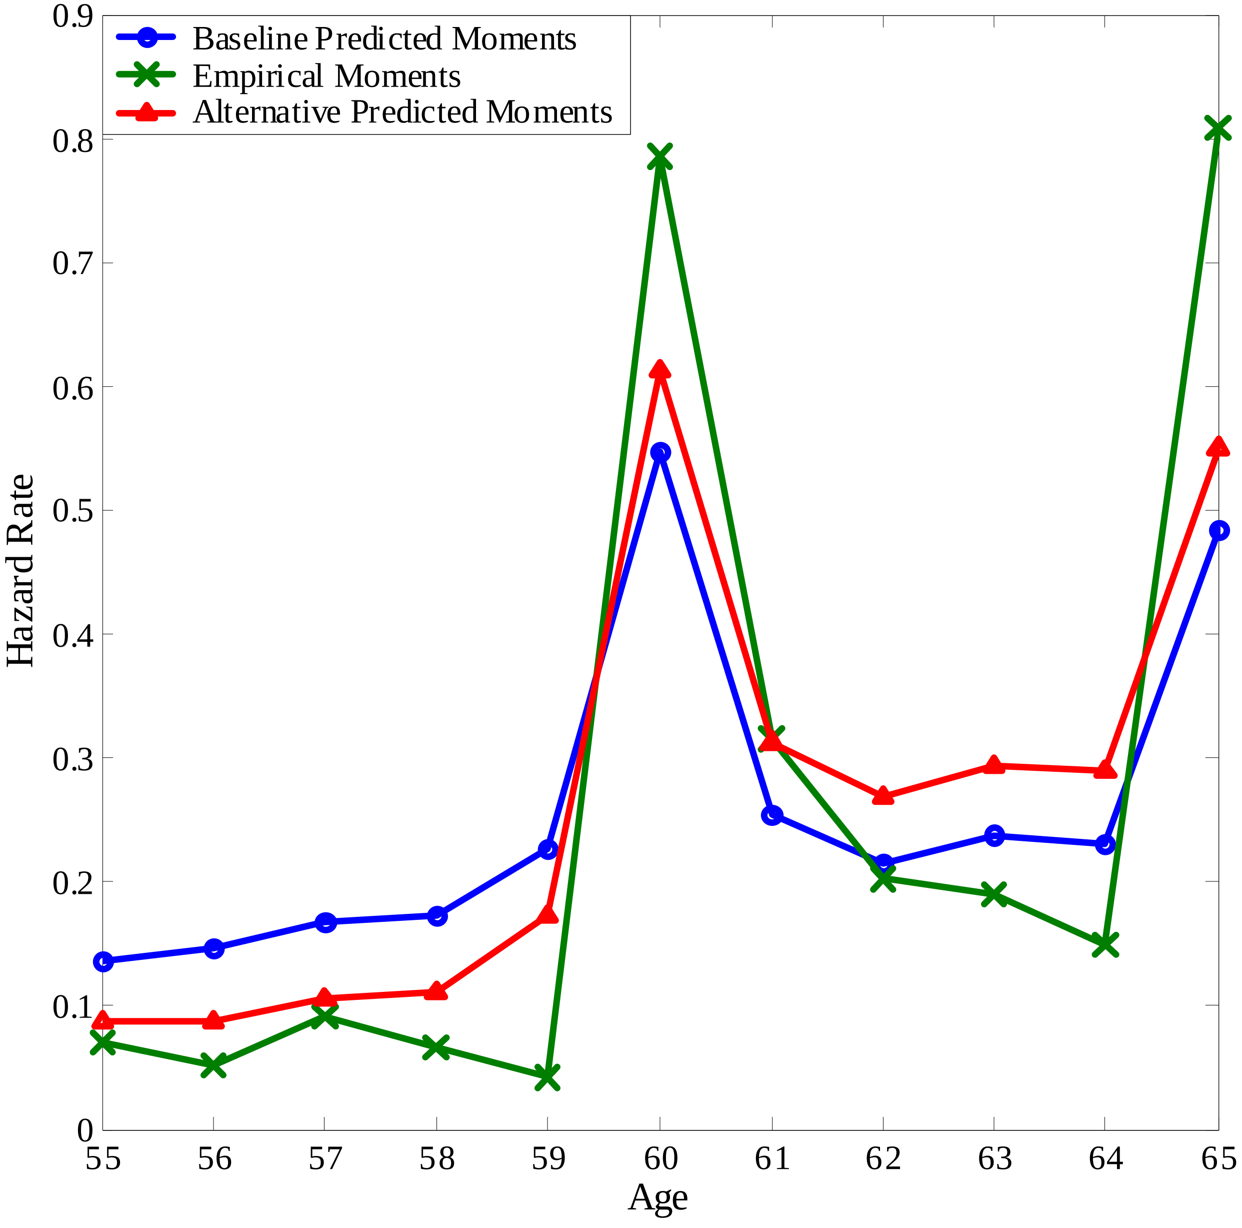

Figure 13A:

Additively Separable Model vs. Non-separable Model Retirement Hazard Rates by Age

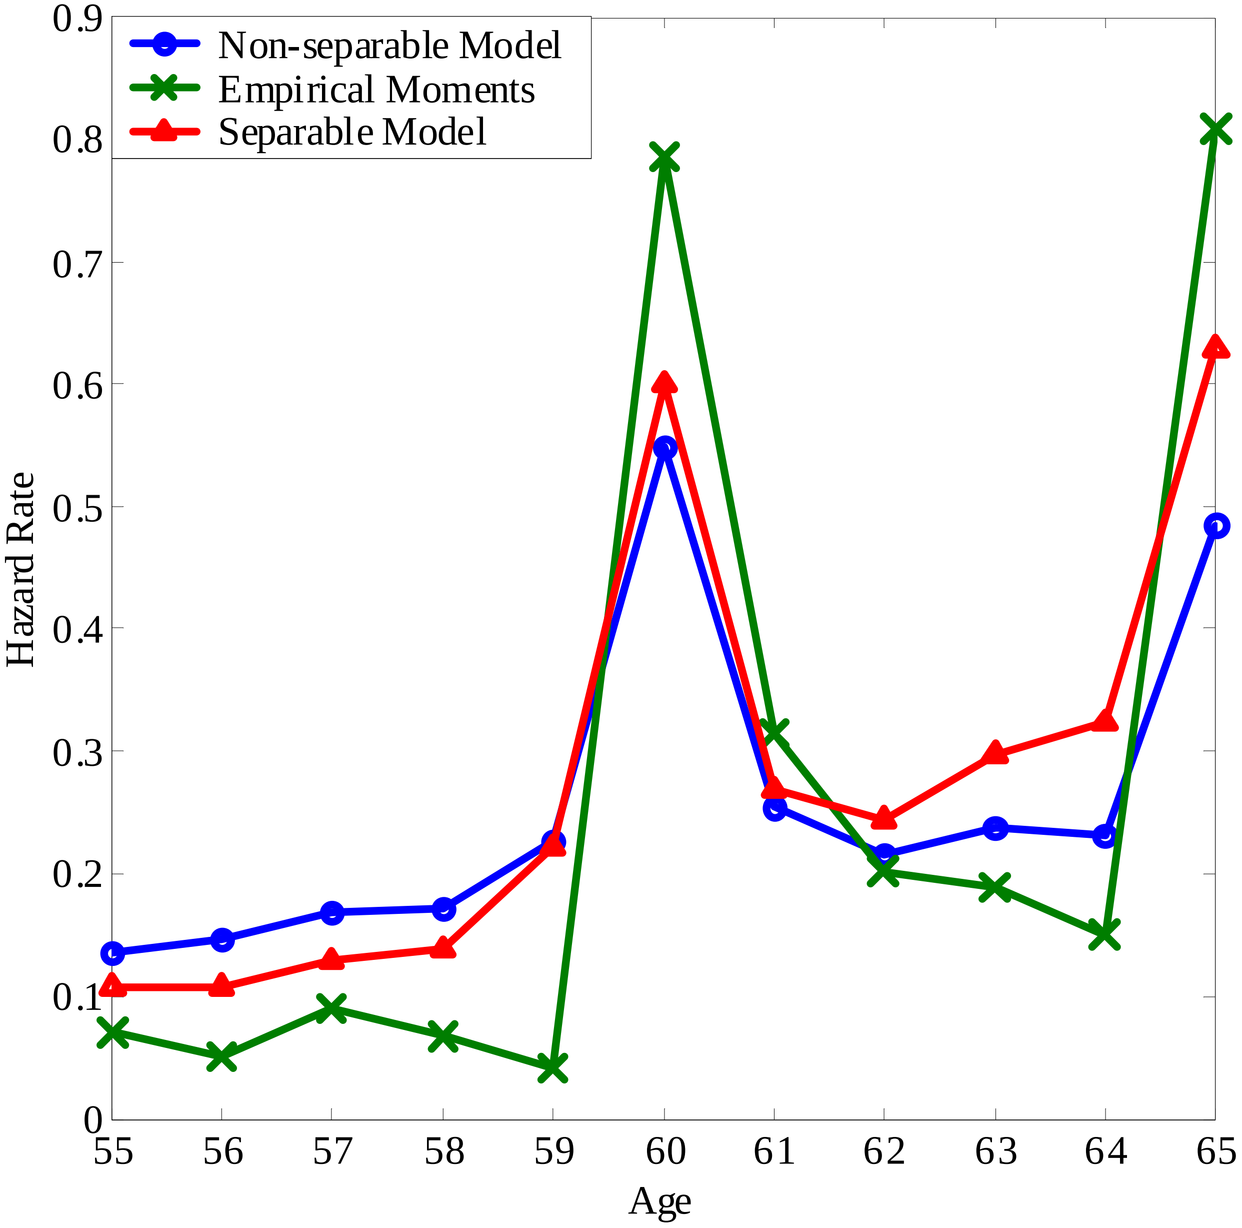

Figure 13B:

Additively Separable Model vs. Non-separable Model Retirement Probabilities by Tenure

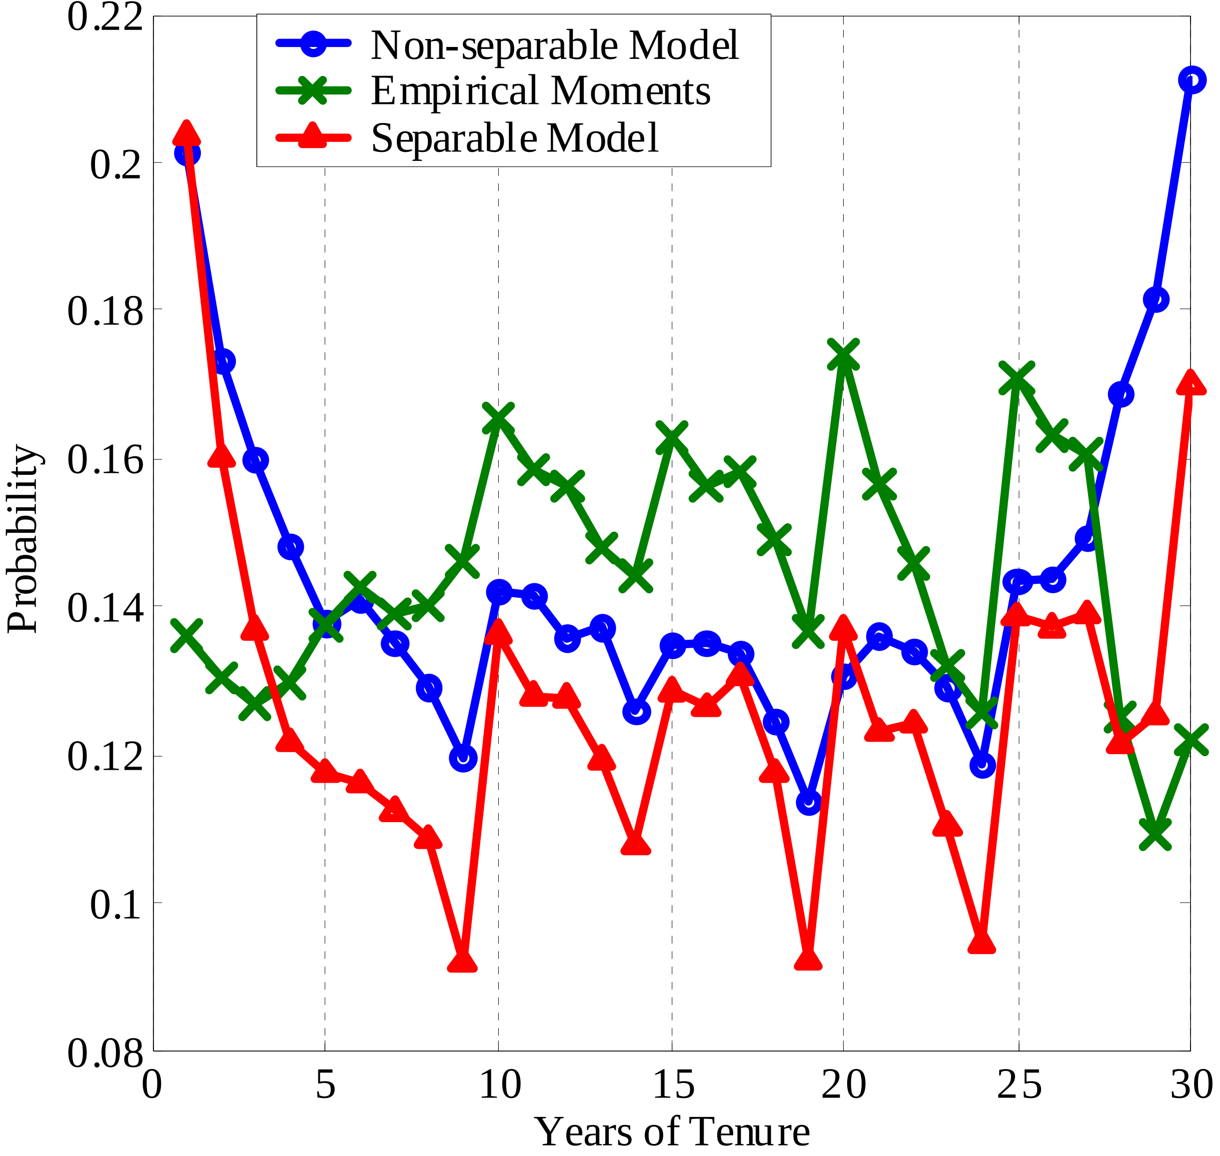

Figure 14: Hypothetical Reforms

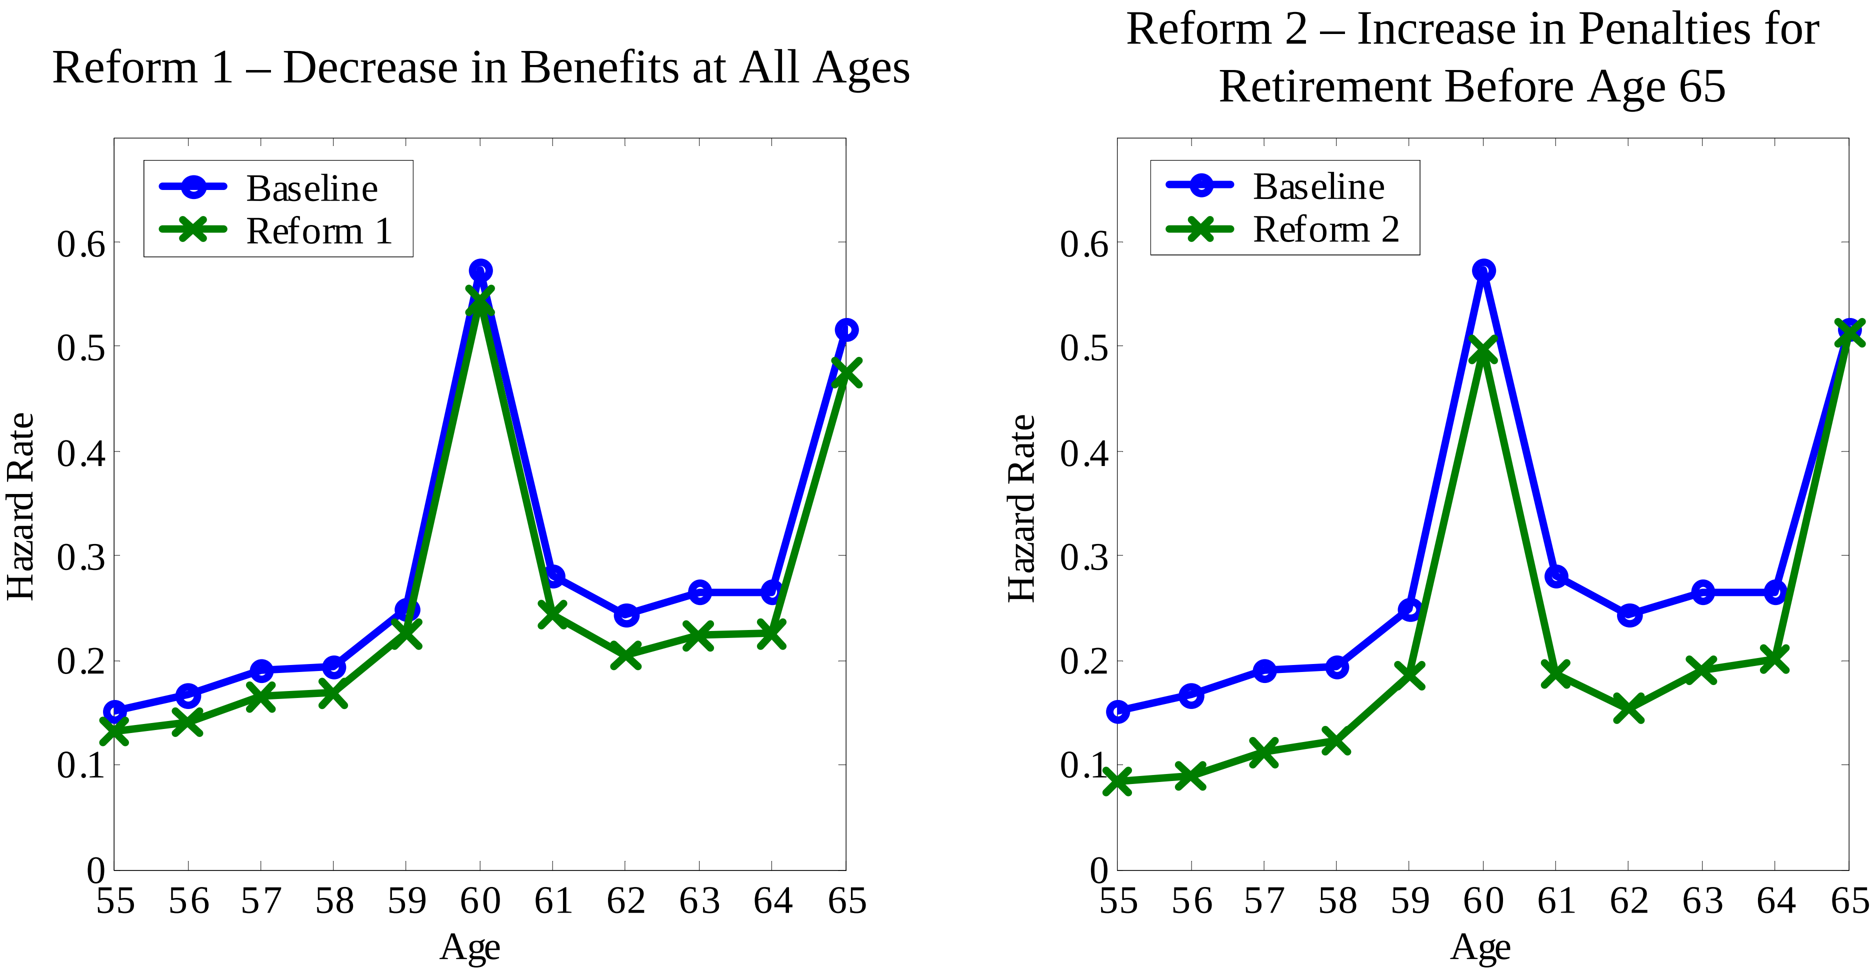

Reform 1 – Decrease in Benefits at All Ages

Reform 2 – Increase in Penalties for Retirement Before Age 65

Figure 15:

Reform 3 – Increase in Benefits at Age 65