Pension Benefits & Retirement Decisions: Income vs. Price Elasticities*

Dayanand Manoli[†] UCLA
Kathleen Mullen RAND
Mathis Wagner Collegio Carlo Alberto

September 28, 2009

Abstract

We separately identify the income and price elasticities in individuals' retirement decisions. We develop a dynamic model of retirement decisions that establishes a precise relationship between a structural parameter and these elasticities. Next, we estimate these elasticities in a proportional hazard specification that is independent of any assumptions of the dynamic model. The elasticities are identified and estimated using policy variation from pension reforms in Austria and administrative data from the Austrian Social Security Database. We then use these elasticities to estimate the dynamic model and examine the labor supply and welfare consequences of potential social security reforms.

*We thank Tajendra Singh and Shao-Ching Huang for excellent computing assistance. We are also grateful for helpful comments from Alan Auerbach, Sandy Black, David Blau, Leah Boustan, Moshe Buchinsky, Raj Chetty, Julie Cullen, Alan Gustman, Justine Hastings, Michael Hurd, Pierre-Carl Michaud, Emmanuel Saez and numerous seminar participants.
†Corresponding author. E-mail: dsmanoli@econ.ucla.edu. Phone: (310) 794-6617.

1 Introduction

In countries around the world, there is increasing pressure for social security reform.[1] Designing effective social security reform requires understanding how changes in retirement benefits affect individuals' retirement decisions. However, the precise mechanisms through which retirement benefits affect retirement decisions have remained unclear. In this paper, we identify the distinct channels through which retirement benefits affect individuals' retirement decisions. Specifically, we exploit policy variation in individuals' retirement benefits to identify and estimate the income and price elasticities in individuals' retirement decisions. We then demonstrate what these elasticities imply for standard economic models of retirement decisions and for the welfare impacts of potential social security reforms.

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. The price effect refers to changes in behavior due to changes in marginal incentives for continued work. Distinguishing between these income and price effects is relevant for pension reform; the magnitude of the price effect relative to the income effect reflects how much effective marginal tax rates from benefit schedules drive retirement behavior rather than the levels of benefits. We exploit variation from multiple pension reforms in Austria and use administrative data from the Austrian Social Security Database to identify and estimate these income and price elasticities. We estimate relatively larger price elasticities and smaller income elasticities indicating that policymakers may want to be particularly conscious of effective marginal tax rates imposed by benefit schedules. Additionally, these estimated elasticities imply that the welfare losses from potential benefits reductions would not be as severe as some previous studies have suggested.

Our analysis contributes to the literature in the following ways. First, we develop a dynamic programming model of retirement decisions that establishes a precise relationship between a structural parameter, the coefficient of relative risk aversion, and the income and price elasticities from changes in retirement benefits. Second, we identify and estimate these income and price elasticities using policy variation and administrative data. Third, we demonstrate how these income and price elasticities can be directly used to estimate a structural model of retirement decisions. Intuitively, observed responses to real pension reforms can be used to estimate these labor supply elasticities, and then the labor supply

[1] See OECD (2007).

elasticities can be used to inform economic models of retirement that are intended to simulate responses to hypothetical pension reforms. Fourth, we demonstrate that the model estimated using labor supply elasticities has different policy implications than some earlier estimated or calibrated models. In particular, the estimation based on labor supply elasticities implies that the marginal utility of consumption diminished more slowly than some earlier studies have estimated or calibrated, and as a result, there are different labor supply and welfare consequences to hypothetical pension reforms.

The dynamic programming model demonstrates a specific relationship between the rate at which the marginal utility of consumption diminishes and the ratio of income and price elasticities in individuals' retirement decisions.[2] In a standard economic model, the coefficient of relative risk aversion (γ) is the structural parameter that reflects how quickly the marginal utility of consumption diminishes. To see the intuition behind the relationship between γ and the income and price elasticities, assume consumption and leisure are known to have zero complementarity in the utility function and consider an example in which an individual receives an exogenous cash grant. The individual can use the additional income to finance additional utility in two ways; first, he could buy additional consumption and second he could buy additional leisure. If one observes that the individual chooses to buy additional leisure, then one can infer that the marginal utility of consumption must decline relatively quickly (i.e. γ must be relatively high) since otherwise the individual would have purchased additional consumption. Thus, a relatively larger income elasticity implies a relatively higher γ, conditional on some a priori knowledge regarding the complementarity between consumption and leisure. This relationship motivates the uses of these labor supply elasticities as moments to directly match in the estimation of the model since the elasticities play a key role in estimating γ and hence other remaining parameters as well.

The empirical analysis is presented in two parts. The first part focuses on the identification of the income and price elasticities based on policy variation from five pension reforms in Austria between 1984 and 2003. Using administrative, social security records data on over 250,000 private sector employees in Austria, we define the income and price measures using social security wealth (the present discounted value of pension benefits) and the one-year accrual (the expected change in social security wealth from delaying retirement by one year). The pension reforms create several independent changes in these measures thereby allowing for separate identification of the respective elasticities. We present multiple pieces

[2] Chetty (2006) provides a discussion of the general relationship between labor supply elasticities and the coefficient of relative risk aversion.

of evidence that highlight individuals' relative sensitivity to changes in marginal incentives for continuing work and their insensitivity to changes in the levels of wealth based on retirement benefits. We start with nonparametric graphical evidence which illustrates individuals' responsiveness to changes in the one-year accrual, and their lack of responsiveness to changes in social security wealth. Using a proportional hazards specification, we quantify these sensitivities and estimate income and price elasticities of 0.43 and 2.90 respectively. The interpretation of these elasticities is that a 1% increase in social security wealth at a given age increases the retirement hazard at that age by 3.34% ($= 0.43 + 2.90$); roughly 13% of this total elasticity $\left(\frac{0.43}{0.43+2.90}\right)$ is due to the income elasticity from increasing lifetime income while the remaining 87% is due to the price elasticity from reducing the effective wage for continuing work. The relatively smaller income elasticity implies a relatively low coefficient of relative risk aversion. This proportional hazards specification is entirely independent from our dynamic model of retirement decisions and hence it is free of any distributional or functional form assumptions of that economic model.

The second part of the empirical analysis focuses on structurally estimating the dynamic model of retirement decisions using the elasticity results from the first part of the empirical analysis. Motivated by the relationship between the income and price elasticities and the coefficient of relative risk aversion that is implied by the model, we estimate the model using a method of simulated moments (MSM) estimation strategy in which the labor supply elasticities are included as moments to match directly. Specifically, for each iteration of the estimation, the algorithm estimates the proportional hazards specification using simulated data from the model. Therefore, the MSM estimation algorithm seeks structural parameters that match the elasticities based on simulated data from the model to the actual income and price elasticities from the first part of the empirical analysis. In addition to these elasticities, the estimation matches the retirement hazard rates conditional on age. Using this strategy, we estimate the coefficient of relative risk aversion to be 0.71 in the baseline specification of the model. This relatively low estimate for γ is consistent with what is implied by the elasticities from the proportional hazards specification; indeed we use the estimated structural model to demonstrate that a higher γ would imply, contrary to the data, relatively larger income elasticities and smaller price elasticities in the proportional hazards specification. Other estimated specifications yield similar degrees of relative risk aversion.

Lastly, we use the estimated structural model to study the labor supply and welfare consequences of a variety of hypothetical pension reforms. The results from the policy simulations generally highlight the consequences of a low degree of relative risk aversion for policy. In particular, smaller income effects are evident from relatively smaller changes in the simulated retirement hazard rates in response to decreases in benefits that reduce total income. Additionally, larger price effects are evident from changes in the hazard rates in response to benefit changes that increase effective wages for continuing to work. In regard to welfare, we measure welfare using expected utility at the initial age in the model. The results from the simulated reforms indicate that reductions in the generosity of benefits do not lead to large reductions in welfare since a low degree of relative risk aversion implies that marginal utility of consumption does not increase sharply with reductions in income. Furthermore, the low γ implies that the variance of expected utility rises relatively quickly when reforms facilitate retirement at older ages using changes in the across-age schedule of benefits. Intuitively, when individuals work longer, they accumulate more work disutility but also more income. When the marginal utility of consumption is nearly constant, consumption rises more so for higher wage earners than for lower wage earners, and thus the additional work disutility is offset more so for higher wage earners.

The remainder of the paper is organized as follows. Section 2 discussed the related literature. Section 3 develops the dynamic programming model of retirement decisions. Second 4 describes the institutional background of the Austrian pension system, the pension reforms and the data used in the empirical analysis. Section 5 presents the first part of the empirical analysis to identify and estimate the income and price elasticities. Section 6 presents the second part of the empirical analysis which structurally estimates the model in section 3 using the elasticity results from section 5. The policy simulations from the estimated structural model are discussed in section 7 and section 8 concludes.

2 Related Literature

Our focus on income and price elasticities from retirement benefits and our estimates of these elasticities relate to several previous studies. Recent research has emphasized the separate identification of income and price effects of benefits from other social insurance programs. These studies have generally found larger income effects and smaller price effects in responses to unemployment benefits (see Chetty (2008) and Card, Chetty and Weber (2007)), disability benefits (see Autor and Duggan (2007)) and health benefits (see Nyman (2003)). The contrast between our results based on pension benefits and these previous results based on other social insurance programs indicates that, in future work, it could be useful to determine what specific factors lead to different income and price elasticities from different social insurance programs. For example, Aguiar and Hurst (2005) discuss differences between retirement and unemployment based on differences in the anticipation of these events; it could be interesting to pursue this difference further and explicitly relate it to the differences in estimated at retirement and unemployment.

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. Consistent with our results, Friedberg presents structural estimates of sizeable uncompensated elasticities with relatively small income elasticities. Other studies have separately focused on either wealth effects or incentive effects from social security benefits. Though it is difficult to map some of the estimates into elasticities that are directly comparable to our results, Costa (1995), Krueger and Pischke (1992), Holtz-Eakin, Joulfaian and Rosen (1999) and Brown, Coile and Weisbenner (2006) present evidence of small wealth effects which is broadly consistent with our results. Additionally, Samwick (1998), Asch, Haider and Zissimopoulos (2005), Song and Manchester (2007), Haider and Loughran (2008) and Liebman, Luttmer and Seif (2008) provide evidence of significant incentive effects of social security benefits which is also consistent with our results.

Our structural estimates also relate to previous results in the literature. In particular, the structural estimation in this paper yields an estimate for γ that is lower than values estimated or calibrated in some earlier studies. French (2005) estimates γ to be between 2.2. and 5.1 while van der Klaauw and Wolpin (2008) estimate γ to be roughly 1.6. In terms of calibration, Blau (2008) calibrates γ to be 2 while Hubbard, Skinner and Zeldes (1995) use a preferred value of 3 for γ. Nonetheless, the estimates of γ from this current analysis are consistent with previous estimates in which γ is identified based on labor supply elasticities. Chetty (2006) surveys estimates of wage and income elasticities from thirty-three previous studies and finds that the mean of the implied values of γ is 0.71 with a range of 0.15 to 1.78.[3]

[3] These values from Chetty (2006) are reported for the case of additive utility. This is the case that corresponds to the model estimated in this study.

3 Theoretical Foundations

In this section, we develop a dynamic model of retirement decisions with uncertainty relating to mortality and job separations.[4] Accounting for these sources of uncertainty in a dynamic setting is an important component of capturing uncertainty relating to retirement benefits. The intuition behind the model is as follows. In each period, an employed individual must choose whether to retire or whether to continue working. A period in the model corresponds to an individual's age. When making this decision at the beginning of a period, the individual knows his assets, retirement benefits, and wage. If he chooses to retire, the individual receives his annuitized retirement benefits and faces no remaining uncertainty from the labor market. Based on his assets and retirement benefits, the retired individual chooses his consumption optimally. If an employed individual chooses to continue working, he receives his wage and his consumption is then determined optimally based on his savings and wage income. The working individual also takes into account the expected continuation value from being able to make a retirement decision at future ages.

3.1 The Model

Consider first the optimization problem for an individual who has chosen to retire. Let $R_a(A_a)$ denote the value of retirement at age a for an individual with assets A_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 that defines the value of retirement:

$$\begin{aligned} R_a(A_a, y_a^R) &= \max_{\{c_t^R\}_{t=a}^T} u(c_a^R) + \sum_{t=a+1}^{T} \beta^{t-a} \pi_{t|a} u(c_t^R) \\ & \quad s.t. \sum_{t=a}^{T} (\tfrac{1}{1+r})^{t-a} c_t^R = A_a - \kappa \mathbf{1}(a < a_{old}) + \sum_{t=a}^{T} (\tfrac{1}{1+r})^{t-a} y_a^R. \end{aligned}$$

The variables r, β and $\pi_{t|a}$ respectively denote the interest rate, the discount factor and the probability of survival to age t conditional on survival to age a. The function $u(.)$ captures utility over consumption with $u' > 0$ and $u'' < 0$. The maximization reflects that the individual chooses his consumption in each period. The term y_a^R denotes the individual's retirement benefits at age a. Based on the institutional setting in Austria which we discuss

[4] The dynamic model that we develop is closely related to previous work in the literature. See Stock and Wise (1990), Lumsdaine, Stock and Wise (1992), Berkovec and Stern (1991) and Rust and Phelan (1997).

in more detail in the next section, these benefits are based on an annual payment from a government-provided pension.[5] The term κ denotes a claiming cost. If the individual retires at an early age when individuals are only eligible for disability pensions, he must pay the one-time cost of claiming κ. This cost can be interpreted as the cost of visiting a physician to be classified as disabled. After age a_{old}, individuals are eligible for old-age pensions and therefore do not face this claiming cost.

Next, consider the problem facing an individual who has chosen to work. As in the case of retirement, the individual must choose his consumption optimally. The optimization problem in the case of continuing to work differs from that in the case of retirement in the following respects. First, the working individual must take into account his disutility of work denoted by v_a. Work disutility is increasing with age and each individual is assumed to know the profile of his work disutility across age with certainty. Specifically, prior to facing the first retirement decision, v_0 is drawn for each individual from a distribution $\Psi(v)$ defined over $(0, \infty)$. The work disutility profile across ages is then given by $v_a = v(a, v_0) \; \forall \; a$. Second, the individual's income is based on his wage income.[6] After-tax work income at age a is denoted by y_a^W. Third, the individual must take into account the continued uncertainty from the labor market. In particular, $E_a[D_{a+1}(.)]$ captures the individual's continuation value from being able to make a retirement decision in the future where the expectation takes uncertainty from mortality and job separations into account. Let $W_a(A_a, y_a^W, v_a)$ denote the value of working at age a with assets A_a and work disutility v_a. This value function is defined as follows

$$\begin{aligned} W_a(A_a, y_a^W, v_a) \quad &= \quad \max_{c_a^W} u(c_a^W) - v_a + \beta \pi_{a+1|a} E_a[D_{a+1}(A_{a+1}, y_{a+1}^R, y_{a+1}^W, v_{a+1})] \\ & \qquad s.t. \; c_a^W + \tfrac{A_{a+1}}{1+r} = y_a^W + A_a. \end{aligned}$$

The individual's consumption while working is denoted by c_a^W. Savings for next period, A_{a+1}, are determined based on the individual's current savings and wage income net of current consumption. The value function $D_a(...)$ captures the value of being in the labor market at age a and having the decision between retiring or continuing to work. When de-

[5] In addition to the government-provided pension benefits, retirement benefits also include employer-provided severance payments. In particular, the Austrian pension system mandates that employers make one-time lump-sum payments to employees at the time of their retirements. The amounts of these severance payments is determined based on individuals' tenure and wages at retirement. For this reason, tenure is also included as a state variable though the notation is suppressed here.

[6] Wage income also includes any severance pay if the individual was separated from his previous employer and had sufficient tenure to qualify for a severance payment.

ciding between retirement and work, the individual simply chooses the option that presents the highest value,

$$D_a(A_a, y_a^R, y_a^W, v_a) = \max_{retire,work} \{R_a(A_a, y_a^R), W_a(A_a, y_a^W, v_a)\}.$$

In regard to heterogeneity, work disutility v is 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, when considering his retirement benefits at any given age, we assume that the individual forecasts benefits at future potential retirement ages based on the current year's legislation. Thus, calendar year enters the value functions as an implicit state variable that determines the legislation under which benefits are computed. In this setting, pension reforms then correspond to unanticipated changes in the legislation and hence unanticipated changes to benefits at current and future potential retirement ages.

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.[7] Formally, the reservation level is defined as the preference for work that leaves the individual indifferent between retiring and continuing to work

$$R_a(A_a, y_a^R) = W_a(A_a, y_a^W, \bar{v}_a) \Rightarrow \bar{v}_a = \bar{v}_a(A_a, y_a^R, y_a^W,).$$

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

$$\text{Retire at age } a \text{ if } v \geq \bar{v}_a.$$

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.

3.2 Labor Supply & Risk Aversion

This section develops the intuition relating labor supply to risk aversion in the dynamic model of retirement decisions. Chetty (2006) provides an extensive discussion of this re-

[7] Since the value functions R_a and W_a depend on age through the different vectors of survival probabilities at each potential retirement age, the reservation work disutility will also depend on age.

lationship. In particular, Chetty emphasizes that the coefficient of relative risk aversion, γ, can be identified based on two components: (1) the ratio of wealth and price elasticities in labor supply and (2) the degree of complementarity between consumption and labor in the utility function. In our setting, we have assumed additively time separable utility. With this assumption, one can fix the degree of complementarity in the per-period utility function to identify γ in terms of labor supply elasticities. We have assumed that work disutility is additively separable from the utility over consumption, so the complementarity between consumption and work disutility in the utility function is assumed to be 0. We discuss this assumption in greater detail in section 2.3.

The wealth and price effects from retirement benefits can be seen by examining how the individual's retirement strategy responds to changes in benefits, wages and wealth at a given age a. Differentiating the equation for $\bar{v}_a$ with respect to y_a^R, y_a^W and A_a yields

$$\begin{array}{rcl} [y_a^R] & : & \frac{d\bar{v}_a}{dy_a^R} = -u'(c_a^R) < 0 \\ [y_a^W] & : & \frac{d\bar{v}_a}{dy_a^W} = u'(c_a^W) > 0 \\ [A_a] & : & \frac{d\bar{v}_a}{dA_a} = u'(c_a^W) - u'(c_a^R) \leq 0 \end{array}$$

$$\Rightarrow \frac{d\bar{v}_a}{dy_a^R} = \underbrace{\frac{d\bar{v}_a}{dA_a}}_{\text{wealth effect}} - \underbrace{\frac{d\bar{v}_a}{dy_a^W}}_{\text{price effect}}.$$

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. 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 wealth 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. These conditions can be combined to decompose the labor supply response to a change in retirement benefits into a wealth effect and a price effect. The wealth effect reflects that the change in benefits creates additional wealth for the individual and the price effect reflects that the change in retirement benefits creates a change in the effective wage from working.

Given this decomposition, risk aversion can now be related to labor supply elasticities.

In particular, taking a second order Taylor approximation, $u'(c_a^R) = u'(c_a^W) + u''(c_a^W)(c_a^R - c_a^W)$, the coefficient of relative risk aversion ($\gamma = \frac{cu''(c)}{u'(c)}$) can be written as

$$\gamma = \left(\frac{c_a^R - c_a^W}{c_a^W}\right)^{-1} \left(\frac{y_a^W}{A_a}\right) \left(\frac{\varepsilon^A}{\varepsilon^w}\right)$$

where $\varepsilon^A = \frac{A_a}{\bar{v}_a}\frac{d\bar{v}_a}{dA_a}$ denotes the wealth elasticity and $\varepsilon^w = \frac{y_a^W}{\bar{v}_a}\frac{d\bar{v}_a}{dy_a^W}$ denotes the price (wage) elasticity.[8][9] This equation highlights that a larger wealth elasticity corresponds to a higher degree of curvature. Intuitively, if an individual changes his labor supply significantly upon receipt of an exogenous cash grant, this indicates that the marginal utility of consumption must decline relatively quickly, since otherwise the individual would have increased his consumption instead. In the Labor Supply & Risk Aversion Appendix below, we discuss this relationship further and emphasize: (1) the distinction between the within-age consumption difference $\frac{c_a^R - c_a^W}{c_a^W}$ and the intertemporal (across-age) consumption drop at retirement; (2) the roles of time separability and additive separability assumptions in allowing for identification of γ based on labor for participation outcomes; and (3) heterogeneity.

4 Institutional Background & Data

This section discusses the Austrian pension system and the pension reforms that create variation used for identification of elasticities and structural parameters. After discussing the institutional setting, we describe the data and construction of key variables used in the empirical analysis.

[8] If $A_a = 0$, γ can still be written in terms of the ratio of wealth and price effects,

$$\gamma = \left(\frac{c_a^R - c_a^W}{c_a^W}\right)^{-1} \left(\frac{d\bar{v}_a/dA_a}{d\bar{v}_a/dy_a^W}\right).$$

In this case, the wealth effect cannot be expressed as an elasticity.

[9] Given that this relationship holds at every age, it is possible to generalize the intuition to relate risk aversion to the ratio of wealth and price elasticities based on changes in the profiles of benefits, assets and wages across *multiple* ages. To consider the intuition behind a change in the profile of wealth, it is useful to consider an annuity that pays a constant amount at each age. The wealth effect based on a change in the profile of benefits across multiple ages can then be defined based on a change in the constant amount paid at every age by the annuity. The price effect is based on a change in the individual's wage profile. Chetty (2008) presents a derivation to consider changes across multiple periods using such an annuity.

4.1 The Austrian Pension System

There are two types of government-provided retirement pensions in Austria: 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 average 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. By claiming a retirement pension, the individual essentially exits the labor market.[10] 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 1A presents retirement hazard rates by age. In this figure, retirement hazard are

[10] 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.

based on claiming either an old-age pension or a disability pension. In this figure, the hazard rates spike at ages 60 and 65 at roughly 80% and 75% respectively; these ages correspond to the minimum early retirement age and the statutory retirement age respectively. To better characterize the population remaining in the labor market, Figure 1B presents the survival function. This figure also highlight the large fraction of individuals leaving the labor market at age 60 and 65 with significant declines at these ages. Importantly, this survival function also highlights a significant amount of retirement prior to age 60. In particular, just under 40% of the sample retires prior to the early retirement age by claiming disability pensions. Figure 2C focuses more directly on disability pensions by presenting the survival function for individuals who claim disability pensions. This figure further emphasizes that individuals enter disability pensions primarily before age 60 and then less so after age 60 since the minimum age for old-age pensions has been passed.

4.2 Pension Reforms

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).[11] Table 1 presents a summary of each reform. Figures 3A & B present benefits-versus-age profiles for different calendar years. These illustrations are meant to demonstrate the variation in pension benefits created by the pension reforms. To make benefits comparable across calendar years, these annual benefits are computed based on an individual who earnings the nominal equivalent of 20000 euros in 2003 in each year that he works, and nominal benefits are then adjusted by the CPI to put all benefits in 2003 euros. In these figures, full insurance years (i.e. experience = age - 15, where 15 is the age corresponding to the end of mandatory schooling) is assumed for each age. As a result of this assumption, individuals reach the maximum insurance years at age 60 (45 insurance years) leading to a kink in the benefit schedule at age 60. Additionally, to emphasize some of the early pension reforms, benefits are computed assuming that the individual has worked for the last 10 years at each age, but not prior to that. Finally, to emphasize the changes in incentives for individuals due to the pension reforms, the benefit-retirement age profiles in each year are computed for a fixed birth cohort (i.e. birth cohort = year - 60). Thus, taking the pension reforms as unanticipated,

[11] 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.

comparisons across retirement ages within a given calendar year reflect the incentives to retire at different ages, and comparisons across the calendar years reflect changes in these incentives die to the reforms.

The pension reforms generally reduced the generosity of the retirement pension system as government officials felt the pension system was not financially sustainable. This trend is evident by the downward trend in benefits across calendar years illustrated in Figures 3A & B. 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 1985 and 1988 pension reforms are illustrated in Figure 2A. In particular, benefits decrease between 1984 and 1988 due to the first increase in the length of the assessment basis from 5 to 10 years. Benefits decrease each year from 1988 to 1992 as the second increases in the assessment basis from 10 to 15 years is phased in. As illustrated, these reforms decreased the levels of benefits across potential retirement ages but left the slopes in the profiles unchanged.[12]

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.[13] 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 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.

[12]The increase in the slope of the benefits profile in 1992 at age 65 corresponds to the introduction of a bonus for retirement at the normal retirement age in 1991.

[13]Since benefits are computed assuming full insurance years at each age, the illustrated profiles already link benefits to insurance years and retirement age. As a result, the change in benefits due to this aspect of the 1993 pension reform are not evident in the figures.

Figure 2B focuses on these later reforms. 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. As illustrated in Figure 2B, the comparison between the 1996 (pre-reform) and 1998 (post-reform) benefit profiles highlights the impact of the introduction of the bonus/malus system on retirement incentives. Specifically, this reform decreased the levels of benefits at early retirement ages (the malus) and then increased the slope in the benefit profiles (bonus) to provide increased incentives for later retirement. 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. As illustrated, nominal adjustments in later years were lower than inflation so that real benefits declined between 1998 and 2002.

4.3 Data

We use social-security records data from the Austrian Social Security Database, provided by *Synthesis Forschung.* Based on this administrative data, our sample consists of private sector employees in the years 1984 through 2003. The data and sample restrictions are described in more detail in the corresponding appendix below. We construct two key variables to capture incentives from the government-provided pensions. The two variables are social security wealth (SSW) and the accrual (ACC). An individual's social security wealth at a given age is defined as the expected present discounted value of his annual pension benefits if he were to retire at the given age. More precisely, we can write SSW as

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

where $b_i(a)$ denotes individual i's annual benefits when retiring at age a, $\pi_{t|a}$ denotes the probability of survival until age t conditional on having survived until age a and $\beta = .93$ captures the individual's discount factor.[14] In this definition, we also assume that the maximum age that individuals can live to is age 100. Each individual's retirement pension

[14] 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.

is calculated based on the rules of the Austrian pension system and the individual's observed earnings history. While the social security wealth variable reflects the levels of benefits, the second pension variable, the accrual, reflects the slope of the benefits schedule across potential retirement ages. In particular, an individual's accrual at a given age a captures the expected change in his social security wealth $SSW_{i,a}$ net of pension contributions from delaying retirement by one additional year. Thus we define the accrual for individual i at age a as

$$ACC_{i,a} = \frac{E_a(SSW_{i,a+1}) - SSW_{i,a}}{SSW_{i,a}}.$$

In calculating the individual's expectation, we assume 1.75% real wage growth to project earnings one year ahead.

Table 2 presents summary statistics by age for key variables from the data used in the empirical analysis. All euro amounts are in 2003 euros; in January 2003, the euro-U.S. dollar exchange rate was 1 euro to roughly 1.06 dollars. The statistics at each age are based on individuals who are not yet retired (i.e. still in the labor market) at the given age, so selection should be taken into account when interpreting profiles across ages within the table. At age 55, the median earnings are roughly $33,000$ euros and the median annual benefits are roughly $21,000$. Median earnings increase across the ages indicating that higher income earners tend to retire later. Annual earnings are computed based on the calendar year that an individual reaches the specified age, and this accounts for the earnings dips at ages 60 and 65 because individuals at these ages work only part of a calendar year and then retire once they reach either age 60 or 65. Based on the annual benefits, survival probabilities, an average inflation rate of 1.5% and a discount factor of $\beta = .93$, median social security wealth ranges from about $260,000$ euros at age 55 to $315,000$ euros at age 65 reflecting that higher earners who have yet to retire at the later ages have higher social security wealth. The accrual is close to -10% at each age reflecting the loss in social security wealth from lack of actuarial adjustments. Additionally, the accrual becomes slightly more negative after age 60 reflecting that higher income earners give up more of their social security wealth when they delay claiming their pension.

Asset data is also important for the empirical analysis. Because such wealth data is not available in the social security records data, we use asset data from the Survey of Health, Ageing, and Retirement in Europe (SHARE).[15] This SHARE dataset has wealth data for individuals in several European countries. We focus on the data collected for Austria in

[15] Information on the SHARE dataset can be found at http://www.share-project.org/.

2005. In particular, we use data on household gross financial assets for 1,391 Austrians ages 50 through 54 in 2005. We present summary statistics characterizing this distribution of assets (in 2003 euros) in the bottom section of Table 2. The data indicates that households have accumulated financial assets roughly equivalent to one-year's earnings. We discuss how this asset data and the other variables summarized in Table 2 are used in the empirical analysis in more detail below.

5 Empirical Analysis I: Identification & Estimation of Elasticities

Before turning to the estimation of the model developed above, we first focus on the identification and estimation of income and price elasticities in individuals' retirement decisions in this section. Specifically, we identify these elasticities using variation in individuals' retirement benefits from the pension reforms in Austria. We present graphical evidence highlighting this identifying variation. We then estimate the elasticities using this identifying variation in the context of a Cox proportional hazards specification. Importantly, this specification is not derived explicitly from the dynamic model above and thus does not impose any distributional or functional form assumptions used in that model. Instead, the use of the proportional hazards specification is motivated by regression specifications commonly used in the literature (for examples, see Lumsdaine, Stock and Wise (1992), Coile and Gruber (2000a,b), Gruber and Wise (2004)).

5.1 Graphical Evidence

Our identification strategy exploits policy variation based on a series of pension reforms in Austria that independently varied the level and slope of pension benefits across ages. Figure 3 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 3, 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 and 1997 from roughly -.096 to -.087, 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 0.05 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 0.10 to 0.03. These graphical results imply very large elasticities. While it is possible that individuals at age 55 are particularly responsive to financial incentives since age 55 is the first possible age for retirement, it is also possible that there are other confounding changes that make a causal interpretation of the implied elasticities tenuous.

The key to the 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.

While Figure 3 focuses on changes at age 55 to avoid complications from survival bias (recall that age 55 is the first age for retirement), we next turn to illustrating the identifying variation from the pension reforms across all ages and years. To do this, we first regress a retirement indicator, the log of social security wealth and the log of the accrual each on earnings history polynomials and age, year, industry, region, blue and white collar, and change-in-eligibility dummies.[16] We then obtain the residual for these three variables and create year-age cell means for each variable. We then plot these cell means in Figure 4. By controlling flexibly for year, age and income groups, the remaining variation in the residuals comes at the level of year-age-income group interactions. This is the level of variation from the pension reforms which differentially impact different income groups at different ages in different years.

Consistent with Figure 3, the plots in Figure 4 indicate relatively smaller wealth effects and larger price effects. Specifically, the scatter plots show a steeper slope with the accrual residuals than with the social security wealth residuals (-4.784 versus 0.717). Furthermore, in a bivariate regression using the cell means, the estimated coefficients (and standard errors clustered at the year level) on the social security wealth and accrual residuals are respectively 0.849 (0.255) and -4.939 (0.799) respectively. These estimates indicate a wealth -to-price elasticity ratio of roughly 0.17. We now turn to estimating these elasticities more directly in the context of a Cox proportional hazards specification.

5.2 Proportional Hazards Specification & Elasticity Results

To determine the income and price elasticities of retirement benefits on retirement age, we estimate the following Cox proportional hazards model on men between the ages of 55 and 65 between 1984 and 2003,

$$R_i(a) = \bar{R}(a) \exp\{\beta_{SSW} \ln(SSW_{i,a}) + \beta_{ACC} \ln(ACC_{i,a}) + \delta X_{i,a}\}.$$

[16] The regressions are of the form

$$Y_{ia} = \delta X_{ia} + \varepsilon_{ia}$$

where the subscripts refer to individual i at age a. The change-in-eligibility dummies capture changes in the mechanical rules governing retirement that are independent from changes in financial incentives. As documented in Table 1, these changes are (1) the introduction of a disability pension at age 57 between 1993 and 2000, the increase in the retirement age from age 60 to 61.5 between 2000 and 2002 and (3) the increased restrictions for claiming disability after 2000.

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. As defined in Section 3, $SSW_{i,a}$ is the expected present value of the individual's retirement pension if he were to retire at age a, and $ACC_{i,a}$ is the individual's expected pension accrual (i.e., the change in $SSW_{i,a}$ from delaying retirement by an additional year). 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 quartic polynomials in calendar year, log annual earnings and log total earnings from the prior 10 years to control for individuals' earnings histories. The full controls include the base controls as well as dummies for education, industry and region, and quartic polynomials in log annual earnings from each of the prior 10 years. We also include a quartic tenure polynomial to control for potential heterogeneity in preferences for work that may be correlated with higher levels of job tenure.

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 primarily estimated 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 are interested in, whereas the alternative models present coefficients that cannot be easily converted into elasticities. In particular, the coefficients β_{SSW} and β_{ACC} relate to the income and price effects that identify the coefficient of relative risk aversion as discussed in Section 2.2. β_{SSW} captures the elasticity of retirement with respect to pension wealth, and β_{ACC} captures the elasticity of retirement with respect to the one-year accrual.

We exploit exogenous variation in retirement benefits created by the five pension reforms in Austria between 1984 and 2003 to identify a causal relationship between retirement benefits and retirement decisions. Specifically, we employ control function methods to use only the variation in pension benefits created by the reforms to identify β_{SSW} and β_{ACC} (Heckman and Robb 1985). Without the exogenous variation from the reforms the identification of causal effects is threatened by unobserved heterogeneity in preferences for work. Intuitively, individuals with greater willingness to work may have higher earnings and hence higher pension benefits, thereby creating a correlation between benefits and retirement decisions. We include polynomials in individuals' earnings histories to control for systematic 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 entirely to the pension reforms, is used to identify the pension wealth and accrual elasticities. In addition, we are able to separately identify both the income and price effects because we observe multiple pension reforms that create independent variation in the level and slope of benefits across retirement ages.

The results from the Cox proportional hazards model are presented in Table 3, Panel (A). The first two columns present estimates of the coefficients on log Social Security Wealth (SSW) and the log accrual rate (ACC) estimated on the entire sample with the base and full controls, respectively. The base results indicate that a 1% increase in pension wealth increases the hazard by 0.44% while a 1% increase in the accrual measure decreases the hazard by roughly 2.9%. After including the full control set, the pension wealth estimate decreases slightly to 0.40% while the estimate for the accrual increases in magnitude to -3.38%. Consistent with the graphical evidence presented above, we estimate much higher price effects than wealth effects, on the order of 6-7 times higher.

Recall the hazard rates into retirement were characterized by spikes at ages 60 and 65. In the next two columns, we estimate the model on the sample of individuals 60 and 65 only in order to examine the importance of the proportionality assumption (i.e., that covariate effects are proportionate across ages). Note that the effect of pension wealth is estimated to be slightly smaller at these ages and the effect of the accrual slightly larger, however these differences are not statistically different from the estimates on all ages. Finally, the fifth and sixth columns present estimates of the model allowing for time-varying covariate effects. Specifically, we allow the effects to vary linearly with age. To obtain the estimated effect of a covariate at a given age, multiply the coefficient by age minus 54. For example, the estimated effect of $\ln(SSW)$ at age 60 is $0.1095 * (60 - 54) = 0.657$. The corresponding estimate of the accrual effect is -2.762. Note that these estimates are similar across all specifications, and in particular the ratio of the wealth to price effect is small. As a result, we will consider the coefficients from the base model estimated on all ages (column 1) to be our baseline estimates.

Before imposing other distributional and functional form assumptions involved in structural estimation, it is useful to put these elasticity estimates from Table 3, Panel (A) in the context of an economic model by examining what the estimates imply about the coefficient of relative risk aversion. Based on the derivation presented in section 2.2, the structural parameter γ can be expressed as the product of three components: (1) the ratio of the income to price elasticities; (2) the wage-to-asset ratio; and (3) the inverse of the consumption difference between work and retirement. Table 3, Panel (B) presents the implied coefficient of relative risk aversion for the range of elasticity ratios reported in Table 3, Panel (A), assuming a wage-to-asset ratio of 1.12 and consumption difference of 0.395. We set the wage-to-asset ratio to 28000/25000, where 28000 is roughly the median wage in our sample and 25000 is the median level of assets in the SHARE data. To benchmark the consumption difference, we use median annual benefits plus 5% of savings while retired (14000+0.05*25000) and median annual benefits net of 10% savings while working (.9*28000). These figures come from the summary statistics shown in Table 2. From Table 3, Panel (B), we see that under the baseline estimates of the income and price elasticities, the implied coefficient of relative risk aversion is relatively small at 0.43. Under a plausible range of elasticity ratios, we see that the formula implies an upper bound of 1 for γ. Thus, when taking these elasticities into account, structural estimates of γ should reflect a relatively low degree of risk aversion.

6 Empirical Analysis II: Structural Estimation

In this section, we focus on the estimation of the model developed in section 2.1. We first discuss the estimation strategy which makes use of the estimated elasticities form the previous section. We then discuss the estimation results.

6.1 Estimation Strategy

Following French (2005), we fix a set of parameters governing the data generating process of the exogenous state variables (χ), and estimate a set of parameters θ conditional on these values. In particular, in the baseline specification we fix the the life span $T = 100$ years, the real wage growth rate $g = 1.75\%$ (estimated from the social security records data on individuals ages 50-54), and the interest rate $r = 7.5\%$ (based on nominal interest rates net of inflation in Austria during the sample period). We obtain mortality probabilities $\pi_{a|a-1}$ from life tables for Austria, and we estimate job separation probabilities π_{sep} directly from the social security record data at ages 50 through 54. Finally, we also fix the discount factor $\beta = \frac{1}{1+r} = 0.93$ since it is difficult to distinguish empirically β from declining work disutility across ages. Thus, $\chi = (T, \bar{s}, g, r, \pi_{a|a-1}, \pi_{sep}, \beta)$. Since we do not observe data on assets, we approximate the initial distribution of assets at age 54 using the Austrian Survey of Health, Ageing and Retirement in Europe (SHARE) data; specifically, we sample initial assets for each individual with replacement from the empirical distribution of assets in the SHARE data.

We parameterize the model presented in Section 2.1 as follows. We assume constant relative risk aversion (CRRA) utility over consumption:

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma},\ \gamma > 0.$$

We assume initial work disutility is drawn from an exponential distribution with mean $\tilde{\eta} = x\eta$, where $\eta > 0$ and $x = u(\bar{c}) - u(\bar{r}\bar{c})$ is a scaling factor for the disutility of work based on income differences between work and retirement (we use $\bar{c} = 30000$ and $\bar{r} = .55$ based on mean wage income and the replacement rate). Work disutility increases linearly with age, with slope $\alpha\tilde{\eta}$ (i.e., $\upsilon_a = \alpha\tilde{\eta}(a-54) + \upsilon_{54}$). Thus, the parameters we are interested in estimating are $\theta = (\gamma, \eta, \alpha, \kappa)$, where κ equals the monetary cost of claiming a disability pension.

For the estimation, we assume that individuals make decisions with complete knowledge of how pension benefits are calculated in a given calendar year. We assume that their projections of future benefits are based on that year's legislation only. Further, we assume that the pension reforms were unanticipated, and that individuals immediately update their calculations based on the new rules. We assume that individuals expect their future earnings to grow at a constant rate per year. In regard to job separations, we assume that the probability of job separation varies only by years of tenure. We assume that separation shocks do not affect wages, so that conditional on separations, wages are still expected to grow at the same constant rate (this is supported by evidence that collective bargaining agreements tend to set wages based on labor market experience rather than tenure). This simplifies the computation of projected pension benefits since we can project pension benefits for individuals at each age based on a single expected earnings path rather than based on multiple paths from different potential histories of job separations.

We estimate the model using a method of simulated moments (MSM) estimation strategy. This estimation strategy matches 13 empirical moments to the corresponding moments in a simulated sample. The empirical moments are the retirement hazard rates at ages 55 through 65 and the income and price elasticities β_{SSW} and β_{ACC} from the proportional hazards specification. The use of retirement hazard rates across ages as moments in the estimation is common in the literature. The use of the additional moments, the income and price elasticities β_{SSW} and β_{ACC}, is motivated by the theoretical relationship implied by the model, namely the relationship between the coefficient of relative risk aversion, a structural parameter, and the labor supply elasticities.

Using m to denote the vector of empirical moments and $\hat{m}(\theta, X)$ to denote the corresponding simulated moments based on parameters θ and data X, the MSM estimator is formally defined by

$$\widehat{\theta}_{MSM} = \arg\min_{\theta} \left[m - \hat{m}(\theta, X)\right]' W \left[m - \hat{m}(\theta, X)\right],$$

where W is the weighting matrix. The hazard rate moments are weighted by the observed survival function at each corresponding age and the proportional hazards coefficients are weighted by the inverse of the corresponding standard errors (see Table 3).

Because we do not observe consumption savings data for individuals in our sample, we are not able to empirically identify the optimal policy function for consumption while working, c_a^W. Thus, we estimate the model using an approximation for this optimal policy function. In particular, we approximate c_a^W using optimal consumption given a fixed, future retirement age that is allowed to vary at each age. In the baseline specification of the model, this fixed retirement age is assumed to be the next age; intuitively, at each age a in which he faces a retirement decision, an individual computes his optimal consumption if he continues working based on the expectation that he will retire at age $a + 1$. We explore alternative approximations for this optimal policy functions and demonstrate that the results are robust to considering these alternatives. The next section discusses these results.

6.2 Structural Estimation Results

Table 4, Panel (A) displays the structural parameter estimates. We focus first on the estimates of the baseline specification of the model presented in Column 1. The estimated coefficient of relative risk aversion in the baseline model is 0.709. This relatively low estimate for the degree of relative risk aversion is driven by the income and price elasticities that are included as moments in the estimation. Specifically, the estimation adjusts γ so as to fit the income and price elasticities since a lower γ implies a relatively smaller income elasticity and a larger price elasticity. Panel (B) of Table 4 displays that estimated income and price elasticities based on the simulated data using the baseline parameter estimates. The simulated income and price elasticities are respectively $\beta_{SSW} = 0.47$ and $\beta_{ACC} = -3.11$. These elasticities predicted by the model are similar to the corresponding estimates using the real data, $\beta_{SSW} = 0.43$ and $\beta_{ACC} = -2.90$ (see Table 3).

To highlight the relationship between the risk aversion parameter and the labor supply elasticities, Table 5 presents the simulated elasticities when varying γ around its baseline estimate while holding the remaining parameters constant at their respective baseline estimates. The results in this table demonstrate that a lower γ implies a lower ratio of income and price elasticities as discussed in Section 2. However, varying γ also demonstrates that, while it may improve the model's prediction regarding the elasticity ratio ($\frac{\beta_{SSW}}{-\beta_{ACC}}$), a lower γ gives the estimation a poorer fit in terms of the magnitudes of the elasticities. Specifically, a lower γ implies generally more responsiveness to financial incentives (a higher overall elasticity $\beta_{SSW} - \beta_{ACC}$) as the magnitudes of both β_{SSW} and β_{ACC} increase.

Another parameter of particular interest is the claiming cost κ. In the baseline specification, this cost is estimated to by roughly 100,000 euros. To put this cost in perspective, the social security wealth (expected presented discounted value) from an annual pension benefit of about 8,300 euros would be roughly 100,000 euros (using the discount rate and survival probabilities used in the model). In terms of benefits then, the claiming cost can be interpreted to be about 40% ($\frac{8300}{21000}$) of one's annual pension benefit. This claiming cost allows the model to fit the spike in the retirement hazard rates at age 60. In particular, given the low estimate of γ, the model predicts that the price effect from a 40% increase in benefits plays a significant role in accounting for the spike in the hazard rates at age 60.

Columns 2 and 3 of Table 4 explore the robustness of the results to alternative approximations for the optimal policy function for consumption while working. The specification in column 2 computes c_a^W using optimal consumption given an expected retirement age that is determined at each age as follows. Prior to age 60, individuals expect to retire at age 60; at age 60 through 64, individuals expect to retire at age 65; and at age 65 individuals expect to retire at age 66. The focus on expected retirement at ages 60 and 65 is motivated by the spikes in the hazard rates at these ages. The results in column 2 of Table 4 indicate that in this specification, γ continues to be estimated to be relatively low at 0.707.

The specification in column 3 assumes that c_a^W is determined according to a fixed savings rule so that individuals save 10% of their wage income while working and consume the remainder. This specification is motivated by earlier studies, such as Stock and Wise

(1990) and Rust and Phelan (1997), that have taken consumption to equal current income and have therefore assumed a fixed savings rate (0% instead of 10%) as well. In this case, c_a^W does not vary with γ and the differences in c_a^W between higher wage and lower wage individuals is greater than in the previous specifications since c_a^W increases linearly with wages. While γ is still estimated to be relatively low, the predictions in terms of both the magnitudes of the elasticities and the ratio of elasticities are poorer than in the baseline specification. The model's relative inability to fit these moments is likely driven by the rigidity of the fixed savings rule. This consumption rule implies higher wage individuals are more likely to continue working than in the other specifications in which consumption did not increase linearly with wages. Thus these higher wage individuals drive the estimation of the elasticities, especially at higher ages. Allowing consumption to vary with γ therefore seems important for fitting the labor supply elasticities.

The last column of Table 4 presents the results from estimating the baseline specification of the model with a lower discount factor ($\beta = 0.88$ instead of 0.93). While the other parameters adjust more noticeably to the lower discount factor, γ continues to be estimated to be relatively low at 0.677. Thus, changes in the discount factor do not seem to affect the overall point; a low γ is necessary to fit the estimate income and price elasticities.

Figure 5 presents actual versus predicted retirement hazard rates by age for each specification. In examining the model's ability to fit the retirement hazard rates, there are a couple of features that standout. First, the model overpredicts retirement at the earliest age, age 55. Intuitively, the model predicts that any individuals that draw a high disutility of work should retire immediately since they will always have high disutility of work at any age. Second, the model underpredicts retirement just before age 60. This is because the model predicts that forward looking agents should simply delay their retirement so as to avoid paying the relatively large fixed cost of claiming a disability pension. Thus, with a discount factor relatively close to 1, it is difficult for the model to fit the hazard rates just before age 60. Lastly, the model underpredicts retirement after age 60. This is because the model predicts that only higher wage earners will continue working beyond age 60. Since these earners have relatively high returns to continue working, they have little incentive to retire at any age. Additionally, the estimation puts less weight on fitting these moments at higher ages since the majority of individuals retire at age 60 and earlier.

The structural results generally emphasize that the model requires a low degree of relative risk aversion to fit the elasticities from the proportional hazard specification. While some previous studies have found or used higher values for γ (for examples, see Hubbard

Skinner and Zeldes (1995), French (2005), Blau (2008), van der Klaauw and Wolpin (2008)), the estimates in Table 4 are entirely consistent with estimates of γ implied by previously estimated income and price elasticities (see Chetty (2006)).

7 Policy Simulations

The estimated structural model permits the examination of a variety of policies intended to facilitate retirement at later ages. In particular, we use the estimated baseline model to simulate four hypothetical pension reforms that mimic features of commonly discussed and previously implemented pension reforms in several countries. The changes for each simulated reform are as follows. The first reform is a 20% reduction in pension benefits at all ages (i.e. $y_a^R \rightarrow (1-.20)y_a^R$). The second reform is a change in benefits by 3% per year from age 65 (i.e. $y_a^R \rightarrow [1-.03(65-a)]\, y_a^R$) ; thus, benefits at all ages prior to 65 are reduced while benefits at age 66 are increased. The third reform is a one year increase in the early retirement age (the age at which individuals can first claim a standard old-age pension). To implement this reform, we increase a_{old} from 60 to 61 so that individuals face the cost of claiming, κ, at age 60 but not at 61 and beyond. The fourth reform is based on having graduated or phased-in eligibility. Rather than having one age like age 60 at which individuals go from qualifying from a low percentage of their benefits to 100% of their pensions, we consider a setting in which eligibility is phased in so that individuals gradually become eligible for larger fractions of their full benefits. To simulate this reform, we simulate the model assuming that individuals must pay the claiming cost according to the schedule presented in Table 6.

To examine the labor supply consequences of these reforms, Figure 6 presents the simulated hazard rates under each reform. Focusing on the effects of the first two pension reforms indicates that the reductions in benefits reduce retirement at the earlier ages and increase the hazard rates at age 60 and beyond. Compared to the first reform, the second reform leads to a greater reduction in the hazard rates at early ages since there is a greater reduction in benefits at the early ages and a smaller reduction at age 60 and beyond. Intuitively, lower benefits lead to later retirement, but once individuals qualify for their old-age pensions at age 60, most individuals retire since their effective wage rate for continuing work falls (after age 60, individuals are passing up their benefits without any claiming cost). These labor supply responses highlight the relatively smaller income effects and larger price effects since reductions in income generate relatively small changes in the hazard rates while the change in the effective wage at age 60 still produces a dramatic spike in the hazard rate at age 60. Additionally, the outcomes from the second reform indicate that reforms which make benefits actuarially fair in Austria will not lead to significant increases in retirement at older ages.

As illustrated in Figure 6, the labor supply responses to the third and fourth pension reforms highlight both individuals' responsiveness to changes in effective wage rates and the importance of understanding multiple pathways into retirement. In regard to the third reform, the increase in the early retirement age (first eligibility age for old-age pensions) leads to a shift in the spike in the hazard rates from age 60 to age 61. By extending the claiming cost to age 60, the effective wage for working at age 60 rises, but it still falls discontinuously at age 61. There is also an increase in early retirement prior to age 60. Once individuals anticipate that they will have to continue working for a longer period to qualify for their old-age pensions, the model predicts that they would rather pay the fixed cost and retire early through disability. Thus, individuals are willing to substitute between multiple path ways to retirement. Similar to the third reform, the phased-in eligibility of the fourth reform creates increases in effective wages for continuing to work beyond age 60. Because this reform eliminates any discontinuous changes in the claiming cost prior to age 65, there are no longer any spikes in the hazard rates prior to age 65 as illustrated in Figure 6. At age 65, however, the effective wage rises sharply since the cost of claiming goes to 0. Intuitively, individuals anticipate this decrease in effective wages and continue working through age 64 and then retire at age 65.

Next, we examine the welfare consequences of each of the simulated reforms. To measure welfare, we use expected utility at the initial age, age 55, so that all individuals are considered at the same age. While Table 7 presents a numerical assessment of the changes in welfare under each of the reforms relative to the baseline, we focus first on Figure 7 which presents the distributions of expected utilities at age 55 for the baseline setting and for each pension reform. Relative to the baseline, the uniform decrease in benefits from the first reform leads to a uniform decrease expected utility as the entire distribution is shifted to the left. The numerical results in Table 7 demonstrate that the 20% reduction in benefits does not lead to a dramatic reduction in welfare. This is because a lower γ implies that the marginal utility of consumption does not rise sharply with a decrease in consumption.

The second pension reform also leads to a reduction in welfare, but the numerical and graphical results both show that variance of the distribution increases relative to the baseline. The intuition for the increase in the variance is as follows. The model predicts that lower wage earners are more likely to retire at early ages than higher wage earners since both groups are equally likely to draw high work disutilities. Because the reductions in benefits are largest at the earlier ages under the second pension reform, low wage earners experience the largest reductions in benefits while higher wage earners can continue working to older ages and thus experience lower reductions in their benefits. Thus, relative to the baseline distribution of expected utilities, there is a larger decrease in expected utilities for the left side of the distribution than the right.

The welfare results for the third pension reform indicate that, while this reform lead to noticeable labor supply changes, the reduction in expected utilities is relatively less dramatic. In particular, the numerical results in Table 7 show that the reduction in welfare from this reform is similar to that from the first two reforms which lead to relatively smaller labor supply changes. Intuitively, expected utility does not change much under this third reform because of two factors. First, individuals can still retire early with no change in benefits if they have a high disutility of work. Second, if individuals continue working an additional year, the loss in benefits is offset by an additional year's wage income. Thus, the only loss in utility comes from the additional work disutility from having to work an additional year. Following this intuition from the third reform, the welfare results for the fourth reform show the largest decreases in expected utilities since several individuals anticipate working multiple years beyond age 60 instead of just one year.

These responses to the simulated pension reforms highlight the general consequences of a low degree of relative risk aversion for policy. First, a lower γ indicates that policies designed to change effective wage rates at older ages will be particularly effective at facilitating retirement at older ages. In contrast, policies that reduce total income through benefit reductions will have relatively small effects on labor supply. Second, a lower γ indicates that reductions in the generosity of benefits do not lead to large reductions in welfare since a low degree of relative risk aversion implies that marginal utility of consumption does not increase sharply with reductions in income. Third, a lower γ implies that the variance of expected utility rises relatively quickly when reforms facilitate retirement at older ages using changes in the across-age schedule of benefits.

8 Conclusion

How do individuals' retirement decisions respond to changes in retirement benefits? What do these responses imply for economic models of retirement and for the consequences of potential social security reforms? 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 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. We are able to put these elasticities in the context of a structural model of retirement decisions and directly use these elasticities in the estimation of the model. Based on a proportional hazards specification, we estimate a relatively smaller income elasticity (0.43) and a larger price elasticity (2.90). These results imply a lower degree of relative risk aversion (0.71) than has been previously estimated or used in some studies. We simulate responses to hypothetical pension reforms to highlight the implications of our estimates for understanding the labor supply and welfare consequences of potential social security reforms.

Our results have important implications for understanding social security reform. The relative importance of price effects indicates that incentive effects from benefit schedules across potential retirement ages greatly impact retirement behavior. Hence, reforms based on changing the across-age slope in the benefit schedule are likely to have larger impacts on retirement behavior compared to reforms using across-the-board changes in benefit levels. Furthermore, the relative importance of the price effects implies that the welfare losses from benefit reductions would not be as severe as some earlier studies have suggested.

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.

Banks, J., R. Blundell, and S. Tanner (1998). Is there a retirement-savings puzzle? *The American Economic Review 88*(4), 769–788.

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

Bernheim, B. D., J. Skinner, and S. Weinberg (2001). What accounts for the variation in retirement wealth among u.s. households? *The American Economic Review 91*(4), 832–857.

Blau, D. (2008). Retirement and consumption in a life cycle model. *Journal of Labor Economics 26*(1), 35–71.

Bloemen, H. and E. Stancanelli (2005). Financial wealth, consumption smoothing and income shocks arising from job loss. *Economica 72*(287), 431–452.

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

Browning, M. and T. Crossley (2001). Unemployment insurance benefit levels and consumption changes. *Journal of Public Economics 80*(1), 1–23.

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. (2008). Moral hazard versus liquidity and optimal unemployment insurance. *Journal of Political Economy 116*(2), 173–234.

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.

French, E. (2005). The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour. *Review of Economic Studies 2*(72), 395–Ű427.

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

Gruber, J. (1997). The consumption smoothing benefits of unemployment insurance. *American Economic Review 87*(1), 192–205.

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

Haider, S. J. and D. S. Loughran (2008). The effect of the social security earnings test on male labor supply: New evidence from survey and administrative data. *J. Human Resources 43*(1), 57–87.

Heckman, J. J. and R. Robb (1985). Alternative methods for evaluating the impact of interventions : An overview. *Journal of Econometrics 30*(1-2), 239–267.

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.

Hubbard, R. G., J. Skinner, and S. Zeldes (1995). Precautionary saving and social insurance. *Journal of Political Economy 103*(2), 360–399.

Hurd, M. D. and S. Rohwedder (2008). The retirement consumption puzzle: Actual spending change in panel data. Working paper 13929, National Bureau of Economic Research.

Hurst, E. (2008). The retirement of a consumption puzzle. Working paper 13789, National Bureau of Economic Research.

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.

Liebman, J. B., E. F. Luttmer, and D. G. Seif (2008, December). Labor supply responses to marginal social security benefits: Evidence from discontinuities. Working paper 14540, National Bureau of Economic Research.

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.

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.

OECD (2007). *Pensions at a Glance: Public Policy Across OECD Countries*. Organisation for Economic Co-Operation and Development.

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.

Song, J. G. and J. Manchester (2007). New evidence on earnings and benefit claims following changes in the retirement earnings test in 2000. *Journal of Public Economics 91*(3-4), 669–700.

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

van der Klaauw, W. and K. I. Wolpin (2008). Social security and the retirement and savings behavior of low income households. *Journal of Econometrics 96*(1-2), 21–42.

A Labor Supply & Risk Aversion Appendix

1. Consumption Differences

The relationship between risk aversion and labor supply elasticities draws on concepts that are familiar in the context of studying insurance. In particular, the consumption difference $\frac{c_a^R - c_a^W}{c_a^W}$ reflects differences in consumption across states of employment at a given age. Even though retirement is an anticipated event in the model, the notion of insurance arises from the consideration of unanticipated changes (shocks) to benefits, wages and savings. Intuitively, the degree to which the individual adjusts his labor supply to move between states and offset the effects of these changes on his consumption will depend on how well insured his initial consumption is across the two states. The discrete nature of the employment states and the fact that wages while working are higher than benefits while retired ($y_a^W > y_a^R$) imply that consumption would always be higher if the individual worked rather than retired ($c_a^W > c_a^R$), though the individual must also take into account the disutility of work. This consumption difference between states does not capture intertemporal consumption differences.[17]

2. Preference Assumptions

Given the parsimonious nature of the theoretical framework above, it is useful to characterize the key assumptions that are necessary for identifying the coefficient of relative risk aversion γ based on the labor force participation outcomes, or more specifically based on the income (wealth) and price elasticities in retirement decisions. Two important assumptions that permit this identification of γ are

$$\begin{aligned} (1) \text{ time separability} &: \quad U(c_a, v_a, c_{a+1}, v_{a+1}, ...) = u(c_a, v_a) + u(c_{a+1}, v_{a+1}) + ... \\ (2) \text{ additive separability} &: \quad u(c_a, v_a) = u(c_a) - v_a \quad \forall a. \end{aligned}$$

The time separability assumption allows the risk aversion parameter to be written in terms of labor supply elasticities as well as an additional term arising due to the complementarity between between c and v (i.e. a term based on the derivative of the marginal utility of consumption with respect to work disutility, $u_{c,v}$). Given this time separability assumption, γ is identified based on these terms once this degree of complementarity $u_{c,v}$ is fixed because

[17] Several papers have focused on "the consumption drop at retirement." See Banks, Blundell and Tanner (1998), Bernheim Skinner and Weinberg (2001), Aguiar and Hurst (2005), Hurd and Rohwedder (2008), Blau (2008) and Hurst (2008) for work on this topic. The phrase "consumption drop at retirement" refers to intertemporal consumption changes, i.e., changes in consumption before and after retirement.

any monotonic transformation of the within-period utility function $u(c, v)$ would change the risk aversion parameter but would also change the degree of complementarity. Next, the additive separability assumption imposes zero complementarity ($u_{c,v} = 0$) between consumption utility and work disutility. This degree of complementarity is based on existing empirical evidence that demonstrates that consumption does not decline with exogenous changes to labor supply due to job loss or other shocks.[18]

Another important assumption regarding preferences relates to state-dependence. To estimate γ based on movements between the states of employment, we must assume that γ is state-independent. In particular, the within-period utility function is assumed to be common across the states of work and retirement (i.e., we assume utility $u(c)$ for work and retirement rather than allowing for $u^W(c)$ while working and $u^R(c)$ while retired). Intuitively, if preferences are state-dependent, the movements between states involve both changes in preferences and changes in the marginal utility of consumption, and the two cannot be disentangled using only participation outcomes.

3. Heterogeneity

Under the preference assumptions discussed above, the model could allow for additional features relating to individual-level heterogeneity in preference parameters or consumption rules. First, in regard to individual-level heterogeneity, consider an example with individual-specific discount factors (β_i for the ith individual). In this case, the ratio of wealth and price elasticities would vary by individual, but each individual's ratio could be used to identify the common risk aversion parameter γ. More generally, one could allow the risk aversion parameter to vary by individual as well. With such heterogeneity, as long as the differences are state-independent, a similar relationship between the risk aversion parameter and the labor supply elasticities holds for each individual. Aggregating across individuals then, one could relate the mean γ to the mean elasticity ratio.

[18] Identification of the degree of complementarity between consumption utility and work disutility is based on examining consumption changes in response to exogenous shocks to labor supply. Intuitively, if individuals do not adjust their consumption in response to exogenous changes in labor supply, then this degree of complementarity must be low. See Gruber (1997), Browning and Crossley (2001) and Bloemen and Stancanelli (2005) for evidence on consumption changes around job loss and Chetty (2006) for a discussion of this evidence on the degree of complementarity between consumption and labor supply. Additionally, see Hurd and Rohwedder (2008) and Hurst (2008) for evidence that anticipated changes in labor supply at retirement do not lead to significant consumption changes.

B Data & Sample Restrictions Appendix

The administrative data from the Austrian Social Security Database, provided by *Synthesis Forschung*, covers nearly all individuals employed in Austria between the years 1972 and 2003, with the exceptions relating to tenured public sector employees and self-employed individuals.[19] 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.

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. Importantly, the social security record data contains all information used in the computation of retirement benefits except insurance years which we are able to impute using the labor market histories.[20]

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 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 tenure.

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

[20]Insurance months are determined using the following imputation for insurance years. Specifically, insurance years are imputed as $InsYrs = Age - Edu - 6 - (\textit{time observed not working})$ where Edu is years of schooling. We observe education for the sample of individuals who experience unemployment and claim benefits during the length of the data. Using this data, we regress education on earnings and quartic polynomials in calendar year and age. We then obtain imputed education using the fitted values from this regression. Using the labor market histories observed in the data, we compute *time observed not working*. Assuming that education begins at age 6, we combine the predicted education with this information from the observed labor market histories and round up to the nearest year to compute insurance years (years of experience). Insurance months are then given by $InsMths = 12 * InsYrs$.

Our main sample consists of men ages 55 through 65 who are first observed at age 55 in the years 1984 through 2003. Our sample restrictions and the reasons for these restrictions are as follows. We start by focusing on men aged 55 or higher in 2003 (birth cohorts 1948 and earlier). 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 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 individuals who claim non-disability or non-old-age pensions at the time of retirement since these claims may not correspond to retirement decisions.[21] We exclude men claiming disability pensions before age 55 on the basis that these individuals are likely to be permanently disabled. We also exclude individuals who retire after age 65 since focusing on ages 55 through 65 simplifies recursive computation of the value functions and most retirees (roughly 99%) retire by 65. Next, we exclude remaining individuals with insufficient earnings histories to compute pension benefits and individuals with outlying observations and missing data. This leaves us with a final sample of 252,907 individuals and 178,997 claimants. Not all individuals are observed to be claimants since some individuals (those at later calendar years) are only observed at younger ages. The sample restrictions are summarized in Table A1.

[21] 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.

Figure 1: Hazard Rates & Survival Function

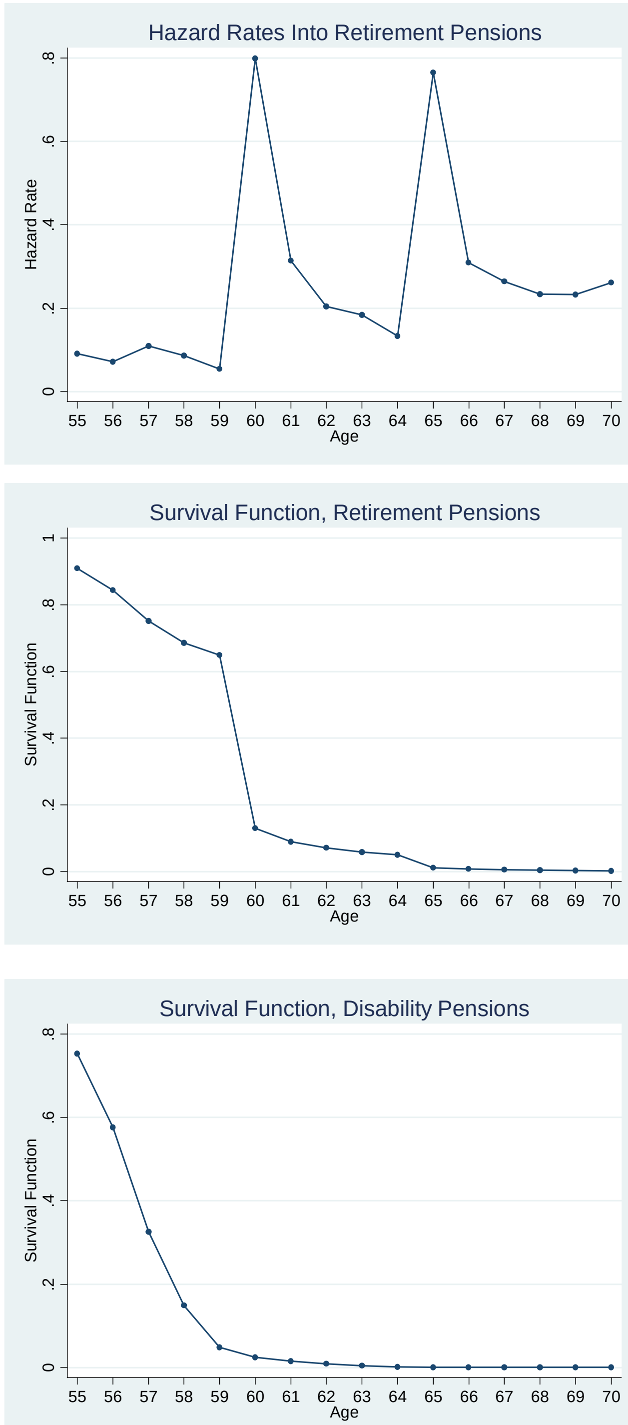

Notes: These figures are based on the sample of all individuals claiming retirement pensions after 1984 (394934 individuals and 275,379 claimants).

Figure 2: Changes in Benefits from Pension Reforms

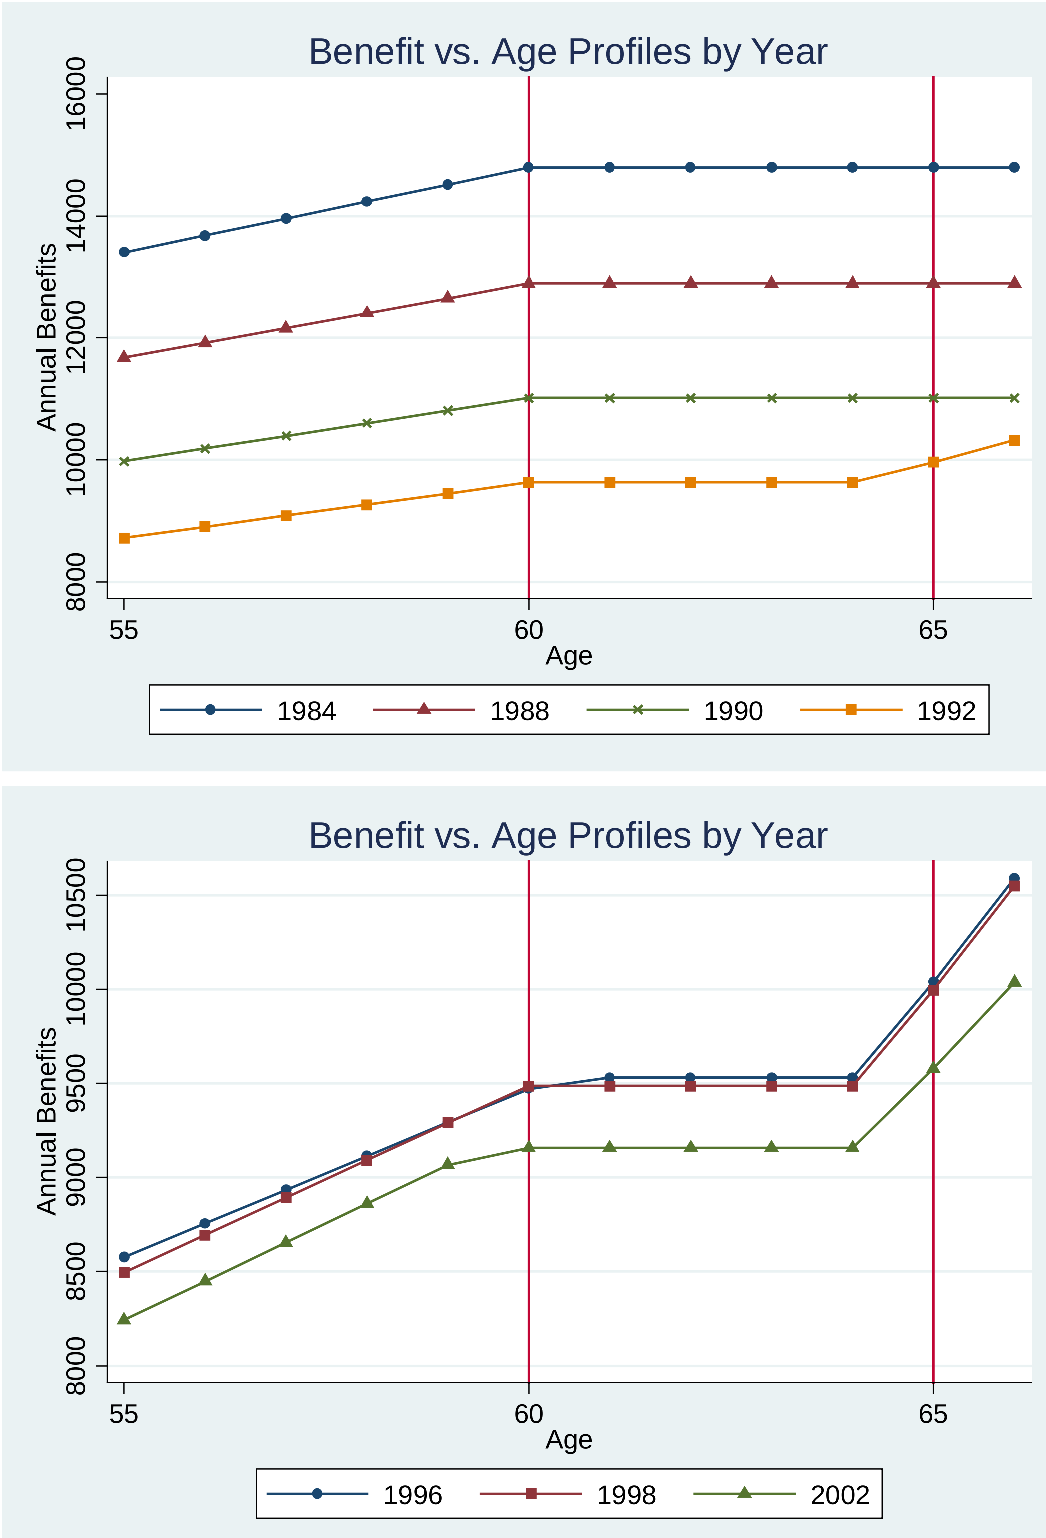

Notes: Benefits are computed under the following assumptions: full insurance years at each age, fixed birth cohort across retirement ages, earnings history with positive earnings in last 10 years, nominal earnings in each year equal to 20000 euros in 2003. All nominal benefits in each calendar year are adjusted to 2003 euros. Please see the text for more details.

Figure 3: Graphical Analysis of Wealth & Price Effects, Age 55

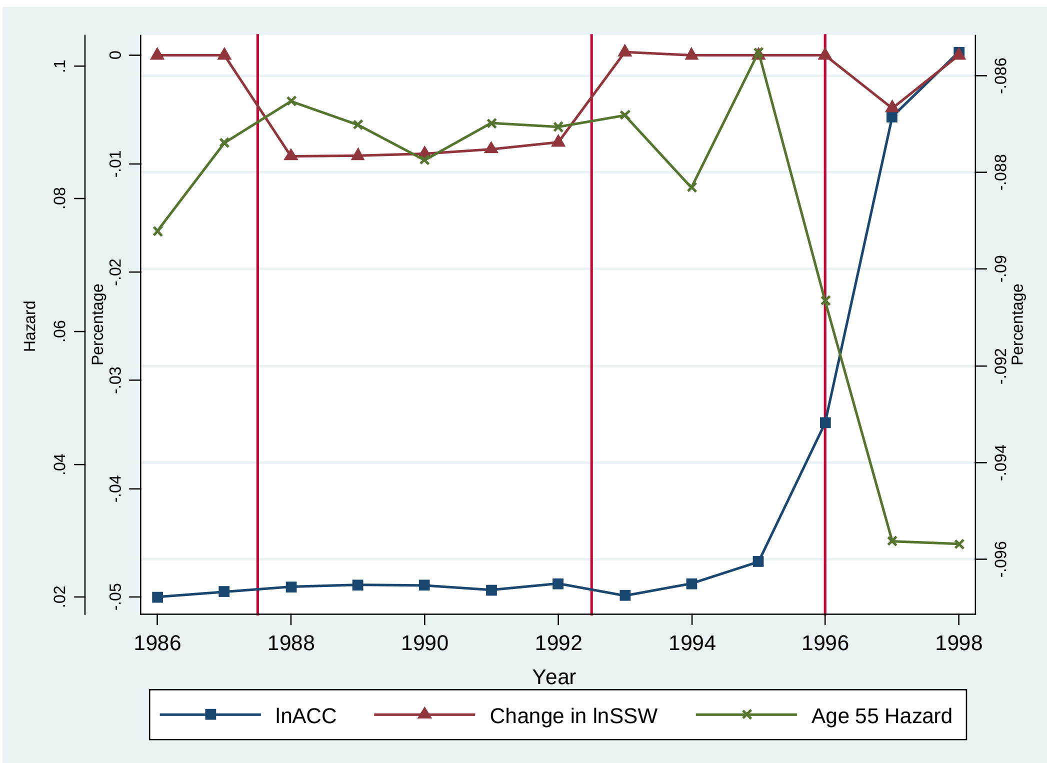

Notes: This figure is based on data at age 55 only. Please see Table 2 for the sample restrictions. The change in log social security wealth is computed relative to the previous year's legislation. Please see the text for more details.

Figure 4: Graphical Analysis of Residuals, Year-Age Cell Means

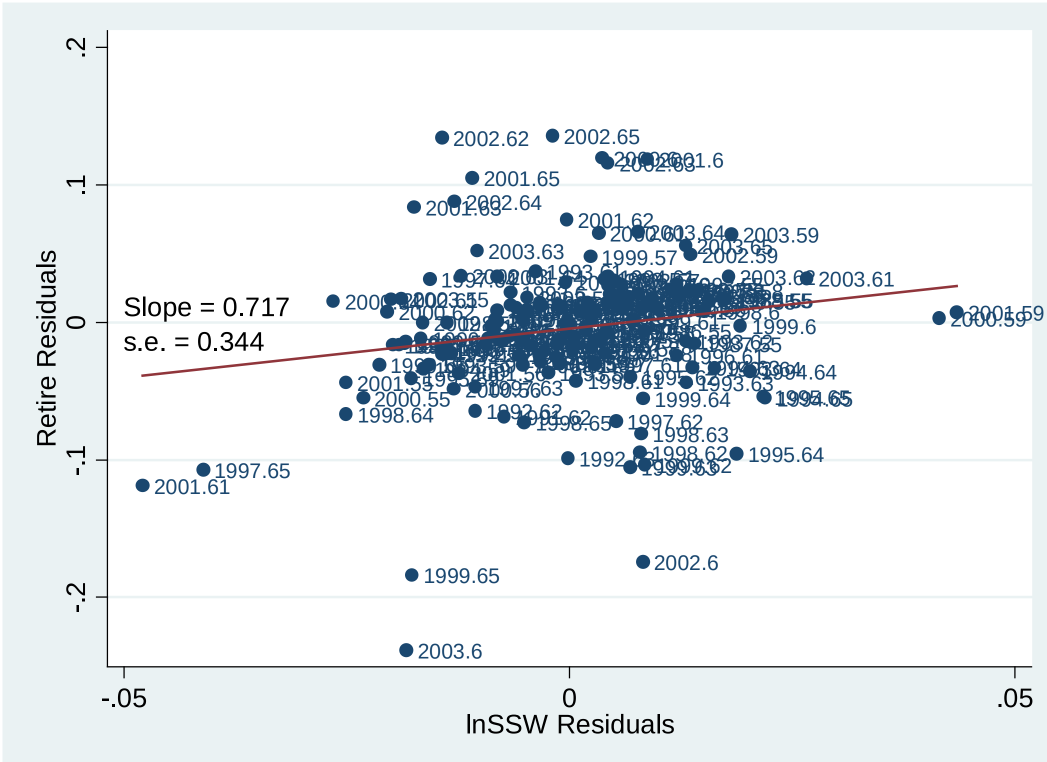

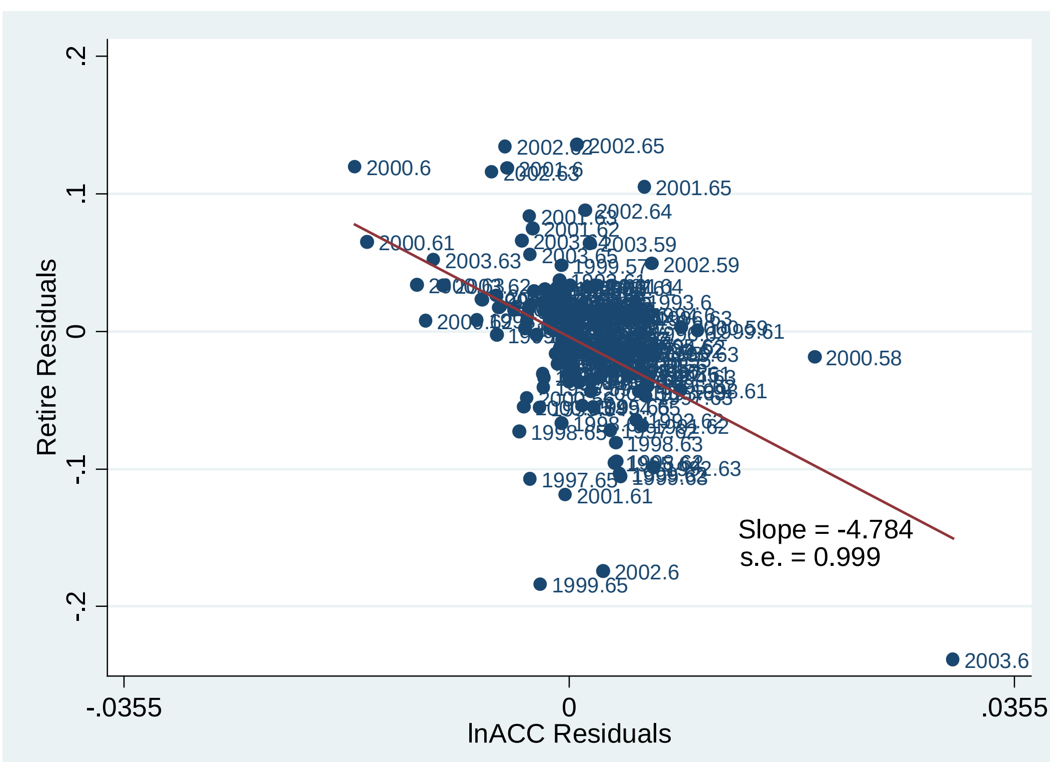

Notes: Standard errors for the slope coefficients are clustered at the year level. Points are labeled by year.age.

Figure 5: Structural Estimation, Predicted vs. Actual Hazard Rates

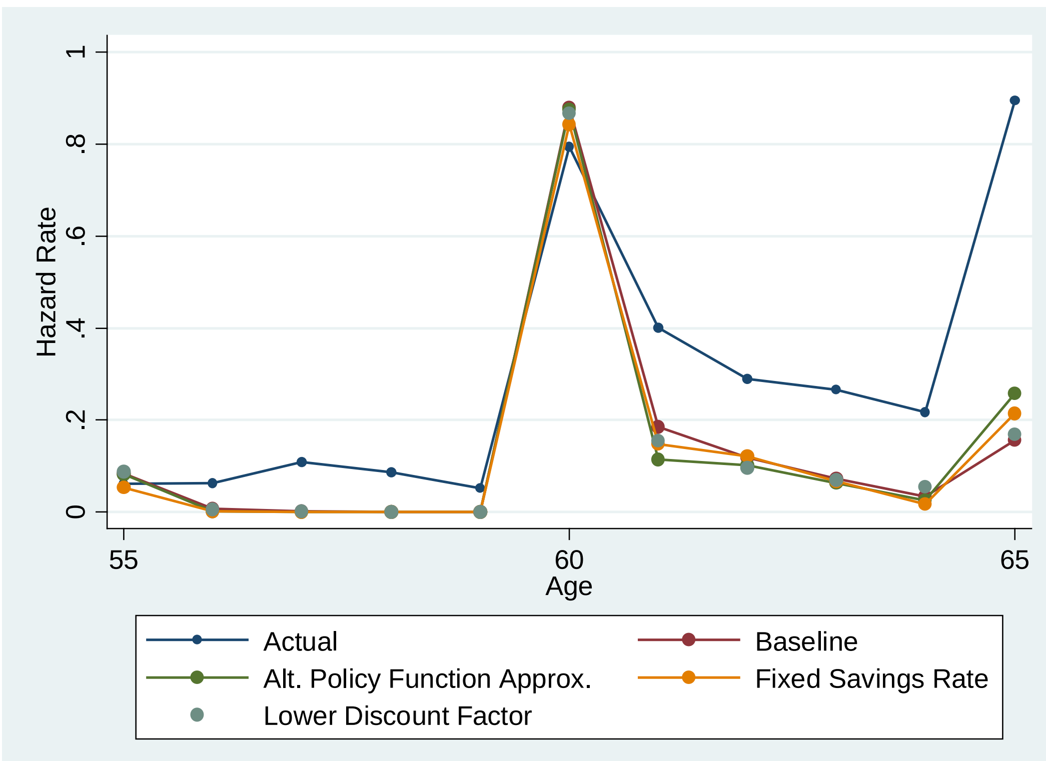

Notes: The baseline model uses a discount factor of β=0.93 and a fixed savings rate of 0.10. The lower discount factor is β=0.88 and the lower savings rate is 0.05. Please see the text for more details.

Figure 6: Labor Supply Responses to Policy Simulations

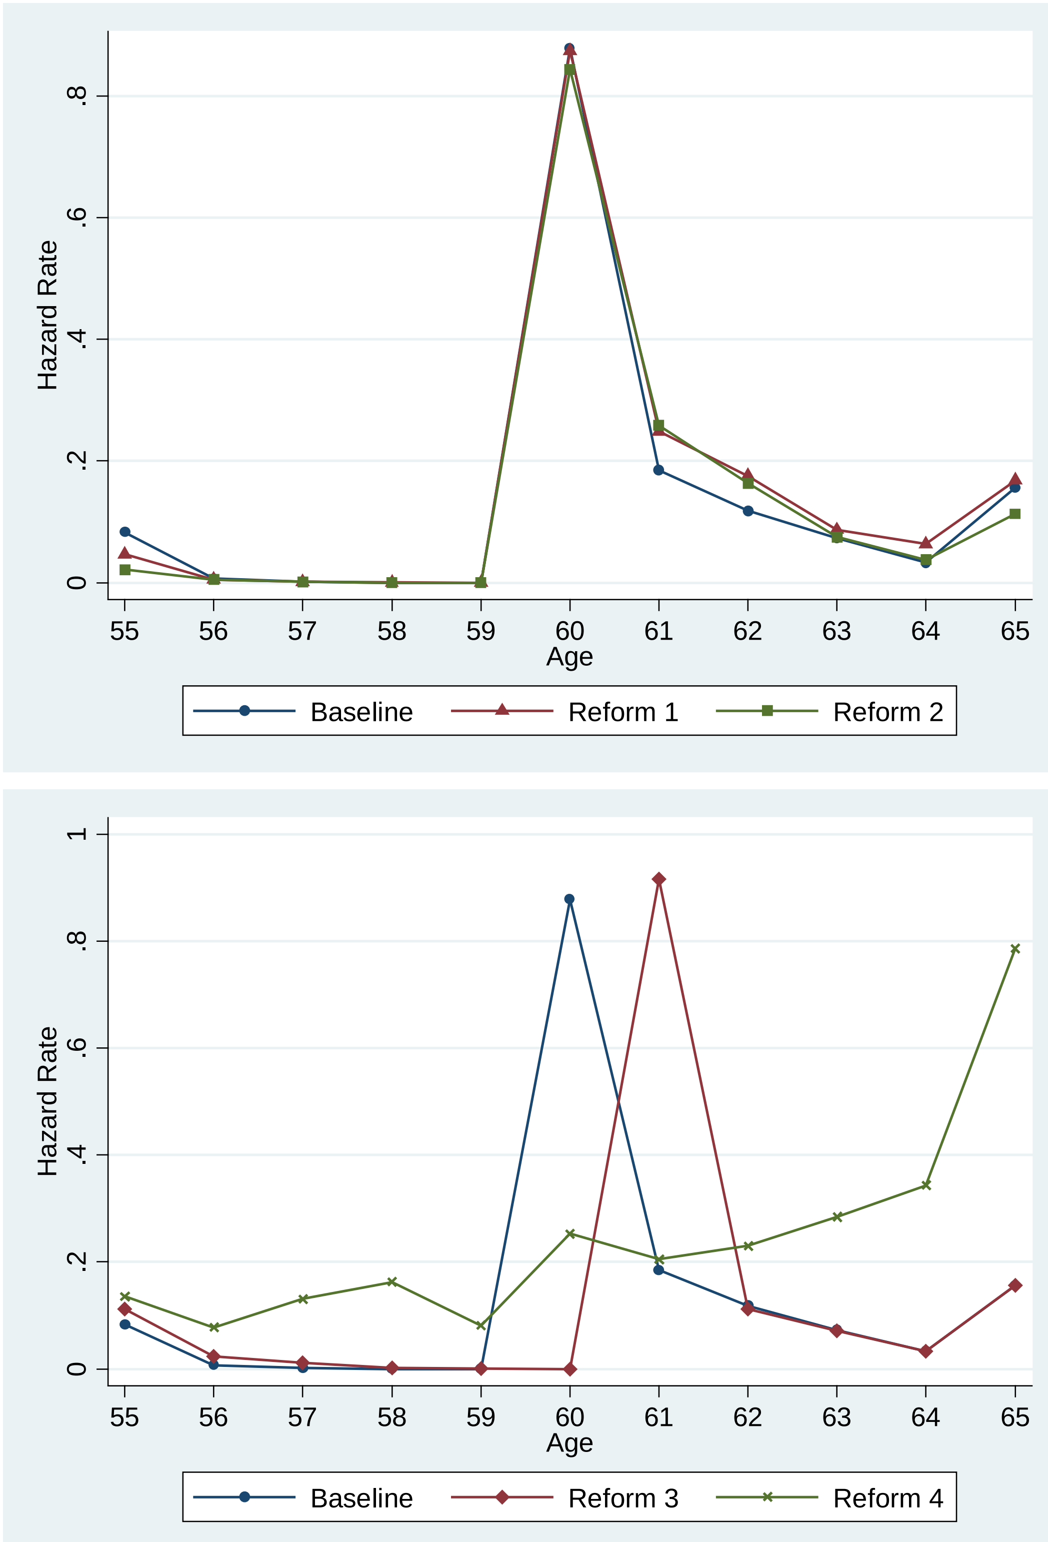

The baseline is computed using the parameter estimates reported in Table 5. Reform 1 decreases benefits across all ages by 20%. Reform 2 changes benefits at a give age based on -3%*(65-age). Reform 3 increases the eligibility age by 1 year. Reform 4 implements phased-in eligibility . Please see the text for more details regarding the reforms.

Figure 7: Welfare Consequences of Policy Simulations

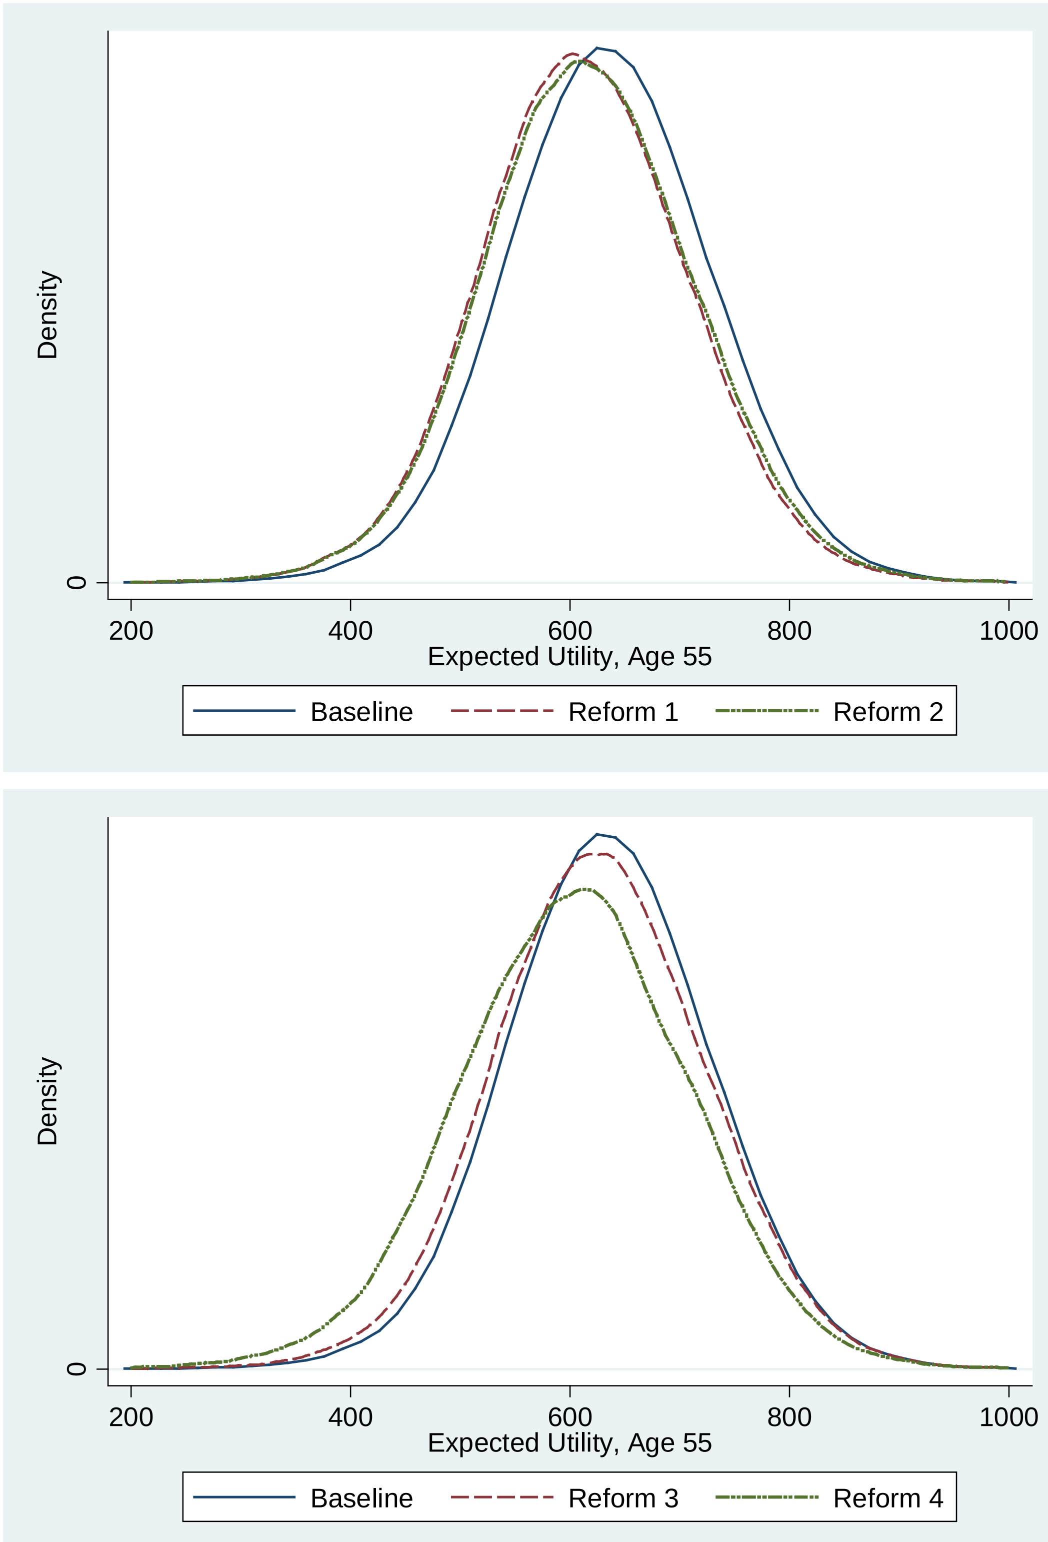

The baseline is computed using the parameter estimates reported in Table 5. Reform 1 decreases benefits across all ages by 20%. Reform 2 changes benefits at a give age based on -3%*(65-age). Reform 3 increases the eligibility age by 1 year. Reform 4 implements phased-in eligibility . Please see the text for more details regarding the reforms.

Table 1
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 2
Summary Statistics by Age

Age		Annual Earnings	Annual Benefits	SSW	ACC
55, N=242,402	mean	39711.99	20959.57	258450.10	-0.089
	median	33127.34	21480.03	264852.50	-0.090
	std. dev.	25620.50	4910.24	62148.53	0.009
56, N=197,959	mean	39822.22	21746.77	264092.70	-0.091
	median	33323.61	22384.38	271606.60	-0.091
	std. dev.	27461.05	5076.34	63352.13	0.008
57, N=172,739	mean	39841.66	22597.94	269996.20	-0.094
	median	33402.45	23367.69	279006.30	-0.093
	std. dev.	28903.82	5224.69	64103.07	0.009
58, N=144,321	mean	40324.68	23465.54	275572.70	-0.094
	median	33715.05	24540.40	287978.70	-0.095
	std. dev.	30562.88	5368.77	64846.93	0.011
59, N=123,954	mean	37798.21	24342.59	281022.50	-0.095
	median	31842.09	25679.08	295721.00	-0.097
	std. dev.	33052.84	5554.38	66215.05	0.012
60, N=108183	mean	26147.84	25150.57	284951.30	-0.102
	median	18919.41	26579.45	299823.80	-0.109
	std. dev.	31470.91	5747.87	67563.28	0.020
61, N=16,268	mean	47094.05	26803.42	300095.00	-0.113
	median	40401.54	29790.58	323255.80	-0.118
	std. dev.	40261.27	6712.42	77020.84	0.022
62, N=7,657	mean	53294.76	26837.35	292728.30	-0.122
	median	45939.05	30263.92	319571.90	-0.123
	std. dev.	42314.43	7054.66	78509.95	0.019
63, N=4,565	mean	52857.83	26715.18	284888.70	-0.126
	median	45611.39	30402.21	313239.80	-0.125
	std. dev.	42908.46	7254.17	78547.99	0.019
64, N=2,793	mean	52823.81	26673.15	277747.00	-0.095
	median	45096.35	30297.56	309250.10	-0.096
	std. dev.	44722.16	7308.58	77052.34	0.018
65, N=1,787	mean	40839.73	27492.06	279203.10	-0.096
	median	31857.40	32011.43	315836.70	-0.096
	std. dev.	38886.03	7735.31	79230.61	0.017

	Percentile				
	10	25	50	75	90
Assets	0.00	4930.6	24884.33	76296.25	160007.4

Notes: The statistics shown for earnings, annual benefits, SSW and assets are in 2003 euros. Annual earnings are computed based on the calendar year that an individual reaches the specified age. SSW is computed assuming β=.93. The asset statistics are based on household gross financial assets from SHARE-Austria data. We use information from 1,465 individuals ages 50 through 54 from the SHARE-Austria data.

Table 3

(A) Hazard Model Estimates

	All Ages		Ages 60 & 65		Time-Varying Covariates	
	Base Controls	Full Controls	Base Controls	Full Controls	Base Controls	Full Controls
β_{SSW}	0.4389	0.4013	0.3253	0.2626	0.1097	0.1018
	(0.0775)	(0.0962)	(0.0402)	(0.0466)	(0.0125)	(0.0138)
β_{ACC}	-2.8972	-3.3815	-2.8575	-3.7060	-0.4804	-0.4334
	(0.8502)	(1.6025)	(1.1183)	(1.5927)	(0.1857)	(0.2683)
$\frac{\beta_{SSW}}{-\beta_{ACC}}$	0.151	0.119	0.114	0.0708	0.228	0.235
	(0.0477)	(0.0557)	(0.0389)	(0.0312)	(0.0857)	(0.142)
$\beta_{SSW\|60}$	0.439	0.401	0.325	0.263	0.658	0.611
$\beta_{ACC\|60}$	-2.897	-3.381	-2.857	-3.706	-2.883	-2.600

(B) Risk Aversion Estimates based on Hazard Model Estimates

	Elasticity Ratio (β_{SSW}/β_{ACC})				
	-0.25	-0.20	-0.15	-0.10	-0.05
γ	0.71	0.57	0.43	0.28	0.14

Notes for panel (A): Estimates are based on a sample of 1,101,444 observations from 252,907 individuals. Standard errors clustered by year are shown in parentheses. All coefficient estimates should be interpreted as changes in the baseline retirement hazard. All specifications include the following base controls: education dummies, a quadratic polynomial in tenure, and quartic polynomials in calendar year, log annual earnings, and log total earnings in the prior 10 years. All specifications also include a censored dummy (current tenure begun in 1972 or earlier) and the interactions between this dummy and each of the severance pay and tenure variables. The full controls specifications include the base controls, industry and region dummies, and quartic polynomials in log earnings from each of the prior 10 years.

Notes for panel (B): These calculations are based on a consumption difference between retirement and employment of 0.395 (=(14000+0.05*25000-.9*28000)/(.9*28000)) and a wage-to-asset ratio of roughly 1.12 (=28000/25000). The numbers used in these calculations come from the medians listed in the summary statistics (Table 2). Please see text for more details.

Table 4
Structural Estimates
Estimation based on Matching Retirement Hazard Rates and Proportional Hazard Coefficients

	Baseline	Alternative Policy Function Approximation	Fixed Savings Rate (s = 0.10)	Lower Discount Factor (β = 0.88)
(A) Parameter Estimates:				
Curvature of Consumption Utility: γ	0.709	0.707	0.953	0.677
	[0.485, 0.814]	[0.678, 0.840]	[0.732, 1.661]	[0.505, 1.027]
Distribution of Work Disutility: η	0.936	1.095	1.159	0.572
	[0.738, 1.155]	[1.085, 1.362]	[0.779, 1.360]	[0.407, 0.790]
Slope of Work Disutility: α	0.307	0.253	0.143	0.484
	[0.230, 0.433]	[0.170, 0.256]	[0.086, 0.197]	[0.304, 0.703]
Disability Pension Fixed Cost: κ	96666.20	110664.44	128703.58	80463.86
	[89184.25, 127346.09]	[108992.13, 132442.70]	[124441.80, 202252.52]	[75420.95, 110506.73]
(B) Proportional Hazard Coefficients:				
Coefficients with All Ages				
β_{SSW}	0.4676	0.3548	0.6486	0.4451
β_{ACC}	-3.1138	-3.0791	-2.5944	-3.117
$\frac{\beta_{SSW}}{-\beta_{ACC}}$	0.1502	0.1152	0.2500	0.1428

Notes: 95% confidence intervals are shown in brackets below the parameter estimates; confidence intervals are based on the bootstrapped distributions of parameter estimates that were calculated using 100 replications in which 75000 individuals were drawn with replacement. Estimates are based on a the same sample used to estimate the proportional hazard specifications in Table 3. The baseline specification is based on a discount factor of β = 0.93, a real interest rate of r = 0.075, and a fixed wage growth rate of 0.0175. Please see the text for more details.

Table 5

Estimation of γ

	0.50(γ)	0.75(γ)	**1.00(γ)**	1.50(γ)	2.00(γ)	3.00(γ)
β_{SSW}	0.5244	0.4922	**0.4676**	0.3996	0.4036	0.4528
β_{ACC}	-4.4805	-3.5717	**-3.114**	-2.3793	-1.7629	-1.2194
$\frac{\beta_{SSW}}{-\beta_{ACC}}$	0.1170	0.1378	**0.1502**	0.1679	0.2289	0.3713

Notes: All values are based on parameters estimates presented in column 1 of Table 4. Please see the test for more details.

Table 6

Policy Simulation: Reform 4 - Phased-In Eligibility

Age	Claiming Cost	Claiming Cost as % of Median Annual Benefits
≤ 59	κ	0.35
60	(0.95)κ	0.31
61	(0.85)κ	0.24
62	(0.70)κ	0.20
63	(0.50)κ	0.14
64	(0.25)κ	0.07
≥ 65	0	0.00

Notes: Please see column 1 of Table 4 for the estimated value for the claiming cost κ. Median annual benefits used to compute the percentages in the third column are based on the summary statistics presented in Table 2.

Table 7
Welfare Analysis of Policy Simulations

	Baseline Welfare: Expected Utility at Age 55		
	Mean	Median	Variance
Baseline	635.3104	634.4412	8940.03
	Changes Relative to Baseline		
	Mean	Median	Variance
Reform 1: 20% Decrease in Benefits	0.9645	0.9635	1.0335
Reform 2: -3%*(65-Age) Change in Benefits	0.9708	0.9699	1.0639
Reform 3: 1 Year Increase in Eligibility Age a_{old}	0.9872	0.9870	1.0783
Reform 4: Increase in Cost of Claiming Disability κ	0.9496	0.9527	1.2767

Notes: The baseline values are computed based on the parameter values presented in the first column of Table 5. Please see the text for more details.

Table A1
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
8. Age < 55, or Age > 65 in 1984 - 2003, & Age > Claim Age (if Claiming)	355805	39129
9. Missing Pension Variables & First Observed at Age > 55	254130	101675
10. Outliers & Missing Earnings and Industry Data	252907	1223
Final Sample, Ages 55-65 in Years 1984-2003, First Observed at Age 55		
# of Individuals	252907	
# of Claimants	178997	
# of Observations	1101443	

Notes: The number of claimants in the final sample is less than the number of individuals in the sample since younger individuals in the later years of the sample have yet to claim pensions. Further details regarding the samples and restrictions are contained in the text.