Economic Inquiry

POLICY VARIATION, LABOR SUPPLY ELASTICITIES, AND A STRUCTURAL MODEL OF RETIREMENT

DAY MANOLI, KATHLEEN J. MULLEN and MATHIS WAGNER*

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* J26, H55)

I. INTRODUCTION

In countries around the world, there is increasing pressure for social security reform (Organisation for Economic Co-operation and Development 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 (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 the individuals' retirement decisions and welfare.

The analysis has two overall objectives. The 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. The 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

*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.

Manoli: Department of Economics, University of Texas at Austin, Austin, TX 78712; NBER, Cambridge, MA 02138. Phone 510-684-8978, Fax 510-684-8978, E-mail dsmanoli@austin.utexas.edu

Mullen: RAND, Santa Monica, CA 90407-2138; IZA, Bonn, 53113, Germany. Phone 310-393-0411 x6265, Fax 310-260-8156, E-mail kmullen@rand.org

Wagner: Department of Economics, Boston College, Boston, MA 02467. Phone 773-396-7897, Fax 773-396-7897, E-mail mathis.wagner@bc.edu

1. Recent 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).

ABBREVIATIONS

II: Indirect Inference
MSM: Method of Simulated Moments
SHARE: Survey of Health, Ageing, and Retirement in Europe
SSW: Social Security Wealth

Economic Inquiry
(ISSN 0095-2583)
Vol. 53, No. 4, October 2015, 1702–1717

doi:10.1111/ecin.12207
Online Early publication March 18, 2015

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 (SSW, the present discounted value of pension benefits) and the one-year accrual (the expected percentage change in SSW from delaying retirement by one year). The pension reforms create several independent changes in these measures thereby allowing for separate identification of the respective elasticities, 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 SSW holding the accrual constant (i.e., an increase in SSW 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 SSW constant (i.e., an increase in next period's SSW holding current SSW 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),who 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 United States.[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 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 that of 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 SSW and accrual elasticities. Furthermore, the estimates of γ from this current analysis are consistent with previous analyses in which γ is identified based on the labor supply elasticities, which is the approach we take. Chetty (2006) surveys estimates of wage and income elasticities from 33 previous studies and finds that the mean of the implied values of γ is 0.71 with a range of 0.15–1.78.[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

2. Other 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 United States; 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).

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

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

to reforms that increase eligibility requirements (i.e., increase the Early Retirement Age).

II. INSTITUTIONAL BACKGROUND AND DATA

A. 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 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 the 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 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 ages 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 because 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 (1985,

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

6. Within 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.

7. We 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 retirement benefits for women.

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

FIGURE 1
Hazard Rates & Survival Function

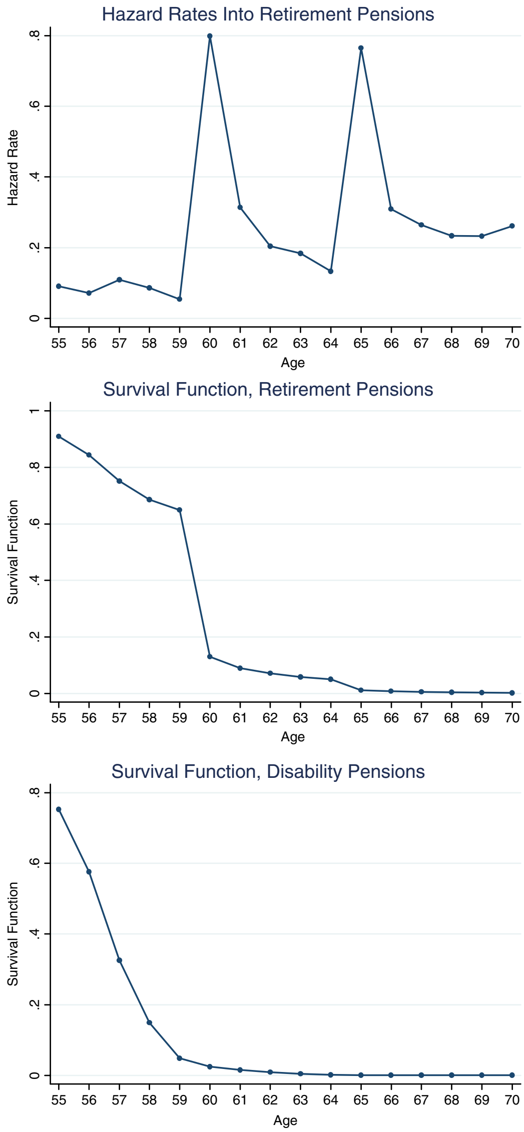

Notes: These figures are based on the sample of all individuals claiming retirement pensions after 1984 (394,934 individuals and 275,379 claimants). In the top plot, the retirement hazard is computed by defining the failure event as claiming a retirement pension (either an old-age pension or a disability 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.

1987–2003).[9] Appendix Table S1 presents a summary of each reform. Figure S1 presents benefits-versus-age profiles for different calendar years that illustrate the variation in pension benefits created by the pension reforms.

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 S1, 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; as 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. The 1996 reform introduced a bonus/malus system to discourage early retirement (before the statutory age) by penalizing early retirees with

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

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 2000 and October 2002. As nominal adjustments in later years were lower than inflation, real benefits declined between 1998 and 2002.

B. 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 (Table S2). We construct two key variables to capture incentives from the government-provided pensions. The two variables are SSW and the accrual (ACC). An individual's SSW 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

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

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 β=.93 captures the individual's discount factor.[10] In this definition, we also assume that the maximum age that individuals can live to is 100 years. Each individual's retirement pension is calculated based on the rules of the Austrian pension system and the individual's observed earnings history. While the SSW 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 $\mathrm{SSW}_{i,a}$ net of pension contributions from delaying retirement by one additional year. Thus we define the accrual for individual i at age a as

$$\mathrm{ACC}_{i,a} = \left(E_a\left(\mathrm{SSW}_{i,a+1}\right) - \mathrm{SSW}_{i,a}\right)/\mathrm{SSW}_{i,a}.$$

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 to roughly $1.06. 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 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 β=.93, median SSW ranges from about €260, 000 at age 55 to €315, 000 at age 65 reflecting that higher earners who have yet to retire at the later ages have higher SSW. The accrual is close to −10% at each age reflecting the loss in SSW 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 SSW 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 indicate that

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

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

TABLE 1
Summary Statistics by Age

Age		Annual Earnings	Annual Benefits	SSW	ACC
55, $N = 242{,}402$	Mean	39,711.99	20,959.57	258,450.10	−0.089
	Median	33,127.34	21,480.03	264,852.50	−0.090
	SD	25,620.50	4,910.24	62,148.53	0.009
56, $N = 197{,}959$	Mean	39,822.22	21,746.77	264,092.70	−0.091
	Median	33,323.61	22,384.38	271,606.60	−0.091
	SD	27,461.05	5,076.34	63,352.13	0.008
57, $N = 172{,}739$	Mean	39,841.66	22,597.94	269,996.20	−0.094
	Median	33,402.45	23,367.69	279,006.30	−0.093
	SD	28,903.82	5,224.69	64,103.07	0.009
58, $N = 144{,}321$	Mean	40,324.68	23,465.54	275,572.70	−0.094
	Median	33,715.05	24,540.40	287,978.70	−0.095
	SD	30,562.88	5,368.77	64,846.93	0.011
59, $N = 123{,}954$	Mean	37,798.21	24,342.59	281,022.50	−0.095
	Median	31,842.09	25,679.08	295,721.00	−0.097
	SD	33,052.84	5,554.38	66,215.05	0.012
60, $N = 108{,}183$	Mean	26,147.84	25,150.57	284,951.30	−0.102
	Median	18,919.41	26,579.45	299,823.80	−0.109
	SD	31,470.91	5,747.87	67,563.28	0.020
61, $N = 16{,}268$	Mean	47,094.05	26,803.42	300,095.00	−0.113
	Median	40,401.54	29,790.58	32,3255.80	−0.118
	SD	40,261.27	6,712.42	77,020.84	0.022
62, $N = 7{,}657$	Mean	53,294.76	26,837.35	292,728.30	−0.122
	Median	45,939.05	3,0263.92	319,571.90	−0.123
	SD	42,314.43	7,054.66	78,509.95	0.019
63, $N = 4{,}565$	Mean	52,857.83	26,715.18	284,888.70	−0.126
	Median	45,611.39	30,402.21	313,239.80	−0.125
	SD	42,908.46	7,254.17	78,547.99	0.019
64, $N = 2{,}793$	Mean	52,823.81	26,673.15	277,747.00	−0.095
	Median	45,096.35	30,297.56	309,250.10	−0.096
	SD	44,722.16	7,308.58	77,052.34	0.018
65, $N = 1{,}787$	Mean	40,839.73	27,492.06	279,203.10	−0.096
	Median	31,857.40	32,011.43	315,836.70	−0.096
	SD	38,886.03	7,735.31	79,230.61	0.017
			Assets		
Percentile	10	25	50	75	90
	0.00	4,930.6	24,884.33	76,296.25	160,007.4

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 50–54 years from the SHARE-Austria data.

households have accumulated financial assets roughly equivalent to one year's earnings.

III. EMPIRICAL ANALYSIS I: ESTIMATION OF ELASTICITIES

A. Proportional Hazards Specification and 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,

$$R_i(a) = \overline{R}(a)\exp\left\{\beta_{\mathrm{SSW}}\ln\left(\mathrm{SSW}_{i,a}\right) + \beta_{\mathrm{ACC}}\ln\left(1+\mathrm{ACC}_{i,a}\right) + \delta X_{i,a}\right\}.$$

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 $\overline{R}(a)$ 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, $\mathrm{SSW}_{i,a}$ is the expected present value of the individual's retirement pension if he were to retire at age a, and $\mathrm{ACC}_{i,a}$ is the individual's expected pension accrual (i.e., the change in $\mathrm{SSW}_{i,a}$ from delaying retirement 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 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.

B. 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 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

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

13. Figures S2 and S3 present graphical evidence highlighting this identifying variation.

TABLE 2
Hazard Model Estimates

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

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 of ages 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.

IV. 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 (Berkovec and Stern 1991; Lumsdaine, Stock, and Wise 1992; Rust and Phelan 1997; Stock and Wise 1990). The intuition behind the model is as follows. In each period, an employed individual must choose whether to retire or 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.

First consider 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

14. It 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 (2014) highlight the impacts of severance payments on individuals' retirement decisions in Austria.

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:

$$R_a\left(\Omega_a\right) = \max_{\{c_t^R\}_{t=a}^T} u\left(c_a^R\right) + \sum_{t=a+1}^{T} \beta^{t-a}\pi_{t|a}u\left(c_t^R\right)$$

$$\text{s.t.}\sum_{t=a}^{T}\left(1/\left[1+r\right]\right)^{t-a}\pi_{t|a}c_t^R = A_a - \kappa\mathbf{1}\left(a < a_{old}\right)$$

$$+\sum_{t=a}^{T}\left(1/\left[1+r\right]\right)^{t-a}\pi_{t|a}y_a^R + s\left(\tau^R\right).$$

The variables r, β, and $\pi_{t|a}$ denote the interest rate, the discount factor, and the probability of survival to age t, respectively, conditional on survival to age a. The function $u(.)$ captures utility over consumption with $u' > 0$ and $u'' < 0$. The maximization reflects that the individual chooses his consumption in each period. The term 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 this article. 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 $\tau^R < 10$, $(1/3)y^W$ if $\tau^R \in [10, 15)$, $(1/2)y^W$ if $\tau^R \in [15, 20)$, $(3/4)y^W$ 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

$$W_a\left(\Omega_a\right) = \max_{c_a^W} u\left(c_a^W\right) - v_a$$

$$+\beta\pi_{a+1|a}E_a\left[D_{a+1}\left(\Omega_{a+1}\right)\right]$$

$$\text{s.t. } c_a^W + \left(\left[A_{a+1}\right]/\left[1+r\right]\right) = y_a^W + A_a.$$

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 $D_a(\ldots)$ 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,

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

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 $u(.)$ 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

15. This 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.

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 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. As collective bargaining agreements effectively determine older workers' wages in Austria, we assume that earnings increase deterministically with age, $y^W_{a+1} = (1+g)\, y^W_a$. 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.

V. EMPIRICAL ANALYSIS II: STRUCTURAL ESTIMATION

A. 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 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 50–54. Finally, we also fix the discount factor $\beta = (1/[1+r]) = .93$ as it is difficult to distinguish empirically β from declining work disutility across ages. Thus, $\chi = \left(T, \overline{s}, g, r, \pi_{a|a-1}, \pi_{\text{sep}}, \beta\right)$. As we do not observe assets, we approximate the initial distribution of assets at age 54 using the Austrian 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 IV as follows. We assume constant relative risk aversion utility over consumption:

$$u(c) = \left(c^{1-\gamma} / \left[1-\gamma\right]\right), \gamma > 0.$$

We assume initial work disutility is drawn from an exponential distribution with mean $\widetilde{\eta} = x\eta$, where $\eta > 0$ and $x = u\left(\overline{c}\right) - u\left(\overline{rc}\right)$ is a scaling factor for the disutility of work based on income differences between work and retirement (we use $\overline{c} = 30,000$ and $\overline{r} = .55$ 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 $\alpha\widetilde{\eta}$ (i.e., $v_a = \alpha\widetilde{\eta}\,(a-54) + v_{54}$). 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 because 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 parameterized by θ. The goal is to find the value of θ which gives the best "fit" of the model, that is, by minimizing the (weighted) distance between the observed and

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

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, that is, 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 II (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 55–65. 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 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 $\widehat{\beta}_N \xrightarrow{p} \beta\left(\theta_0\right)$ 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 β(θ). 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:

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

for some weighting matrix W. Gourieroux, Monfort, and Renault (1993) show that the II is a consistent estimator of θ for $N \to \infty$. 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.

B. 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 as 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

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 ($\beta = .88$)
(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	102,018.94	116,292.34	134,748.04	68,491.32
(B) Proportional Hazard Coefficients:				
Coefficients with All Ages				
β_{SSW}	.4069	.3473	.5855	.4575
β_{ACC}	−3.0064	−3.1644	−2.5491	−2.8995
$\beta_{SSW}/{-\beta_{ACC}}$	.1353	.1098	.2297	.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 the same sample used to estimate the proportional hazard specifications in Table 3. The baseline specification is based on a discount factor of $\beta = .93$, 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 γ (Using Baseline Estimate $\gamma = 0.7064$)

	0.50(γ)	0.75(γ)	1.00(γ)	1.50(γ)	2.00(γ)	3.00(γ)
β_{SSW}	0.5292	0.5095	0.4863	0.4236	0.4296	0.4786
β_{ACC}	−4.3906	−3.5462	−3.0179	−2.3442	−1.647	−0.9957
$\beta_{SSW}/{-\beta_{ACC}}$	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 text for more details.

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 ($\beta_{SSW}/{-\beta_{ACC}}$), 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. To put this cost in perspective, the SSW (expected presented discounted value) from an annual pension benefit of about €8,300 would be roughly €100,000 (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% (=8, 300/21, 000) 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 as 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 ($\beta = .88$ instead of .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

Structural Estimation, Predicted versus Actual Hazard Rates.

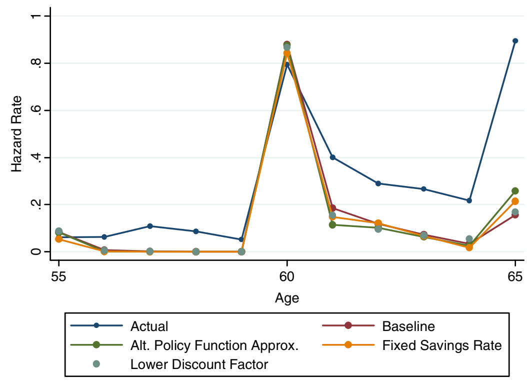

Notes: The baseline model uses a discount factor of β=0.93 and a fixed savings rate of 0.10. The lower discount factor is β = .88 and the lower savings rate is .05. Please see the text for more details

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 standout. First, the model overpredicts retirement at the earliest age, age 55. Intuitively, the model predicts that any individuals that draw a high disutility of work should retire immediately as 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. As 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 (e.g., see Blau (2008); French (2005); Hubbard Skinner and Zeldes (1995); 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)).

VI. 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., $y_a^R \to (1 - .20)\, y_a^R$). The second reform is a change in benefits by 3% per year from age 65 (i.e., $y_a^R \to [1 - .03\,(65 - a)]\, y_a^R$); 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 as 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 as 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 as 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

FIGURE 3

(A) Labor Supply Responses to Policy Simulations and (B) Welfare Consequences of Policy Simulations.

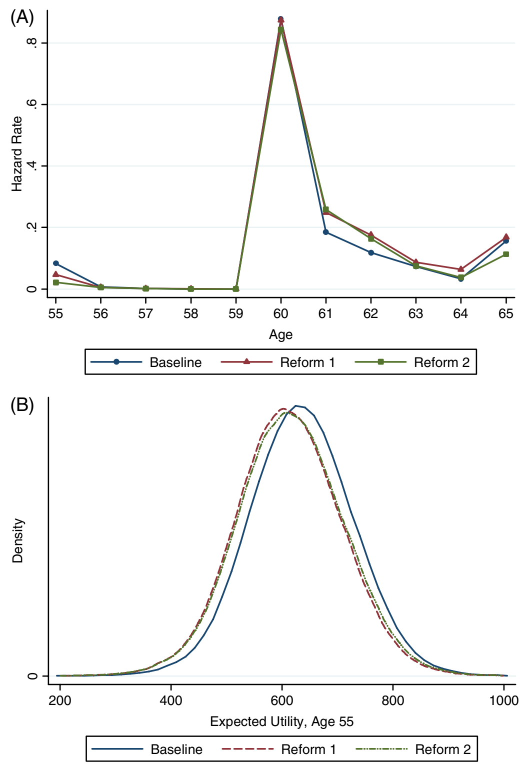

Notes: The baseline is computed using the parameter estimates reported in Table 4. Reform 1 decreases benefits across all ages by 20%. Reform 2 changes benefits at a give age based on 3%*(65-age).

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 versus 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 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.

VII. 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 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. "Distinguishing Income from Substitution Effects in Disability Insurance." *American Economic Review Papers and Proceedings*, 97(2), 2007, 119–24.

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

Blau, D. "Retirement and Consumption in a Life Cycle Model." *Journal of Labor Economics*, 26(1), 2008, 35–71.

Boskin, M. J. "Social Security and Retirement Decisions." *Economic Inquiry*, 15, 1977, 1–25.

Brown, K. "The Link Between Pensions and Retirement Timing: Lessons from California Teachers." *Journal of Public Economics*, 98, 2013, 1–14.

Chetty, R. "A New Method of Estimating Risk Aversion." *American Economic Review*, 96(5), 2006, 1821–34.

———. "Moral Hazard Versus Liquidity and Optimal Unemployment Insurance." *Journal of Political Economy*, 116(2), 2008, 173–234.

Coile, C., and J. Gruber. "Future Social Security Entitlements and the Retirement Decision." *Review of Economics and Statistics*, 89(2), 2007, 234–46.

French, E. "The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour." *Review of Economic Studies*, 2(72), 2005, 395–427.

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

Gourieroux, C., A. Monfort, and E. Renault. "Indirect Inference." *Journal of Applied Econometrics*, 8(S1), 1993, S85–S118.

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

Heckman, J. J., and R. Robb. "Alternative Methods for Evaluating the Impact of Interventions: An Overview." *Journal of Econometrics*, 30(1–2), 1985, 239–67.

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

Hubbard, R. G., J. Skinner, and S. Zeldes. "Precautionary Saving and Social Insurance." *Journal of Political Economy*, 103(2), 1995, 360–99.

van der Klaauw, W., and K. I. Wolpin. "Social Security and the Retirement and Savings Behavior of Low Income Households." *Journal of Econometrics*, 96(1–2), 2008, 21–42.

Koman, R., U. Schuh, and A. Weber. "The Austrian Severance Pay Reform: Toward a Funded Pension Pillar." *Empirica*, 32(3), 2005, 255–74.

Krueger, A., and J. Pischke. "The Effect of Social Security on Labor Supply: A Cohort Analysis of the Notch Generation." *Journal of Labor Economics*, 10(4), 1992, 412–37.

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

Lumsdaine, R., J. Stock, and D. Wise. "Three Models of Retirement: Computational Complexity Versus Predictive Validity," in *Topics in the Economics of Aging*, edited by D. Wise. Chicago: University of Chicago Press, 1992.

Manoli, D., and A. Weber. "Nonparametric Evidence on the Effects of Financial Incentives on Retirement Decisions." Working Paper, University of Texas at Austin, 2014.

Marek, E. Pensionsberechnung Leicht Gemacht. Vienna: Weiss Verlag, 1985, 1987–2003.

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

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

Organisation for Economic Co-operation and Development. *Pensions at a Glance: Public Policy Across OECD Countries*. Paris: OECD, 2007.

Pencavel, J. "The Response of Employees to Severance Incentives: The University of California's Faculty, 1991–94." *Journal of Human Resources*, 36(1), 2001, 58–84.

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

Staubli, S. "The Impact of Stricter Criteria for Disability Insurance on Labor Force Participation." *Journal of Public Economics*, 95(9–10), 2011, 1223–35.

Stock, J., and D. Wise. "Pensions, the Option Value of Work, and Retirement." *Econometrica*, 58(5), 1990, 1151–80.

SUPPORTING INFORMATION

Additional Supporting Information may be found in the online version of this article:

Figure S1. Changes in Benefits from Pension Reforms.

Figure S2. Graphical Analysis of Wealth and Price Effects, Age 55.

Figure S3. Graphical Analysis of Residuals, Year-Age Cell Means.

Table S1. Summary of Austrian Pension Reforms—1984–2003.

Table S2. Sample Restrictions, Initial Sample (Males, Birth Cohorts ≥ 1948): 2403454.