

Manoli, Dayanand; Mullen, Kathleen; Wagner, Mathis

Working Paper

Policy Variation, Labor Supply Elasticities, and a Structural Model of Retirement

IZA Discussion Papers, No. 8659

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Manoli, Dayanand; Mullen, Kathleen; Wagner, Mathis (2014) : Policy Variation, Labor Supply Elasticities, and a Structural Model of Retirement, IZA Discussion Papers, No. 8659, Institute for the Study of Labor (IZA), Bonn

This Version is available at:

<https://hdl.handle.net/10419/106591>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

IZA DP No. 8659

**Policy Variation, Labor Supply Elasticities,
and a Structural Model of Retirement**

Day Manoli
Kathleen J. Mullen
Mathis Wagner

November 2014

Policy Variation, Labor Supply Elasticities, and a Structural Model of Retirement

Day Manoli

UT-Austin and NBER

Kathleen J. Mullen

RAND and IZA

Mathis Wagner

Boston College

Discussion Paper No. 8659
November 2014

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0

Fax: +49-228-3894-180

E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Policy Variation, Labor Supply Elasticities, and a Structural Model of Retirement^{*}

This paper exploits a combination of policy variation from multiple pension reforms in Austria and administrative data from the Austrian Social Security Database. Using the policy changes for identification, we estimate social security wealth and accrual elasticities in individuals' retirement decisions. Next, we use these elasticities to estimate a dynamic programming model of retirement decisions. Finally, we use the estimated model to examine the labor supply and welfare consequences of potential social security reforms.

JEL Classification: J26, H55

Keywords: policy variation, retirement, labor supply elasticities

Corresponding author:

Kathleen Mullen
RAND Corporation
1776 Main Street
P.O. Box 2138
Santa Monica, CA 90407-2138
USA
E-mail: kmullen@rand.org

^{*} 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. This paper was supported by the National Institute on Aging through grant number 1 R03AG033833. Its contents are solely the responsibility of the authors and do not represent the official views of the National Institute on Aging.

1 Introduction

In countries around the world, there is increasing pressure for social security reform (OECD 2007). Designing effective social security reform requires understanding how changes in retirement benefits affect individuals' retirement decisions. In this paper, 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 labor supply responses to potential social security reforms.¹

Retirement benefits are traditionally thought to affect individuals' behavior through two channels: an income effect and a price effect (see, e.g., Boskin (1977)). The income effect refers to changes in behavior due to changes in lifetime income. The price effect (or implicit tax on earnings) refers to changes in behavior due to changes in marginal incentives for continued work. The magnitudes of these effects are relevant for understanding how potential social security reforms are likely to affect individuals' retirement decisions and welfare.

The analysis has two overall objectives. A first objective is to identify and estimate these income and price elasticities. We use administrative data from the Austrian Social Security Database and exploit variation from multiple pension reforms in Austria to identify and estimate these elasticities. A second objective is to demonstrate how these income and price elasticities can be used to estimate a structural model of retirement decisions.

The empirical analysis is therefore 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 percentage change in social security wealth from delaying

¹Recent research has emphasized the identification of income and price effects of benefits from other social insurance programs; for unemployment insurance see Chetty (2008), for disability insurance see Autor and Duggan (2007), and for health insurance see Nyman (2003).

retirement by one year). The pension reforms create several independent changes in these measures thereby allowing for separate identification of the respective elasticities, which we estimate using a proportional hazards specification. We follow commonly used terminology and refer to the proportional hazards results as reduced-form elasticities. However, 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.

We estimate that a 1% increase in social security wealth holding the accrual constant (i.e. an increase in social security wealth at all ages) increases the hazard rate at a given age by 0.44%. At the same time, we estimate that a 1% increase in the accrual holding social security wealth constant (i.e. an increase in next period's social security wealth holding current social security wealth constant) decreases the hazard rate at a given age by 2.90%. Our estimates point to a much larger role for price effects than has previously been found in the literature. The studies by Gruber and Wise (2004) and Coile and Gruber (2007) provide the most directly comparable labor supply estimates (although their incentive (price) measures are parameterized differently) and they tend to find small and insignificant price effects, whereas we find large and significant accrual elasticities, both in absolute terms and relative to the estimated income elasticity. An important difference between their study and ours is that they rely on observational variation whereas we exploit policy variation from multiple pension reforms that independently vary the income and price measures. Interestingly, Friedberg (2000) which exploits multiple changes to the Social Security earnings test also finds a large role for price effects in determining hours worked among retirees in the U.S.²

The second part of the empirical analysis focuses on structurally estimating a dynamic model of retirement decisions using the elasticity results from the first part of the empirical analysis. We estimate the model using an Indirect Inference (II) estimation strategy in which

²Other studies which rely on policy variation tend to exploit a single reform (e.g., Krueger and Pischke (1992) exploit the Social Security “notch” in the U.S.; Lumsdaine, Stock and Wise (1992) and Pencavel (2001) examined responses of employees who were offered a temporary incentive to retire early; Brown (2013) studies a reform to the California teachers pension system which created kinks in return to work).

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 retirement outcomes from the model. Therefore, the II estimation algorithm seeks structural parameters that match the elasticities based on the model’s simulated data to the actual 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 (γ) equal to 0.66 in the baseline specification of the model. While this is a relatively low estimate of γ compared to the previous literature (e.g., French (2005) and van der Klaauw and Wolpin (2008) estimate γ to be between 2.2. and 5.1 and roughly 1.6, respectively)³, we demonstrate that in our model a higher value of γ inhibits the model’s ability to predict the estimated social security wealth and accrual elasticities. Furthermore, the estimates of γ from this current analysis are consistent with previous analyses in which γ is identified based on labor supply elasticities, which is the approach we take. 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 1.78.⁴

Finally, 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 that individuals appear sensitive to changes in effective wages that arise due to changes in eligibility for benefits. Additionally, the simulations indicate that marginal changes to benefits conditional on being eligible do not seem to have significant impacts on retirement decisions. Thus, reforms that introduce or refine actuarial adjustments to benefits are likely to have smaller effects on retirement decisions relative to reforms that increase eligibility requirements (i.e. increase the Early Retirement Age).

³Similarly, Hubbard, Skinner and Zeldes (1995) use a preferred value of $\gamma = 3$ in their calibration exercise, while Blau (2008) calibrates $\gamma = 2$.

⁴These 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.

2 Institutional Background & Data

2.1 The Austrian Pension System and Reforms

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 earnings caps to earnings in each year, the capped annual earnings 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. Contribution years correspond only to periods of employment. Prior to 2001, disability pensions were computed identically to old-age pensions; in 2001 and after, the pension coefficient used in the disability pension was reduced relative to that of the old-age pension.⁵

By claiming a retirement pension, the individual essentially exits the labor market.⁶ Men are first eligible for old-age pensions at age 60 which is therefore referred to as the early retirement age.⁷ In addition to being at least age 60, an individual who claims an

⁵The reduction in the disability pension coefficient is based on insurance years with lower insurance years receiving larger reductions.

⁶Within one year after claiming a pension, nearly all men exit the labor market and never work again. As a result, we focus on the pension claiming decision as an exit from the labor market.

⁷We focus on men for two reasons. First, women have different statutory retirement ages in the time period we study. Second, because maternity spells are not fully observed in our data, we cannot calculate

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). However, a disability pension can be claimed prior to eligibility for a retirement pension, provided that the claimant can be classified as disabled.⁸

Figure 1A presents retirement hazard rates by age. When computing the hazard rates, failure is defined as 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 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. The survival rate at a given age measures the fraction of the population that has not yet claimed a pension at that age and hence remains in the labor market. This figure also highlights 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 1C focuses more directly on disability pensions by presenting the survival function for individuals who claim disability pensions. In particular, we restrict the sample to those who ultimately claim disability pensions and then compute survival rates. In this case, the survival rate measures the fraction of disability claimants that have not yet claimed a disability pension at that age. 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.

Between 1984 and 2003, there were five significant pension reforms in Austria in 1985, 1988, 1993, 1996 and 2000 which generally reduced the generosity of the retirement pension system as government officials felt the pension system was not financially sustainable. Our detailed knowledge of these reforms and the computation of the pensions is based on Marek

retirement benefits for women.

⁸The requirements for claiming a disability pension are relaxed for those age 55 and older. This was raised to age 57 in 1997; see Staubli (2011) for more details.

(1985, 1987-2003).⁹ Table A1 presents a summary of each reform. Figure A1 presents benefits-versus-age profiles for different calendar years that illustrate the variation in pension benefits created by the pension reforms. See Appendix A for details.

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. 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 in Figure A1, these reforms decreased the levels of benefits across potential retirement ages but left the slopes in the profiles unchanged.

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

⁹Ney (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.

The 1996 reform introduced a bonus/malus system to discourage early retirement (before the statutory age) by penalizing early retirees with reduced pension coefficients. 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. Since nominal adjustments in later years were lower than inflation, real benefits declined between 1998 and 2002.

2.2 Data and Variable Construction

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. See Appendix B for more details. 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

$$SSW_{i,a} = \sum_{t=a}^{100} \beta^{t-a} \pi_{t|a} y_{i,a}^R$$

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.¹⁰ In this definition, we also assume that the maximum age that individuals can live to is age 100. Each individual's retirement pension is

¹⁰The 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.

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 1 presents summary statistics by age for key variables 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), 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).¹¹ This SHARE dataset has wealth data for individuals in several European countries. We focus on the data collected for Austria in 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 1. The data indicates that households have accumulated financial assets roughly equivalent to one-year’s earnings.

3 Empirical Analysis I: Estimation of Elasticities

3.1 Proportional Hazards Specification & Identification Strategy

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,

$$R_i(a) = \bar{R}(a) \exp\{\beta_{SSW} \ln(SSW_{i,a}) + \beta_{ACC} \ln(1 + ACC_{i,a}) + \delta X_{i,a}\}.$$

In this specification, $R_i(a)$ denotes the relative hazard for individual i at age a . The relative hazard is the probability that individual i retires at age a conditional on not having retired at an earlier age relative a baseline probability across all individuals at age a . The term $\bar{R}(a)$ denotes the baseline hazard rate at age a . This baseline hazard is common across individuals at each age and thus the intuition regarding the baseline hazard closely follows the intuition of age fixed effects in a linear model. As defined above, $SSW_{i,a}$ 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

¹¹Information on the SHARE dataset can be found at <http://www.share-project.org/>.

by an additional year).¹² 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 (2007), 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. 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 rate.

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.¹³ 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. In a setting with observational variation alone, including controls for inputs into the benefit formula, such as polynomials in individuals' earnings histories, will control for any unobserved heterogeneity that is correlated with these variables; this is the widely used

¹²Because the accrual is often negative, it is necessary to add 1 when taking logs.

¹³Appendix C presents graphical evidence highlighting this identifying variation.

control function approach (Heckman and Robb 1985). In this case, identification of β_{SSW} and β_{ACC} relies on precise measurement of these variables and, importantly, the existence of nonlinearities in the benefit formula. If benefits are close to linear in earnings history variables, as is the case in many countries including Austria and the United States, then elasticity estimates using observational variation may be fragile. A solution to this problem - which we employ in this paper - is to find and exploit exogenous variation in pension benefits with respect to earnings history and other control variables. Specifically, the five pension reforms in Austria created five new benefit *formulas*. Importantly, these changes to the benefit formula are independent of individuals' past work decisions, so that individuals with identical earnings histories (and, presumably, identical preferences for work) could have different retirement benefits and incentives if they were born in different years.

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 mostly to the pension reforms (but also to some extent to nonlinearities in the benefit formula), 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.

3.2 Hazard Model Results

The results from the Cox proportional hazards model are presented in Table 2. 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% . We estimate much higher price effects than wealth effects, on the order of 6-7 times higher. While these elasticities are not directly comparable to one another due to differences in units, even when scaling the accrual elasticity by the average accrual rate (roughly 0.10) to reflect changes in the social security accrual itself (i.e., euros), our estimates point to a much larger role for price effects than has been previously found. For example, Coile and Gruber (2007) estimate elasticities of 0.16 and -0.003 with respect to SSW and accrual (in levels), respectively; even their estimated elasticity with respect to their “peak value” incentive measure of -0.07 points to a smaller role for incentive effects than our estimates do, both in absolute terms and relative to the SSW elasticity.

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). 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. As a result, we will consider the coefficients from the model estimated on all ages (column 1) to be our baseline estimates.

4 A Structural Model of Retirement

In this section, we develop a simple dynamic programming model of retirement decisions with uncertainty relating to mortality and job separations.¹⁴ The dynamic model that we develop is closely related to previous work in the literature (Stock and Wise 1990, Lumsdaine, Stock and Wise 1992, Berkovec and Stern 1991, Rust and Phelan 1997). 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. At the beginning of a period, the individual knows his assets, retirement benefits, wage and disutility of work. If he chooses to retire, the individual receives his annuitized retirement benefits and faces no remaining uncertainty from the labor market. If he chooses to continue working, the individual receives his wage, experiences his disutility of work and takes into account the value of retirement decisions at future ages.

Consider first the optimization problem for an individual who has chosen to retire. Let $R_a(\Omega_a)$ denote the value of retirement at age a for an individual who enters the period with state variables Ω_a . The state variables reflect all information that is known when the individual enters the period; we will describe these state variables as we describe the value functions. Once an individual has chosen to retire, the individual solves the following optimization problem that defines the value of retirement:

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

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 $u(\cdot)$

¹⁴It is important that we capture uncertainty related to job separations because of the Austrian severance payment system, which confers one-time lump-sum payments to employees at the time of retirement. Manoli and Weber (2011) highlight the impacts of severance payments on individuals' retirement decisions in Austria.

captures utility over consumption with $u' > 0$ and $u'' < 0$. The maximization reflects that the individual chooses his consumption in each period. 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 monetized psychic cost of claiming disability (e.g., stigma) and/or the effort cost of proving one's qualifications as disabled. After age a_{old} , individuals are eligible for old-age pensions and therefore do not face this claiming cost.

The term y_a^R denotes the individual's retirement benefits at age a . These benefits are based on an annual payment from a government-provided pension, which is the focus of the series of pension reforms that we exploit for exogenous variation in benefits later in the paper. Following the institutional setting in Austria, upon retirement individuals also receive a one-time, employer-provided, lump-sum severance payment. Using τ^R to denote tenure at retirement and y^W to denote salary income from employment, the amounts of the severance payments $s(\tau^R)$ are as follows: 0 if $\tau^R < 10$, $\frac{1}{3}y^W$ if $\tau^R \in [10, 15)$, $\frac{1}{2}y^W$ if $\tau^R \in [15, 20)$, $\frac{3}{4}y^W$ if $\tau^R \in [20, 25)$ and y^W if $\tau^R \geq 25$.

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. 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, $E_a[D_{a+1}(\Omega_{a+1})]$ 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 $W_a(\Omega_a)$ denote the value of working at age a with assets A_a and work disutility v_a . This value function is defined as

$$\begin{aligned} W_a(\Omega_a) &= \max_{c_a^W} u(c_a^W) - v_a + \beta \pi_{a+1|a} E_a[D_{a+1}(\Omega_{a+1})] \\ \text{s.t. } c_a^W + \frac{A_{a+1}}{1+r} &= y_a^W + A_a. \end{aligned}$$

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

$$D_a(\Omega_a) = \max_{\text{retire, work}} \{R_a(\Omega_a), W_a(\Omega_a)\}.$$

In regard to heterogeneity, work disutility v is allowed to vary across individuals.¹⁵ The interest rate r , discount rate β , survival probabilities $\pi_{a+1|a}$, and consumption-utility function $u(\cdot)$ are restricted to be common across individuals.

The laws of motion related to retirement benefits are as follows. First, retirement benefits are fixed once an individual chooses his retirement date. Second, while continuing to work, 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. Additionally, we assume that job tenure evolves stochastically while the individual remains in the labor market. Let τ_a denote years of tenure at age a . With probability π_{sep} , the individual experiences a job separation so that $\tau_{a+1} = 0$; with probability $1 - \pi_{sep}$ the

¹⁵This parameter may also reflect differences in health across individuals, but unfortunately we do not observe health in our data set and so we do not model health transitions.

individual remains in his current job so that $\tau_{a+1} = \tau_a + 1$.¹⁶

Next we specify the laws of motion for earnings and the disutility of work. Since collective bargaining agreements effectively determine older workers' wages in Austria, we assume that earnings increase deterministically with age, $y_{a+1}^W = (1 + g)y_a^W$. With this assumption, we effectively assume that job separations affect an individual's tenure, but not his wages. Intuitively, if an individual becomes separated, then we assume that he can find a job that pays the same amount in wages, but he will have zero tenure at his new job and his potential severance payment will be reset to zero.

5 Empirical Analysis II: Structural Estimation

5.1 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 assets, we approximate the initial distribution of assets at age 54 using the Austrian Survey of Health, Ageing and Retirement in Europe (SHARE); specifically, we randomly sample initial assets for each individual with replacement from the empirical distribution of assets in SHARE.

We parameterize the model presented in Section 4 as follows. We assume constant relative

¹⁶In our data set we cannot distinguish between voluntary and involuntary job separations.

risk aversion (CRRA) utility over consumption:

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

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). We assume initial work disutility is unrelated to observed characteristics. Work disutility increases linearly with age, with slope $\alpha\tilde{\eta}$ (i.e., $v_a = \alpha\tilde{\eta}(a - 54) + v_{54}$). 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.

A common method for estimating θ is the Method of Simulated Moments (MSM). Moment-based estimation strategies match key moments (e.g., retirement hazard rates by age) observed in the sample data with the analogous moments implied by a model param-

eterized by θ . The goal is to find the value of θ which gives the best “fit” of the model, i.e., by minimizing the (weighted) distance between the observed and predicted moments. Where estimation of the predicted moments is computationally intractable by conventional methods, simulation methods must be employed. MSM approximates such moments using Monte Carlo integration, i.e., by averaging over simulations of the model. Assuming the moment conditions are correctly specified, MSM is consistent for a fixed number of simulations. Typical moments used to estimate structural retirement models are retirement hazard rates by age, which exploit differences in incentives for retirement across ages.

The approach we propose is a modified version of MSM based on the method of Indirect Inference (Gourieroux, Monfort and Renault 1993). Specifically, we include the estimated income and price elasticities β_{SSW} and β_{ACC} from the proportional hazards specification as empirical moments to be matched, along with the more traditional retirement hazard rates at ages 55 through 65. By matching the labor supply elasticities directly, we ensure that predicted labor supply responses from the estimated structural model are consistent with reduced form estimates. Moreover, our approach explicitly exploits quasi-experimental variation stemming from the pension reforms as an important source of identification. As described earlier, the reforms allow us to break apart the relationship between past earnings and benefits. Thus, we are able to observe individuals with identical earnings histories facing different incentives for retirement.

The method of Indirect Inference (II) can be described as follows. First, specify an auxiliary model. This model is ‘incorrect’ in the sense that optimization of the likelihood does not yield consistent estimates of the parameters of interest, θ . Instead, it provides consistent estimates of some auxiliary parameters β such that $\widehat{\beta}_N \xrightarrow{p} \beta(\theta_0)$ when the data are generated under the true value of $\theta = \theta_0$. This binding function $\beta = \beta(\theta)$ links the auxiliary (reduced-form) and structural parameters. II consists of finding the value of θ that minimizes the distance between $\widehat{\beta}_N$ (the estimates on the observed data) and an estimate of $\beta(\theta)$. This estimate is obtained by simulating S paths of the model conditional on θ and

estimating the auxiliary model on the simulated retirement decisions of the pooled sample (of size SN). Thus, the II estimator is the following minimum distance estimator:

$$\hat{\theta}_{II} = \arg \min_{\theta} \left(\hat{\beta}_N - \tilde{\beta}_{SN}(\theta) \right)' W \left(\hat{\beta}_N - \tilde{\beta}_{SN}(\theta) \right),$$

for some weighting matrix W . Gouriéroux et al. (1993) show that the II is a consistent estimator of θ for $N \rightarrow \infty$. Identification requires $\dim(\beta) \geq \dim(\theta)$, and the binding function should be a one-to-one function mapping θ to β .

We apply a variation on the II estimator by combining it with the more traditional MSM estimator in order to identify and estimate all the parameters in θ . Specifically, our II estimator matches the following moments: retirement hazard rates by age, and the coefficients β_{SSW} and β_{ACC} from the proportional hazard model above. 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 2).

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 . Additionally, because the data set is so large, it is extremely computationally burdensome to solve jointly for optimal consumption and retirement age. 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.

5.2 Structural Estimation Results

Table 3, 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 0.66. 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 3 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 $\beta_{SSW} = 0.41$ and $\beta_{ACC} = -3.01$. These elasticities predicted by the model are similar to the corresponding estimates using the real data, $\beta_{SSW} = 0.44$ and $\beta_{ACC} = -2.90$ (see Table 3).

As mentioned earlier, Chetty (2006) derives a general relationship between risk aversion and the ratio of labor supply elasticities with respect to wage and income, respectively. This motivates us to examine the relationship between the coefficient of relative risk aversion parameter γ in the structural model and the estimated labor supply elasticities β_{SSW} and β_{ACC} in our setting. Table 4 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. 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 be roughly 102,000 euros. To put this cost in perspective, the social security wealth (expected present 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 3 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 3 indicate that in this specification, γ continues to be estimated to be relatively low at 0.75.

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 3 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 0.57. 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 2 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 stand out. 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 3 are entirely consistent with estimates of γ implied by previously estimated income and price elasticities (see Chetty (2006)).

6 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 two hypothetical pension reforms that mimic features of commonly discussed and previously implemented pension reforms in several countries. 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.

To examine the labor supply consequences of these reforms, Figure 3A presents the simulated hazard rates under each reform. The simulations indicate 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 to the second reform highlight that reforms which make benefits actuarially fair in Austria will not lead to significant increases in retirement at older ages.

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. Figure 3B 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. Thus, the 20% reduction in benefits does not lead to a dramatic reduction in welfare. We estimate that average expected utility at age 55 is reduced by approximately 3.6% as a result of reform 1. 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, by about 3%, but the variance of the distribution increases by 7% relative to the baseline (compared with a 3.7% increase in variance for reform 1). 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.

Finally, to examine the effect of changes in eligibility vs. the benefit formula, we simulate a third reform in which there 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. Not surprisingly, the increase in the early retirement age leads to a shift in the spike in the hazard rates from age 60 to age 61 (not shown). However, the welfare results for the the third pension reform indicate that, while this reform leads to noticeable labor supply changes, the reduction in expected utilities is less dramatic. We find that shifting the early retirement age reduces average expected utility by only 1.3% relative to baseline, with the variance in expected utility increasing by 7.7%—comparable to the second reform which shifts benefits from before to after the full retirement age. 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 very 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.

7 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. Based on a proportional hazards specification, we estimate a Social Security Wealth (SSW) elasticity of 0.43 and accrual elasticity of -2.90 , which points to a larger role for price (marginal incentive) effects than previous estimates. 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. 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.

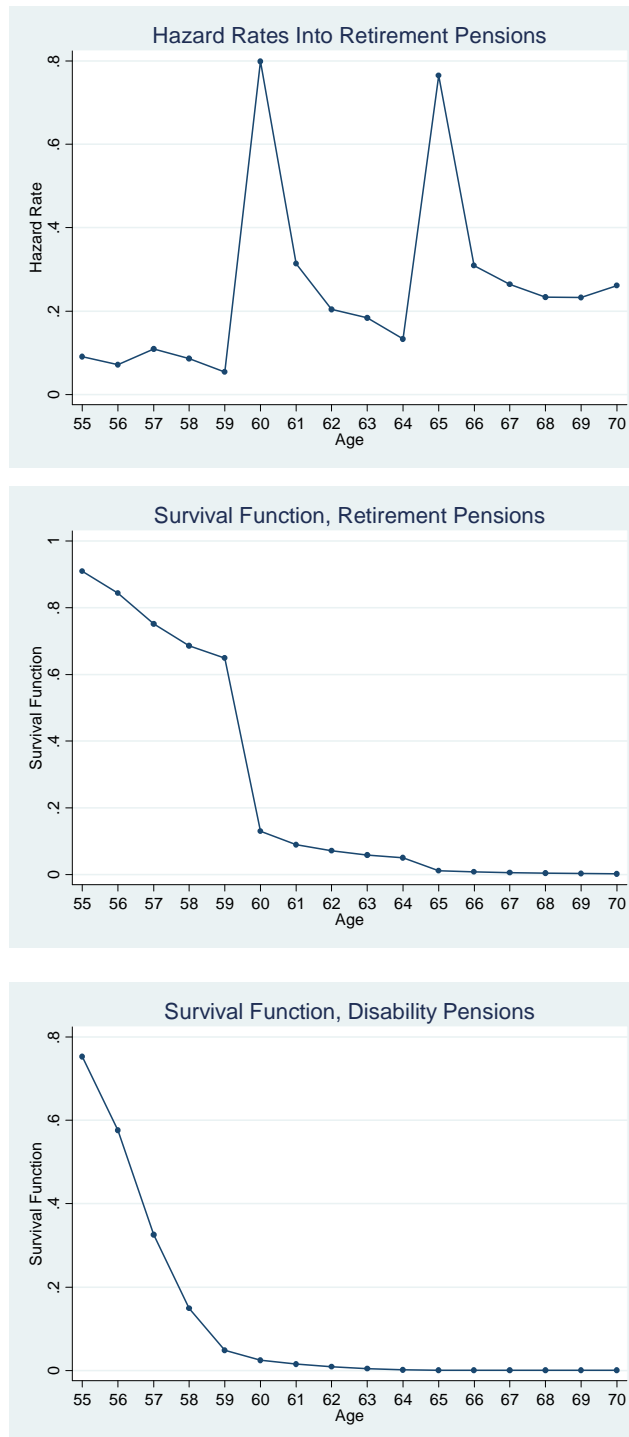
Our results have important implications for understanding social security reform. The relative importance of price effects indicates that incentive effects from benefit schedules across potential retirement ages greatly impact retirement behavior for inframarginal individuals (i.e., individuals not constrained by eligibility requirements). Policy simulations from the estimated structural model based on matching these labor supply elasticities imply a much larger role for policies that affect early eligibility ages. This is likely because, given the generosity and benefit structure of the Austrian pension system, many individuals would like to retire as early as feasible, given their budget constraint. Thus, changes in eligibility requirements have substantial impacts on individuals' effective net wages and individuals' retirement decisions appear sensitive to these changes.

References

- Autor, D. and M. Duggan (2007). Distinguishing Income from Substitution Effects in Disability Insurance. *American Economic Review Papers and Proceedings* 97(2), 119.
- Berkovec, J. and S. Stern (1991). Job Exit Behavior of Older Men. *Econometrica* 59(1), 189-210.
- Blau, D. (2008). Retirement and consumption in a life cycle model. *Journal of Labor Economics* 26(1), 35-71.
- Boskin, M. J. (1977), Social Security and Retirement Decisions. *Economic Inquiry*, 15: 1–25.
- Brown, K. (2013). The link between pensions and retirement timing: Lessons from California teachers. *Journal of Public Economics*, 98, 1-14.
- 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 (2007). Future social security entitlements and the retirement decision. *Review of Economics and Statistics* 89(2), 234-246.
- French, E. (2005). The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour. *Review of Economic Studies* 2(72), 395.
- Friedberg, L. (2000). The Labor Supply Effects of the Social Security Earnings Test. *Review of Economics and Statistics* 82(1), 48-63.
- Gourieroux, C., A. Monfort, and E. Renault (1993). Indirect inference. *Journal of Applied Econometrics* 8(S1), S85-S118.
- Gruber, J. and D. Wise (2004). Social Security Programs and Retirement around the World: Micro-Estimation. University of Chicago Press.
- 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.
- Hubbard, R. G., J. Skinner, and S. Zeldes (1995). Precautionary saving and social insurance. *Journal of Political Economy* 103(2), 360-399.
- Koman, R., U. Schuh, and A. Weber (2005). The Austrian Severance Pay Reform: Toward a Funded Pension Pillar. *Empirica* 32(3), 255-274.
- Krueger, A. and J. Pischke (1992). The Effect of Social Security on Labor Supply: A Cohort Analysis of the Notch Generation. *Journal of Labor Economics* 10(4), 412-437.
- Linnerooth-Bayer, J. (2001). Pension Reform in Austria. International Institute for Applied System Analysis. Report to the PEN-REF project.
- Lumsdaine, R., J. Stock, and D. Wise (1992). Three Models of Retirement: Computational Complexity Versus Predictive Validity. in D. Wise, ed. Topics in the Economics of Aging, University of Chicago Press.
- 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.

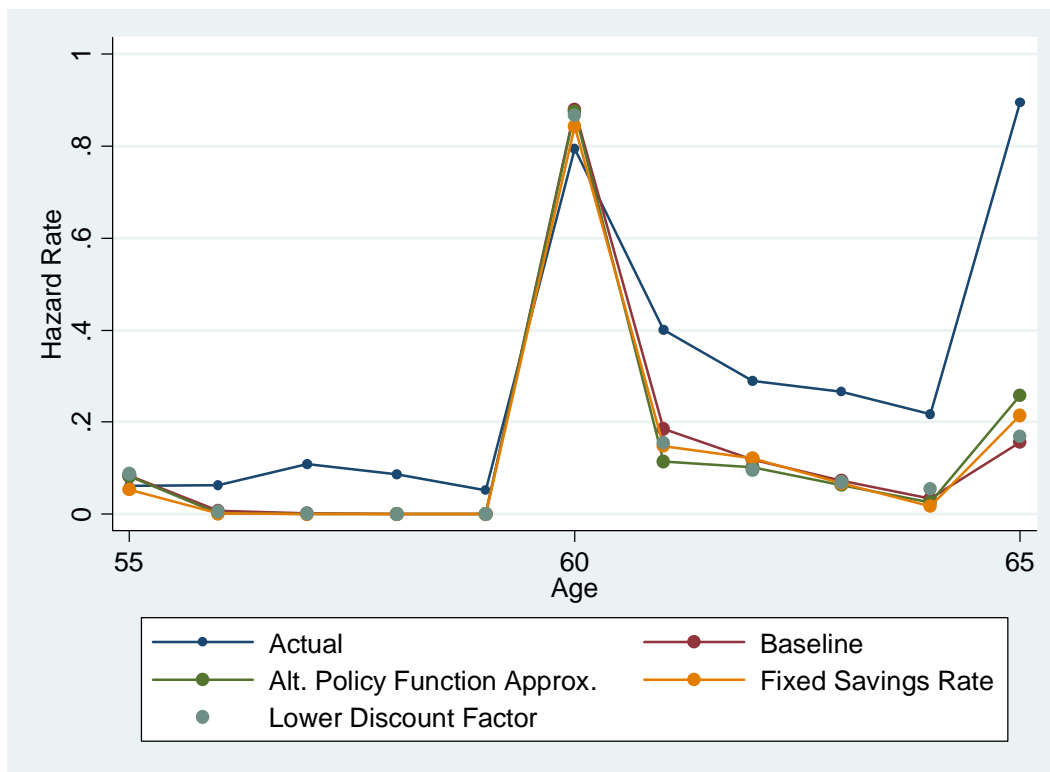
- Pencavel, John. The Response of Employees to Severance Incentives: The University of California's Faculty, 1991-94. *The Journal of Human Resources*, 2001, 36(1), pp.58-84.
- Rust, J. and C. Phelan (1997). How Social Security and Medicare Affect Retirement Behavior In a World of Incomplete Markets. *Econometrica* 65(4), 781-831.
- Staubli, S. (2011). "The Impact of Stricter Criteria for Disability Insurance on Labor Force Participation." *Journal of Public Economics*, 95 (9-10), 1223–1235.
- 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.

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). In the top plot, the retirement hazard is computed by defining the failure event as claiming a retirement pension (either an old-age pension or a disability pension). In the middle plot, the survival function at a given age measures the fraction of individuals who have not yet claimed a retirement pension by that age. In the bottom plot, the sample is restricted to individuals who ultimately claim a disability pension. In this case, the survival function at a given age measures the fraction of disability claimants that have not yet claimed a disability pension by that age.

Figure 2: Structural Estimation, Predicted vs. Actual Hazard Rates



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

Figure 3A: Labor Supply Responses to Policy Simulations

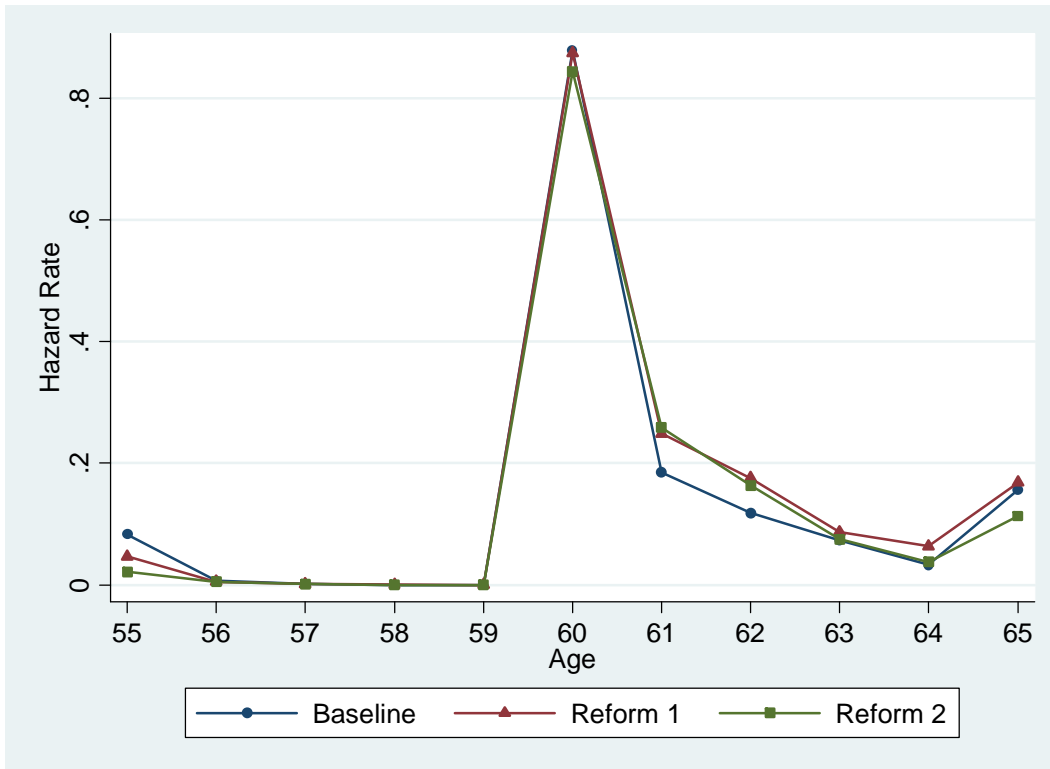
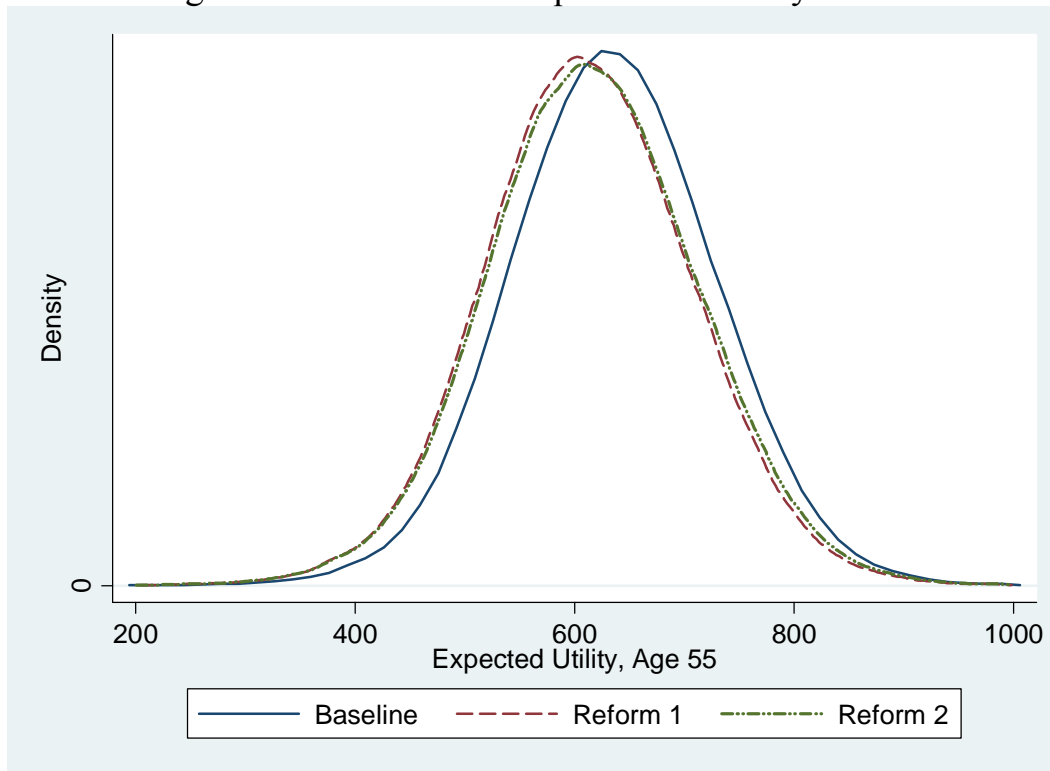


Figure 3B: Welfare Consequences of Policy Simulations



The baseline is computed using the parameter estimates reported in Table 4. Reform 1 decreases benefits across all ages by 20%. Reform 2 changes benefits at a give age based on $-3\% \cdot (65 - \text{age})$.

Table 1
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 $\beta=.93$. 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 2
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 (0.0775)	0.4013 (0.0962)	0.3253 (0.0402)	0.2626 (0.0466)	0.1097 (0.0125)	0.1018 (0.0138)
β_{ACC}	-2.8972 (0.8502)	-3.3815 (1.6025)	-2.8575 (1.1183)	-3.7060 (1.5927)	-0.4804 (0.1857)	-0.4334 (0.2683)
$\frac{\beta_{SSW}}{-\beta_{ACC}}$	0.151 (0.0477)	0.119 (0.0557)	0.114 (0.0389)	0.0708 (0.0312)	0.228 (0.0857)	0.235 (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

Notes: 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. Please see text for more details.

Table 3
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 ($\beta = 0.88$)
(A) Parameter Estimates:				
Curvature of Consumption Utility: γ	0.6649	0.7493	0.7957	0.5714
Distribution of Work Disutility: η	0.9611	1.2288	1.2867	0.4361
Slope of Work Disutility: α	0.3008	0.2087	0.1344	0.6982
Disability Pension Fixed Cost: κ	102018.94	116292.34	134748.04	68491.32
(B) Proportional Hazard Coefficients:				
Coefficients with All Ages				
β_{SSW}	0.4069	0.3473	0.5855	0.4575
β_{ACC}	-3.0064	-3.1644	-2.5491	-2.8995
$\frac{\beta_{SSW}}{-\beta_{ACC}}$	0.1353	0.1098	0.2297	0.1578

Notes: 95% confidence intervals are shown in parentheses below the parameter estimates; confidence intervals are based on the bootstrapped distributions of parameter estimates that were calculated using 100 replications in which 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 $\beta = 0.93$, 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 4
 Estimation of Ψ
 (Using Baseline Estimate $\Psi = 0.7064$)

	0.50(Ψ)	0.75(Ψ)	1.00(Ψ)	1.50(Ψ)	2.00(Ψ)	3.00(Ψ)
β_{SSW}	0.5292	0.5095	0.4863	0.4236	0.4296	0.4786
β_{ACC}	-4.3906	-3.5462	-3.0179	-2.3442	-1.6467	-0.9957
$\frac{\beta_{SSW}}{-\beta_{ACC}}$	0.1205	0.1437	0.1611	0.1807	0.2609	0.4807

Notes: All other parameter values are held constant at the corresponding baseline parameter estimates. Please see the test for more details.

A Details of Pension Reforms

This Appendix illustrates in detail how the pension reforms affected retirement benefits in Austria, summarized in Table A1. Figure A1 presents benefits-versus-age profiles for different calendar years that illustrate 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 earns the nominal equivalent of 20,000 euros in 2003 in the last 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. 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, 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 due 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 Figure A1. 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 the top panel of Figure A1. 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 increase 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.¹

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, 65.² 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

¹The 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.

²Since 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.

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. The bottom panel of Figure A1 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 the bottom panel of Figure A1, 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.

B Data & Sample Restrictions

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.³ 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.⁴

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.

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 re-

³Tenured public sector employees are observed only starting in 1988 or in some cases 1995, and income is not observed for self-employed individuals.

⁴Insurance months are determined using the following imputation for insurance years. Specifically, insurance years are imputed as $InsYrs = Age - Edu - 6 - (time\ observed\ not\ working)$ 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 $InsMths = 12 * InsYrs$.

tirement since these claims may not correspond to retirement decisions.⁵ 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. See Table B1.

⁵The 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.

C 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 C1 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 C1, 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 -0.096 to -0.087 , 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,⁶ 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

⁶Age 55 is the first age for relaxed entry into disability pensions for retirement. Thus, at age 55, it is possible to retire by claiming a disability pension, but age 60 is the first age an individual can claim an old-age pension.

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 C1 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.⁷ 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 C2. 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.


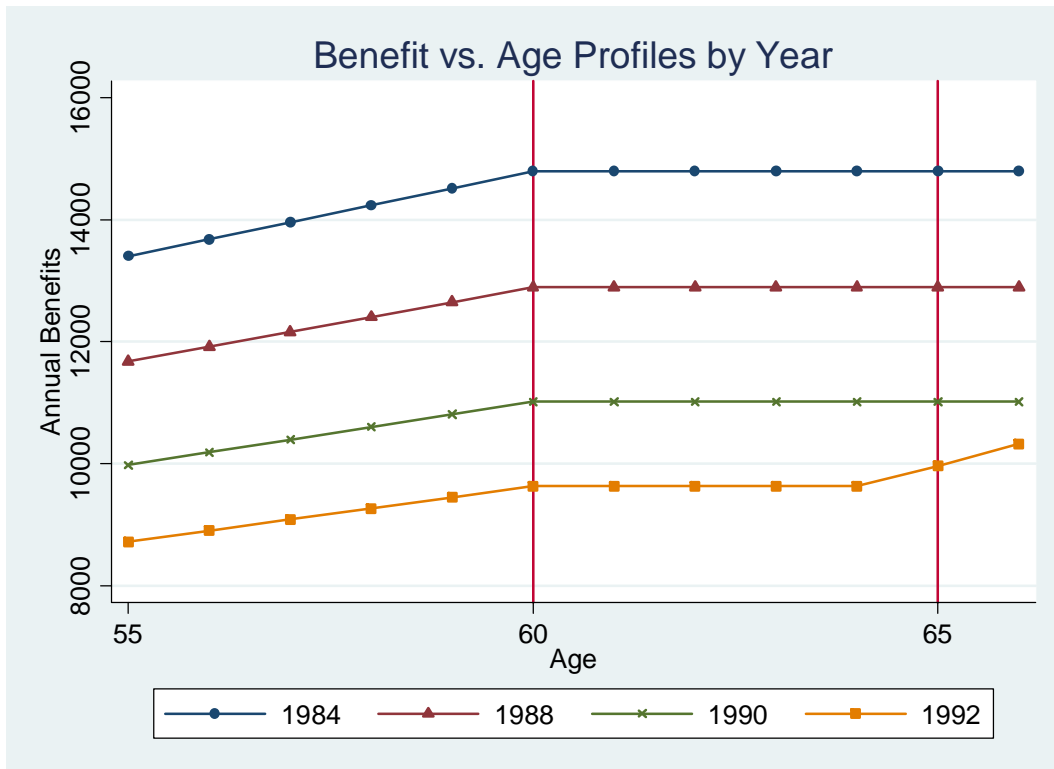
Consistent with Figure C1, the plots in Figure C2 indicate relatively smaller wealth effects and larger price effects. Specifically, the scatter plots show a steeper slope with the accrual residuals than with the social security wealth residuals (-4.784 versus 0.717). Furthermore, in a bivariate regression 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 estimate these elasticities more directly in the context of a Cox proportional hazards specification in Section 3.

⁷The regressions are of the form

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

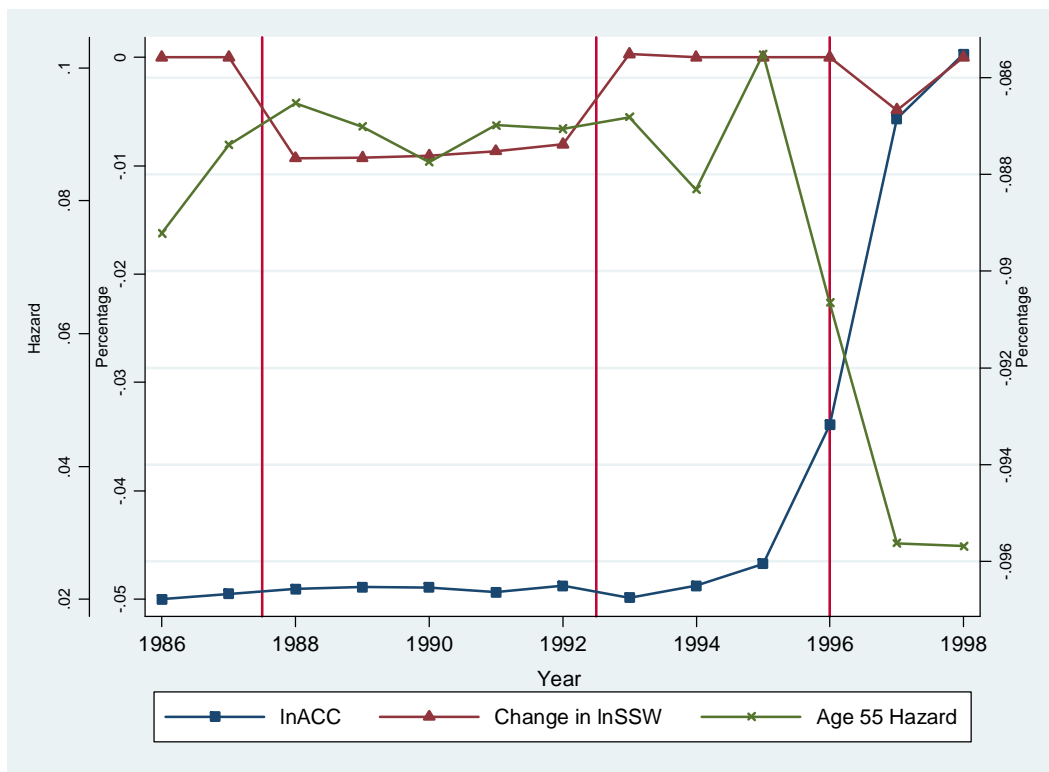
where the subscripts refer to individual i at age a . The change-in-eligibility dummies capture changes in the mechanical rules governing retirement that are independent from changes in financial incentives. As documented in Table A1, these changes are (1) the introduction of a disability pension at 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.

Figure A1: 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, nominal earnings at retirement 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 C1: Graphical Analysis of Wealth & Price Effects, Age 55



Notes: This figure is based on data at age 55 only. Please see Table B1 for the sample restrictions. The Age 55 Hazard (left-most vertical axis) measures the fraction of individuals claiming a pension at age 55. The Change in lnSSW (2nd vertical axis) measures the mean of the differences between SSW_t computed under the current year's legislation and SSW_t computed under the previous year's legislation, i.e.

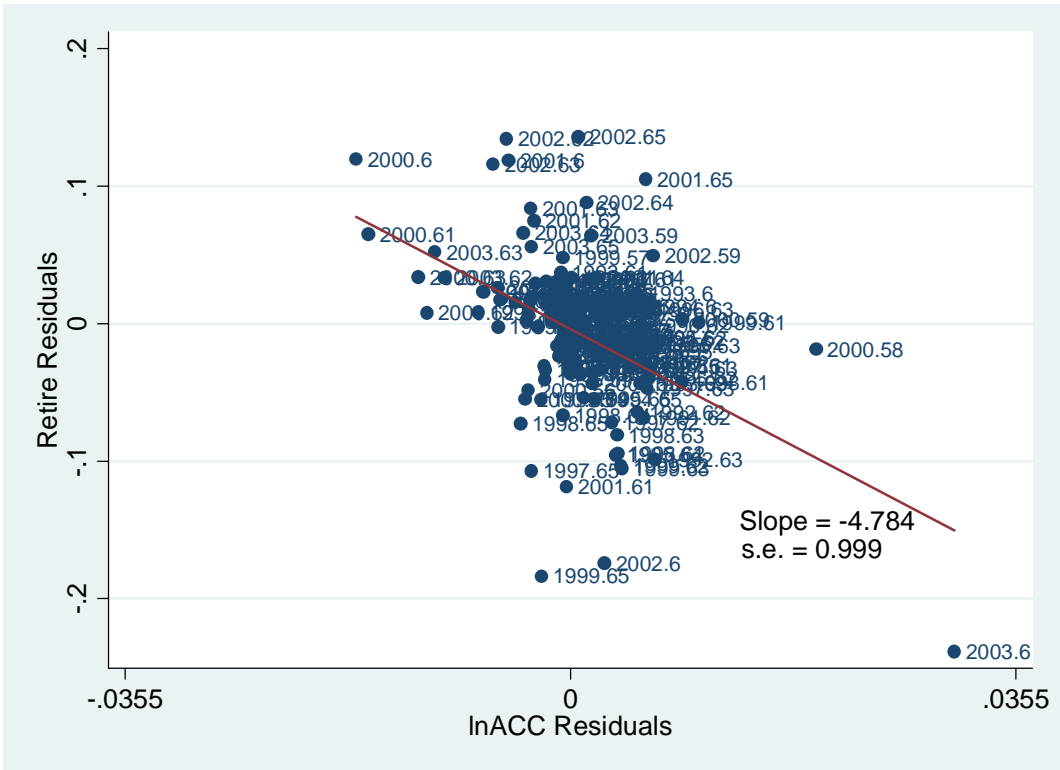
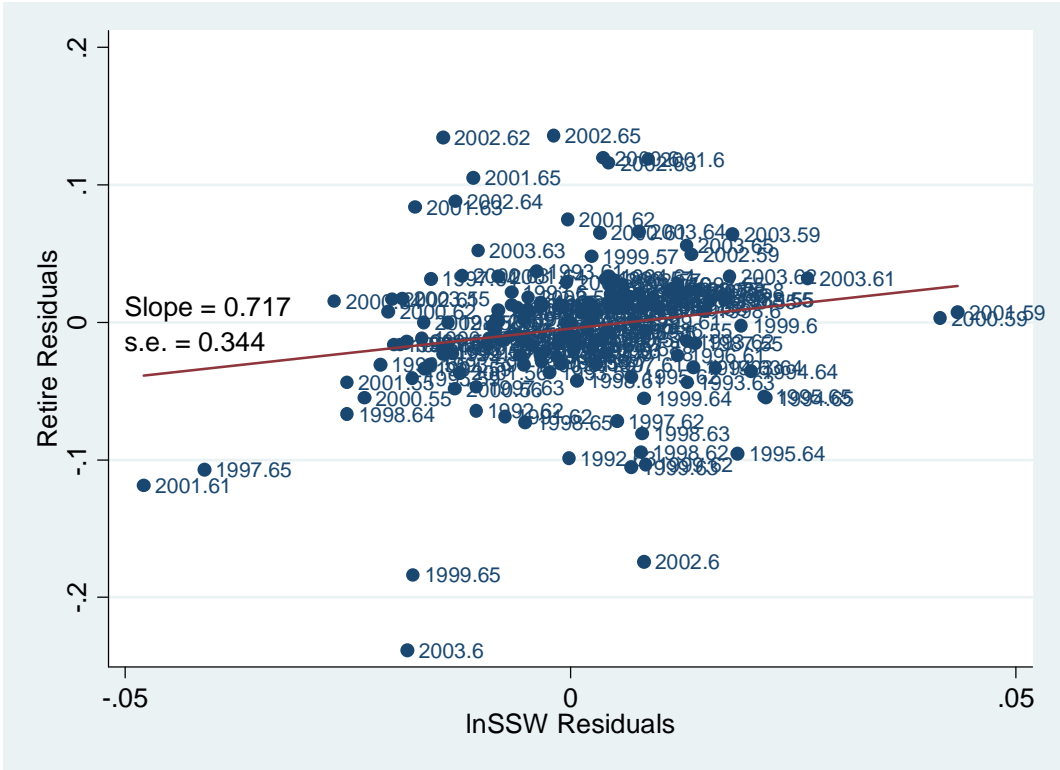
$$\text{Change in } \ln SSW_t = E[\ln(SSW_t(\text{legislation at time } t)) - \ln(SSW_t(\text{legislation at time } t-1))]$$

The lnACC variable (right vertical axis) measures the mean of accrual within a given year; the accrual captures the change in SSW_t if the individual were to delay retirement one additional year, i.e.

$$\ln ACC_t = E\left[\ln\left(1 + \frac{E_t(SSW_{i,t+1}) - SSW_{i,t}}{SSW_{i,t}}\right)\right]$$

To construct the plot, means of each of the variables are computed over all observations within a given calendar year. The means are then plotted across calendar years (horizontal axis).

Figure C2: 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.

Table A1
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 penalties 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 B1
Sample Restrictions, Initial Sample (Males, Birth Cohorts \geq 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
<hr/>		
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.