Early Retirement as Insurance against Automation: Evidence from Italian Couples*

Claudio Costanzo

September 16, 2026

Abstract

Occupational exposure to industrial robots pushed older Italian men out of the labor force in the late 1990s, but only through the pension door: where the seniority pension was available, a one-standard-deviation increase in exposure raised the annual flow from work into retirement by 6.0 percentage points, while unemployment, inactivity, and pay never moved. The pension replaced 82 percent of a husband's income and unemployment benefit 14 percent, and neither spouse adjusted labor supply when the other stopped working. A life-cycle model of the couple, estimated on employee households, explains both facts with one taste for joint leisure set against those two replacement rates: the wife withdraws when he retires and barely moves when he loses his job, as in the data. Because exposed men used the early exit most, raising the eligibility age by five years, as the Amato and Dini reforms did, removes 45 percentage points of retirement from above-median-exposure men against 33 from the rest, and costs exposed households two to four times more than unexposed ones. Public transfers fall, and between one half and one fifth of the removed retirement returns to work. Who pays for such a reform therefore depends on which workers held the early exit. Two alternative designs spread the burden differently: allowing men to retire early with a reduced pension costs households less while raising nearly as much revenue, and exempting the most exposed lowers the cost further while giving up more than half of the revenue.

(JEL D15, H55, J14, J22, J26, O33)

*ECARES, Université libre de Bruxelles. Contact: claudio.costanzo@ulb.be

I. Introduction

Industrial robots displace some workers and reshape the jobs of others. Where a displaced worker goes next depends on the alternatives open to him, and near the end of a working life the leading alternative is to leave the labor force altogether. Acemoglu and Restrepo (2020) and Graetz and Michaels (2018) measure the aggregate effects of robot adoption on employment and productivity, and Dauth et al. (2021) show that in Germany incumbent manufacturing workers are insulated by within-firm reallocation while the young bear the adjustment through entry, estimates that, by construction, come from workers still far from the end of their careers. Work at the other end of the age distribution documents that automation raises the rate at which older workers leave their jobs. Bessen et al. (2025) show in Dutch administrative data that automating firms shed incumbents, that the five-year earnings loss is about nine percent of a year's pay, and that the loss falls disproportionately on older workers; Yashiro et al. (2022) show for Finland that technology-exposed older workers exit faster where an extended unemployment benefit bridges them to a pension; and Casas and Román (2023) find in European survey data covering 26 countries that workers in occupations at higher automation risk retire early more often. None of the three asks what the exit costs the household, and none models the choice between exits that insure the worker differently. For an older worker whose occupation is being automated, the relevant choice is between his job and an exit route, and which routes are open to him is a matter of pension law. The paper's organizing observation is that the most generous exit was used disproportionately by the men whose occupations were being automated, so that the incidence of closing it depends on how exposure and eligibility are distributed together.

This paper studies that choice in Italy between 1994 and 2001 using data from the European Community Household Panel. Over those years the robot stock roughly doubled, and two pension reforms, Amato in 1992 and Dini in 1995, began to close the most generous of the exit routes. Three features make the setting informative. Italy had by far the most permissive early-retirement system in the OECD, built around a *pensione di anzianità* that turned on thirty-five years of contributions alone, with no age requirement, so eligibility arrived at different ages for observationally similar men. Italian women's employment fell steeply over the ages at which men retire, so for most couples the second earner was out of the labor force, and her margin had to be opened before it could adjust. Finally, the panel follows both spouses, with activity status, personal income by source, and occupation, for up to eight consecutive years.

I document five facts and then develop a model to account for them. First, exposure to automation raises the flow of older men into retirement and moves neither unemployment nor inactivity, a pattern the design identifies more sharply than its size. The response is present among the private employees the seniority pension covered and absent among the self-employed, who contributed to separate funds. Second, exposure does not change what exposed men are paid, so it does not act through the wage. Third, the two exits differ in how well they are insured, the pension replacing 82 percent of the husband's income and unemployment benefit 14 percent. Fourth, neither spouse increases labor supply when the other stops working: a household losing 86 percent of the husband's earnings to unemployment does not send the wife to work. Fifth, couples retire together, and retirement is the only exit they share: his job loss triggers no joint move. A question underneath the first fact is whether the exit is chosen or imposed. The routes men take, the reasons they give, and the absence of shedding at ineligible ages point to a chosen exit, and the small involuntary residual reappears in the counterfactual as an imposed bound.

The absence of an added-worker effect is not new (Lundberg, 1985; Cullen and Gruber, 2000; Stephens, 2002), but its size here is hard to reconcile with a frictionless household. The model explains it with one taste and two replacement rates, the share of his working income that each exit pays him. The couple values shared free time, and that pull is the same whether the husband is retired or unemployed. The insurance motive pushes the other way, and its strength depends on the replacement rate: the pension pays most of his income, unemployment benefit almost none of it. When he retires the insurance motive is therefore weak and the taste for joint leisure wins, so the wife withdraws; when he loses his job the two forces roughly cancel, so she barely moves. Nothing in the preferences distinguishes his two exits; the replacement rates, measured directly in the income data, do all the work. The data show the same ordering, and the difference between her two responses is not statistically significant.

Rationing of the wife's entry margin is the obvious explanation, but it accounts for the level of her entry and almost none of its response. Making offers as common in her fifties as at thirty would more than triple how often she enters, and would leave her response to his job loss essentially unchanged.

The counterfactual then closes the insured exit and asks who bears the closure. Delaying eligibility by five years removes 45 percentage points of retirement among above-median-exposure men at ages 55–59, against 33 among the rest, and costs exposed households between two and four times more than unexposed ones. That gap is the paper's most robust quantity: a modeling choice described below moves every other number in the exercise and leaves this one alone. Public transfers fall, and on a common present-value scale the government gains less than the household loses. Automation changes who pays far more than how much is saved: in a version of the model where every occupation has the same exposure, the reform saves and costs almost the same in total, but the burden is no longer concentrated on the exposed. Two alternative designs spread the burden differently. Letting men retire early with a reduced pension raises nearly as much revenue at a lower cost to households, and exempting the most exposed lowers the cost further while giving up more than half of the revenue. Section VI prices both.

The destination of the removed retirement is the least well determined part of the exercise. Between one half and one fifth returns to work, and the rest splits between unemployment and inactivity. The two exposure groups land differently: the unexposed mostly return to work, and the exposed divert toward inactivity. The width of that range turns on one question, whether entry into long-term inactivity is a shock or a choice, and the estimation sample is close to silent on it: it sees men who leave work and come back, and hardly any who leave for good. Section VI reports every quantity under both readings, with the evidence that bears on them. Under either reading the household absorbs the reform on the husband's margin: the wife supplies a tenth of the adjustment, and opening her entry margin shrinks her response further.

The paper connects three literatures. The structural retirement literature (Stock and Wise, 1990; Rust and Phelan, 1997; French, 2005; French and Jones, 2011) models the exit decision of a single worker facing a pension system. This paper adds a spouse whose margin is rationed and an occupational shock that acts on the cost of continuing and leaves the wage alone. The literature on couples' retirement (Blau, 1998; Gustman and Steinmeier, 2000; Coile, 2004; Casanova, 2010) documents joint retirement and models it through complementarities in leisure. Here the taste for joint leisure is the same whether the husband is retired or unemployed, and the wife's different responses to the two events follow from the budget. The third is the literature on family labor supply as insurance (Blundell et al., 2016; Halla et al., 2020; Fadlon and Nielsen, 2021), which estimates the response of one spouse to the other's shock. The model here reproduces the near-zero response and attributes it to the interaction of a taste parameter with two measured replacement rates. On the Italian institutions the paper draws on Brugiavini (1999), Brugiavini and Peracchi (2004), Bottazzi et al. (2006), and Battistin et al. (2009).

Three older literatures anticipate the mechanism and none measures it against a pension system with two exits. Bartel and Sicherman (1993) show that a worker in an industry whose rate of technical change accelerates unexpectedly retires sooner, because the retraining the acceleration requires is a bad investment at the end of a career, while a worker in an industry with a high but anticipated rate retires later; Friedberg (2003) and Aubert et al. (2006) document the same asymmetry between older and younger workers in computer use and in firm-level technology adoption; and Ahituv and Zeira (2011) formalize it as human-capital obsolescence that raises the retirement hazard. The direction is not universal: Blanas et al. (2019) find across ten high-income countries that software and robots raised relative demand for older workers, mainly in services, while lowering it for the young and the medium-skilled. The obsolescence channel explains why an exposed man wants to stop. It is silent on which exit he takes, and in Italy in these years the two available exits differed by a factor of nearly six in what they replaced.

The Italian evidence on robots is recent, studies local labor markets, occupations, and individual workers, and does not reach the household. Dottori (2021) finds no statistically significant negative effect of robot diffusion on total local employment in Italy between the early 1990s and 2016, nor on the employment and wages of incumbent manufacturing workers; Caselli et al. (2025) decompose local employment changes over 2011–2018 by the activities occupations perform and recover effects that differ across occupations; Cirillo et al. (2021) show on Italian task data that digital-intensive occupations gain employment, except where their tasks are routine. Anelli et al. (2021) price the political consequence of individual robot exposure in Western Europe. A growing literature follows automation into the family. Anelli et al. (2024) link robot exposure to marriage and fertility across US commuting zones; Matysiak et al. (2023) find that robot adoption lowered regional fertility in European countries, an effect concentrated at young ages and significant in Italy; Lu et al. (2025) report comparable fertility effects for China; Giuntella et al. (2025) show, also in China, that exposed workers leave the labor force earlier while households adjust through borrowing and investment in children's education; and Costanzo (2025) studies how automation shifts the optimal timing of fertility. This paper differs in following the older worker and his spouse through an automation-driven exit into retirement, and in placing that exit inside the Italian pension system, on panel data that observe both spouses over the years the seniority pension was still open.

Section II describes the setting and the data. Section III presents the five facts. Section IV sets out the model; Section V estimates it and decomposes the household's response. Section VI reports the pension counterfactual. Section VII concludes.

II. Setting and Data

A. Data Sources

The paper combines four sources. The household data are the Italian waves of the European Community Household Panel, a harmonized survey run by Eurostat between 1994 and 2001 in which the same households are interviewed once a year for up to eight consecutive waves. The Italian component is administered by ISTAT. It covers 7,115 households and 17,729 individuals in the first wave and, after attrition, 5,606 and 13,392 in the eighth. It records, for every adult member separately, self-reported main activity in the week of interview, usual weekly hours, occupation, educational attainment, region of residence, and annual income in the calendar year preceding the interview, broken down by source. Two features make it well suited to this question. Where many surveys follow a designated head, it follows both spouses, so cross-spouse labor supply responses are observed within couple. It also separates income into personal components, so a pension can be attributed to the person who receives it and need not be inferred from a household residual. The panel's construction, sampling design and known attrition problems are described by Peracchi (2002).

The occupational exposure measure is the task-based robot-exposure index of Webb (2020), computed for 341 US Census occupations. The aggregate diffusion measure uses operational robot stocks for Italy by industry from the International Federation of Robotics, combined with Italian sectoral employment. Mapping the first of these onto the survey requires an occupational crosswalk, and mapping the second onto regions requires a sectoral one. Both are described in Section II.D.

B. Sample, Activity States, and Income

Two samples are used, and the distinction between them is institutional. The couple panel keeps couples in which both spouses are aged 25–69, the region of residence is observed, and the husband's first-observed occupation is recorded. It contains 28,290 couple-waves on 4,936 couples. The estimation sample restricts that panel further to couples in which the husband is a paid employee in the first wave he is observed working, and contains 19,611 couple-waves on 3,375 couples observed for 5.8 waves on average. Both are described in Table 1, which reports the estimation sample.

The reason for the restriction is that the institution the paper is about applied to employees. The *pensione di anzianità* was a scheme of the private-employee fund. The self-employed contributed to the separate artisan and trader funds, and the same age threshold does not describe their seniority conditions. Fact 1 below shows that the retirement response to exposure is present among employees and absent among the self-employed, and Section III reports the split. The model of Section IV is therefore estimated on employee households, and its counterfactual speaks to them. The reduced-form facts are reported on whichever of the two samples the object belongs to, and each table names its own.

Table 1. Descriptive statistics, estimation sample

	Husband	Wife
Panel A. Demographics		
Age, mean	45.1	41.9
Age, standard deviation	9.8	9.5
Lower secondary education or less	0.5150	0.5301
Upper secondary	0.3704	0.3784
Tertiary	0.1093	0.0865
Panel B. Activity status		
Working / employed	0.8635	0.4379
Unemployed	0.0303	
Inactive / non-employed	0.0183	0.5182
Retired	0.0879	0.0440
Weekly hours, if working	40.7	34.8
Panel C. Annual net personal income, EUR		
All person-waves	14,168	5,245
Conditional on working	14,746	10,323
Panel D. Household and sample		
Spousal age gap, mean (s.d.)	3.21 (3.66)	
Any child under 16 present	0.5667	
Household net income, EUR	41,097	
Region: North / Centre / South and Islands	0.3727 / 0.1893 / 0.4379	
Couples	3,375	
Couple-waves	19,611	
Waves observed per couple, mean (median)	5.8 (7)	

Notes: Italian waves of the European Community Household Panel, 1994–2001. The estimation sample is couples in which the husband is an employee in the first wave he is observed working. Panels A–C are means over couple-waves except for education, which is measured at the couple's first observation. The wife's non-employment pools her unemployment with her other inactivity. Income is annual net personal income converted from thousands of lire at 1,936.27 lire to the euro left in nominal terms; household net income is winsorized at the 95th percentile.

Activity status is self-reported main activity, and the classification is not self-evident. The survey distinguishes paid employment from self-employment, and self-employment is a working state: men who report it are recorded as employed on every other activity field in the panel. Its position in the response list invites treating it as non-employment, which would place men who work full time and draw earnings from work into the non-employment states. Paid employment, self-employment and other work are therefore grouped into a single working state, and unemployment, retirement, and other inactivity are kept separate. The last of these is kept apart from unemployment because the counterfactual asks whether delaying pension eligibility diverts men into disability. The classification is validated against the panel's own employment, participation and hours variables, which are constructed from different underlying questions; Appendix A reports the validation.

Income is measured on personal components. The husband's net personal income when working averages 15,940 EUR per year on the state-stable households from which the budget is built, and is the numeraire throughout. Replacement rates are computed on those households: the husband is in the same state at $t-1$ and t, so the annual income measure, which refers to the previous calendar year, matches the state actually occupied. The alternative is a residual measure, which attributes the whole change in household income net of the couple's own earnings to the husband's exit. The two diverge sharply: the residual gives a replacement rate on retirement of 1.0431 against 0.8156 on the personal-income measure. The gap is the earnings of co-resident adult children, which rise over exactly the ages at which Italian men retire, and which a household residual cannot separate from a pension.

C. Descriptive Statistics

Table 1 describes the estimation sample. Husbands average 45.1 years and wives 41.9, with a mean spousal age gap of 3.21 years and a standard deviation of 3.66. That dispersion bears on the model, because a couple's joint retirement problem depends on how far apart the two spouses are from their respective eligibility ages. Just over half of each spouse has lower secondary education or less, eleven percent of husbands and nine percent of wives have a degree, and 56.7 percent of couple-waves have a child under 16 in the household. The sample is 37.3 percent northern, 18.9 percent central and 43.8 percent southern.

The labor market rows carry the two features the model is built around. Husbands are at work in 86.4 percent of couple-waves and retired in 8.8; wives are employed in 43.8 percent and non-employed in 51.8, with retirement at 4.4. Conditional on working, husbands report 40.7 hours a week and wives 34.8. The income rows show how large a gap the second earner could close: a husband earns 14,746 EUR when working and an employed wife 10,323, so a wife who moves from non-employment into work replaces about seven-tenths of her husband's earnings.

D. Measuring Occupational Exposure to Automation

The treatment is the husband's exposure to robots through the tasks his occupation requires, measured with the index of Webb (2020). The index is built by parsing the task descriptions in the US Department of Labor's O*NET into verb–noun pairs—the action a job involves and the object it acts on, such as "weld frame" or "diagnose patient"—parsing the text of technology patents the same way, and scoring each occupation by how far the capabilities the patents describe cover its tasks. Webb (2020) applies this procedure separately to three bodies of patents, yielding one index each for software, industrial robots, and artificial intelligence. This paper uses the industrial-robot index, which scores an occupation by the share of its tasks that fall within the capabilities claimed in robot patents and expresses it as a percentile rank from 0 to 100. The index is defined over 341 harmonized 1990 US Census occupations, weighted by 2010 US employment. It is built from patent text and task content, with no labor-market outcome anywhere in its construction, so it cannot inherit the employment or wage variation it is used to explain. A robot patent describes a machine's physical capabilities, with no reference to any firm's decision to install one, so the score is a property of the occupation itself, the same in every local labor market.

Bringing the index onto the survey takes two steps, because it is defined over US Census occupations while the panel records the international classification at a coarser resolution. First, no direct map between the two exists, so the construction composes three public crosswalks: from the international classification to the 2000 US Standard Occupational Classification, from that to the 2000 US Census classification, and from that to the harmonized 1990 Census codes on which the index is defined, with the digit-dropping and residual-category conventions each map requires. Second, the index sits at a finer resolution than the survey carries, so within each survey occupation cell the Census-occupation scores that map into it are averaged with 2010 US employment weights. This produces a robot-exposure score for each of the seventeen two-digit occupation groups the panel records and, as a fallback, for each of the nine one-digit major groups. Each husband is assigned the score of the two-digit occupation he holds in the first wave he is observed working, or of his one-digit group where the two-digit code is missing. The survey's missing codes are screened out before that first observation is taken. On the estimation sample 3,332 of the 3,375 couples take a two-digit score and 43 the one-digit fallback, so the treatment takes 18 distinct values, reported in Table 2, and it is standardized to mean zero and unit variance across the estimation sample.

In symbols, survey occupation o receives the score

$$w_o = \frac{\sum_{k \in \mathcal{K}(o)} \ell_k \, x_k}{\sum_{k \in \mathcal{K}(o)} \ell_k}, \tag{1}$$

where $\mathcal{K}(o)$ is the set of harmonized 1990 Census occupations that the composed crosswalk maps into o, x_k is the robot-exposure percentile of Census occupation k, and ℓ_k is its 2010 US employment. The treatment of couple c is

$$z_c = \frac{w_{o(c)} - \bar{w}}{s_w}, \tag{2}$$

where $o(c)$ is the occupation assigned above and $\bar{w}$ and s_w are the mean and standard deviation of the assigned scores over the estimation sample.

Table 2. Occupational robot-exposure scores

Occupation group	Webb score	Standardized score	Husbands Couples	Share
Teaching professionals	3.0	−2.13	100	0.030
Other professionals	19.7	−1.42	33	0.010
Legislators, senior officials and managers	20.7	−1.37	110	0.033
Salespersons and demonstrators	26.7	−1.12	116	0.034
Science, engineering and health professionals	28.8	−1.03	88	0.026
Clerks	29.7	−0.99	710	0.210
Elementary sales and service occupations	41.6	−0.48	193	0.057
Teaching and other associate professionals	47.3	−0.24	216	0.064
Science, engineering and health associate professionals	58.1	+0.22	150	0.044
Personal and protective service workers	62.0	+0.39	162	0.048
Metal, machinery and precision trades	66.6	+0.59	411	0.122
Agricultural labourers	70.9	+0.77	85	0.025
Extraction, building and other craft trades	71.4	+0.79	393	0.116
Skilled agricultural and fishery workers	74.0	+0.91	92	0.027
Machine operators and assemblers	76.6	+1.02	70	0.021
Plant operators and drivers	86.7	+1.45	227	0.067
Labourers in mining, construction and manufacturing	90.6	+1.62	176	0.052

Notes: The Webb score is the employment-weighted mean of the occupation-level robot-exposure measure of Webb (2020), aggregated from US Census occupations to these groups through the crosswalk described in Section II.D; it runs from 0 to 100 and is the percentile rank of an occupation's task overlap with the text of robot patents. The standardized score subtracts the mean and divides by the standard deviation over the estimation sample, and is the treatment variable throughout; all counts and shares in this table are on that sample. The 3,332 couples listed here are those whose husband's first-observed occupation carries a detailed code, of the 3,375 in the estimation sample; the remaining 43 are assigned the score of their one-digit major group, so the treatment takes 18 distinct values in all. Survey missing codes on the detailed occupation field are screened before the first observation is taken.

The ordering in Table 2 is the one the measure is meant to produce. Labourers in manufacturing and construction, plant operators and drivers, and machine operators sit at the top; teaching professionals, other professionals and managers at the bottom. The raw score runs from 3.0 to 90.6 with a mean of 52.8 and a standard deviation of 23.3. The distribution is not concentrated: the tenth percentile of the standardized score is −1.12 and the ninetieth is +1.45.

The score is taken at first observation, so it is predetermined with respect to everything that happens afterwards, but it is not a perfect measure of the occupation a man spends his career in. Of the 4,249 husbands with two or more usable occupation reports, half report more than one cell, a mean of 1.61 cells over 5.5 waves. For 342 couples of the couple panel the first non-null entry in the detailed occupation field is a survey missing code. All of them report a valid detailed occupation later, at a median lag of two waves, which is why the codes are screened before the first observation is taken. Rebuilding the score on the husband's modal occupation over all his valid waves instead gives a coefficient of +0.0544 against +0.0599, an attenuation of about a tenth, which is the direction and roughly the size that classical measurement error in a treatment with this many cells would produce; Section III reports the three constructions side by side.

It is worth being concrete about which couples identify the effect. Because exposure is fixed within a couple, the design learns only from couples observed both before and after the husband turns 55. There are 263 of them among the 1,252 in the risk set, contributing 1,472 of the 4,337 couple-waves and 81 of the 432 retirement events, and they carry essentially all of the identifying variation. Born between 1941 and 1946, they are the cohorts the Amato and Dini phase-in reached, and the ones the counterfactual of Section VI is about. Estimated on those couples alone the coefficient is +0.0480 (0.0163), against +0.0599 on the full risk set.

Because the score takes only eighteen distinct values, treatment varies across occupation cells and across couples, and Section III reports standard errors at both levels.

One concern is reverse causality: Acemoglu and Restrepo (2022) show that an ageing workforce itself induces firms to automate, in which case exposure and the age structure are jointly determined and any relationship between them is uninterpretable. Their variation is across countries and industries and over decades. The variation used here is across occupations, inside one country, inside eight years, and it enters the estimating equation only through the interaction with an institutional age threshold, with the couple's own exposure absorbed by a couple effect and national adoption absorbed by a year effect. The exposure score is a property of an occupation's task content, taken from patent text at the man's first observation, so it cannot move in response to the ageing of the workforce it is used to explain. The design cannot rule out that firms employing exposed occupations were adopting robots because their own workforces were old; that channel would push in the same direction as the estimate.

The second measure is aggregate, and its role is to date the diffusion. It is a shift-share index of regional robot exposure in the manner of Acemoglu and Restrepo (2020), built from the operational stock of industrial robots in Italy that the International Federation of Robotics reports by industry in its *World Robotics* series. For each of the eleven survey regions and each year,

$$R_{rt} \;=\; 1{,}000 \times \sum_{s} \frac{\ell_{rs}^{1991}}{\ell_{r}^{1991}} \cdot \frac{B_{st}}{\ell_{s}^{1991}}, \tag{3}$$

where B_{st} is the IFR operational robot stock of sector s in year t, ℓ_{rs}^{1991} is region r's employment in sector s in the 1991 base year, ℓ_{r}^{1991} its total, and ℓ_{s}^{1991} the sector's national employment, so that the index counts robots per thousand base-year workers; the national index R_t averages the regional ones. The shift, B_{st}/ℓ_{s}^{1991}, is common to every region, and the only cross-regional variation is the fixed 1991 industry mix, so a region is more exposed when its base-year employment sat in sectors that later robotized. The national index runs from 1.77 robots per thousand workers in 1994 to 3.62 in 2001 and correlates 0.9932 with the calendar year. That near-collinearity is why the paper does not interact exposure with the robot stock in its main specification: such a specification is a difference-in-differences in trends and cannot be separated from any other secular change over the same eight years. Identification comes instead from the interaction of predetermined occupational exposure with the crossing of an institutional age, within couple and within year. The aggregate index serves to date the period as one in which robot adoption in Italy roughly doubled and to supply the within-couple dose-response design used where a single flow is traced across ages.

E. The Italian Pension System

Two routes out of work existed for older Italian men. The old-age pension (*pensione di vecchiaia*) required reaching a statutory age, 60 for men before Amato. The seniority pension (*pensione di anzianità*) required thirty-five years of contributions and no age at all, so a man who entered the labor force at 18 could claim at 53 and a man with an interrupted career might never qualify. Replacement rates on final salary were high, and there was no actuarial adjustment for claiming early.

The Amato reform of 1992 raised the old-age age to 65 for men on a phased schedule and lengthened the minimum contribution requirement. The Dini reform of 1995 replaced the seniority rule with a combination of thirty-five contribution years and a minimum age, and moved new entrants onto a notional defined-contribution formula. Both reforms moved the date at which a given man could stop working, which is how the counterfactual in Section VI represents them.

The contribution rule has an implication for what the exposure coefficient can be. A man who left school at fourteen reached thirty-five contribution years five years before a man who left at nineteen and ten years before a graduate, and the occupations at the top of the exposure distribution are the ones such men enter: mean standardized exposure is $+0.30$ among husbands with lower-secondary schooling or less and -1.40 among the tertiary-educated, a correlation of 0.52 with the low-schooling indicator. Eligibility and exposure therefore arrive together, and a specification with only an age threshold cannot tell them apart; the survey carries no contribution histories with which to separate them directly, so Section III allows the threshold to differ by schooling and reports what that costs the estimate.

The age threshold used below is a reduced-form stand-in for an institution that did not run on birthdays, and the survey allows that stand-in to be checked. Every respondent is asked, in every wave, how old he was when he began his working life. The answer is available for 96.9 percent of the men in the risk set and is identical across waves for 97.1 percent of respondents, so potential contribution years can be measured directly. Among men of the same age in the same calendar year they have a standard deviation of 6.57 years. Applying the statutory schedule to them yields a predicted eligibility that behaves as the institution says it should. Men predicted eligible retire six to seven percentage points per year more often than ineligible men in the same narrow age band, and with couple and year effects the eligibility coefficient is $+0.0368$ (0.0178 clustered on the occupation cell). Tightening is visible inside the observation window: among men aged 53 to 56 the predicted-eligible share falls from 0.60 in 1995 to 0.33 in 2001. Eligibility in the mid-fifties is therefore an institution, visible in the data themselves.

III. Five Facts

The five facts fall into two groups with different regressors. The facts about automation use the two constructed measures of Section II.D: the husband's couple-fixed occupational exposure z_c, standardized to mean zero and unit variance, and the year-varying national robot stock R_t. These enter through two designs. The primary one is a threshold difference-in-differences that interacts z_c with an indicator for the husband's having reached the institutional retirement age, absorbing the couple and the year with fixed effects (equation (4) below, where the notation is defined). Where a single flow must be traced across ages and the couple cannot be held fixed, exposure enters instead through a dose-response term $z_c \times R_t$ under couple and year effects. The facts about the household do not use exposure: they regress each spouse's labor supply on the other's activity state within couple, again with couple and year effects. The dependent variable changes from fact to fact, from a transition hazard to log earnings to a spouse's employment, and each is defined where it is used.

FACT 1: *Exposure to automation raises the flow of older men into retirement, and does not move unemployment or inactivity.*

Table 3 reports the annual hazards out of work for husbands aged 50–69. The retirement hazard rises from 0.0436 at ages 50–54 to 0.4537 at 65–69, and pooled over the band it is 0.0996. The two competing exits are flat but not negligible: job loss averages 0.0168 and never exceeds 0.0194 in any band, and the flow into inactivity averages 0.0173, which is if anything the larger of the two and overtakes job loss from age 60. At 50–54 the retirement hazard is 0.0436 against 0.0171 into unemployment; by 60–64 it is 0.1918 against 0.0136. Retirement is distinguished by the slope of its age profile. The difference in level is small.

TABLE 3. ANNUAL TRANSITION HAZARDS OUT OF WORK, HUSBANDS AGED 50–69

		Hazard into		
Husband's age	Working at $t-1$	Retirement	Unemployment	Inactivity
50–54	2,227	0.0436	0.0171	0.0135
55–59	1,340	0.1187	0.0194	0.0172
60–64	662	0.1918	0.0136	0.0302
65–69	108	0.4537	0.0000	0.0185
50–69	4,337	0.0996	0.0168	0.0173

Notes: ECHP Italy, 1994–2001, estimation sample. The second column counts the couple-waves in which the husband was working at the previous wave and is observed in the current one; each hazard is the share of those couple-waves in which he is in the named state at t. Activity status is the respondent's self-reported main activity in the week of interview. The pooled row is the risk-set-weighted average of the four bands for retirement and the pooled hazard for the two other destinations. Re-employment out of unemployment is 0.3871 and out of inactivity 0.3154.

Exit at these ages is overwhelmingly a direct move from work into retirement. Table 4 follows husbands over 50 of the couple panel observed in three consecutive waves. In 5,367 couple-waves the man is working at both preceding waves, and 422 of these end in retirement. The route through unemployment carries almost no mass, with 100 couple-waves at risk and 7 retirements in the entire panel. Of the 692 men over 50 ever observed retired, 21 were ever observed unemployed at an earlier wave. The exposure regression cannot be estimated on the indirect routes: 100 and 63 observations spread over almost as many couples leave a couple fixed effect nothing to work with, so the layoff reading is instead ruled out below by other means.

Table 4. The route out of work, husbands aged 50 and over

State at $t-2$ and $t-1$	Couple-waves	Retired at t	Rate
Working, working (direct)	5,367	422	0.0786
Working, unemployed	100	7	0.0700
Working, inactive	63	11	0.1746
Working, working, same industry	4,971	381	0.0766
Working, working, changed industry	282	27	0.0957

Notes: Husbands over 50 observed in three consecutive waves, on the full couple panel. Each row conditions on the husband's states at $t-2$ and $t-1$: Couple-waves counts the observations with that history, Retired at t the ones that end in retirement, and the rate is their ratio. The industry comparison is between the two waves in work and is defined only for men working at both. Of the 692 men over 50 ever observed retired, 21 were ever observed unemployed at an earlier wave (55 of 2,426 couple-waves). The exposure regression is not estimable on the two indirect routes: with couple fixed effects, 100 and 63 observations spread over almost as many couples leave nothing within couple to identify it.

The effect of exposure is identified from the interaction of the predetermined score with the crossing of the institutional age, holding the couple and the year fixed:

$$y_{ct} \;=\; \alpha_c + \delta_t + \beta\,(z_c \times \mathbf{1}[a_{ct} \geq 55]) + \varepsilon_{ct}. \tag{4}$$

Here y_{ct} is the outcome for couple c in year t, in Fact 1 the indicator that a husband working at $t-1$ has retired by t; α_c and δ_t are couple and year fixed effects; z_c is the husband's standardized occupational exposure and a_{ct} his age, so $\mathbf{1}[a_{ct} \geq 55]$ marks his having reached the mid-fifties eligibility age; and ε_{ct} is an idiosyncratic error. The couple effect absorbs z_c and the year effect absorbs the national robot stock, so β is identified from the within-couple change as a man of given exposure crosses the threshold: it compares two men of different exposure crossing age 55 in the same calendar year, and a positive β means the more exposed retire sooner once they reach eligibility. Table 5 reports it for the three destinations on the estimation sample. A one-standard-deviation increase in exposure raises the retirement hazard by 6.0 percentage points at the threshold, against a base of 11.9 percentage points at ages 55–59, and has no effect on unemployment or inactivity.

Table 5. The effect of occupational robot exposure on the route out of work

	Retirement	Unemployment	Inactivity
Exposure $\times \mathbf{1}[\text{age} \geq 55]$	+0.0599***	−0.0005	+0.0032
	(0.0170)	(0.0061)	(0.0064)
Couple-waves	4,337	4,337	4,337
Couples straddling age 55	263	263	263
Retirement events, straddling couples	81	81	81

Notes: Linear probability models on the risk set of Table 3, husbands aged 50–69, with couple and year fixed effects; standard errors clustered on the couple. Exposure is the husband's first-observed occupational robot-exposure score, standardized. The couple fixed effect absorbs exposure itself; identification comes from the same man crossing age 55 in the same calendar year as men of different exposure, so only the 263 couples observed on both sides of the threshold contribute, and they carry 81 of the 432 retirement events. Table B1 reports the estimate on those couples alone and under four departures from the specification, and Table 6 reports the retirement coefficient under four inference schemes. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$, on the couple-clustered standard error.

On the couple panel, pooling employees with the self-employed, the same equation gives +0.0374, and Table 6 reports that coefficient under four inference schemes. Clustered on the couple its standard error is 0.0126; clustered on the occupation cells at which treatment actually varies it is 0.0265, and a wild cluster bootstrap-t at that level returns a p-value of 0.1240. Randomization inference that permutes the occupation scores across occupation cells and re-estimates the same equation returns 0.0579. The point estimate is unaffected by the choice of scheme, and the case for the effect rests on the design and on the placebo pattern. Both placebos are flat under every scheme: the coefficient on the flow into unemployment is −0.0000 and on the flow into other inactivity +0.0007, with design-based p-values above 0.88.

Table 6. Equation (4) under four inference schemes

Flow out of work	Coefficient	Clustered standard error: Couple ($G = 1{,}880$)	Clustered standard error: Occupation ($G = 18$)	Wild cluster bootstrap-t p	Randomization inference p
Into retirement	+0.0374	0.0126	0.0265	0.1240	0.0579
		[0.0031]	[0.1770]		
Into unemployment	−0.0000	0.0044	0.0030	0.9930	0.9421
		[0.9944]	[0.9919]		
Into other inactivity	+0.0007	0.0050	0.0052	0.8810	0.9381
		[0.8890]	[0.8954]		

Notes: All four schemes are applied to the same regression, equation (4) on 6,817 couple-waves with couple and year fixed effects. p-values from the clustered standard errors are in square brackets. Exposure is an occupational score taking 18 distinct values, so the occupation cell is the level at which treatment varies. The wild cluster bootstrap-t is at the occupation level with Rademacher weights, the null imposed, and 999 replications. Randomization inference permutes the 18 occupation scores across occupation cells, rebuilds the standardized exposure and the interaction, and re-estimates the same equation, 1,000 replications; the permutation standard deviation of the retirement coefficient is 0.0193, larger than the couple-clustered standard error, which is why the design-based p-value is the larger one. Clustering on the 11 regions instead gives a standard error of 0.0116 on the retirement row ($p = 0.0090$), with fewer clusters still.

The pooled coefficient averages two populations with different institutions. Table 7 splits the risk set on the husband's status in the previous wave. Among employees the coefficient is +0.0608 (standard error 0.0172, 434 retirements in 4,283 couple-waves); among the self-employed it is +0.0042 (0.0126, 144 retirements in 2,534 couple-waves), which the data cannot distinguish from zero. The two groups partition the couple panel's risk set, and the split has a direct institutional reading: the seniority pension was a scheme for private-sector employees and the self-employed contributed to separate funds, so the men with no response are the men for whom the threshold in equation (4) describes no institution. It is also why the model is estimated on employee households.

Table 7. The retirement response by pension scheme

	Employee at $t-1$	Self-employed at $t-1$
Exposure $\times\mathbf{1}[\text{age} \geq 55]$	+0.0608***	+0.0042
	(0.0172)	(0.0126)
Base retirement hazard	0.1013	0.0568
Retirements	434	144
Couple-waves	4,283	2,534
Couples	1,311	703

Notes: Linear probability models on the full couple panel at ages 50–69, split on the husband's activity status in the previous wave; couple and year fixed effects, standard errors clustered on the couple. The two columns partition 6,817 couple-waves and 578 retirements. The seniority pension was a scheme for private-sector employees; the self-employed contributed to the separate artisan and trader funds, whose seniority conditions the same threshold does not describe. This is why the model is estimated on employee couples. Stars as in Table 5.

The response is also specific to the husband. The same score built from the wife's occupation, interacted with her crossing of 55, moves her retirement hazard from work by -0.0144 (standard error 0.0227) on 740 working wives with 152 retirements. Working wives of these cohorts rarely hold the thirty-five contribution years the seniority rule turns on, so the option his design exploits mostly does not exist for her at these ages.

Panel B of Table B1, in Appendix B, reports the coefficient under four departures from equation (4). Adding the age-55 main effect, which the couple and year effects do not absorb and which is correlated with the interaction because mean exposure in the risk set is -0.16, gives +0.0587 (0.0183). Allowing the threshold to differ by schooling gives +0.0443 (0.0218). Rebuilding the score from the husband's modal occupation in place of his first valid one gives +0.0544 (0.0168), and leaving the survey's missing codes unscreened gives +0.0604 (0.0170). The demanding control is the remaining one: giving each occupation cell its own year effect leaves +0.0235 (0.0203), so that more than half of the estimate is common to an occupation in a year and is not specific to the men crossing the threshold in it. Imposing all four at once leaves +0.0092 (0.0246). The sign is positive in every column and the nulls on unemployment and inactivity are undisturbed by any of them. The magnitude is less settled than the sign: six percentage points is the estimate under the baseline specification, and the demanding controls bring it down while leaving it positive.

The schooling departure exists because eligibility and exposure arrive together, as Section II.E documents: a tertiary-educated husband's retirement hazard steps down by 14.8 points at 55 relative to one with lower-secondary schooling or less, and entering that interaction alongside exposure moves the coefficient from +0.0601, its value on the subsample with observed schooling, to the table's +0.0443. A further check replaces the age threshold with the predicted eligibility that Section II.E builds from the reported age of first work: the interaction is +0.0237 (0.0162), and once the exposure-by-age profile is absorbed non-parametrically it is +0.0067 (0.0143), a null too imprecise to refute the threshold estimate. The eligibility interaction has within-couple variation for 281 couples and 78 events against 393 couples and 93 events for the age threshold, and its minimum detectable effect, 0.0402, exceeds the estimate it was meant to test.

Figure 1 shows the result without a regression, in the schooling group where the comparison is cleanest. Upper-secondary men started work at similar ages, so their thirty-five contribution years arrive together, near 55. Within this group the retirement hazard rises with age for both exposure groups and the gap between them opens at 55, where the seniority pension becomes available. Below that age exposed and unexposed men behave identically, because exposure raises the value of stopping and there is nowhere to stop. Appendix B draws the same figure for the other schooling groups: among the least schooled, eligibility arrives before the early fifties for exposed and unexposed alike and both retire early, while among the tertiary-educated retirement begins only in the late fifties and the cells are thin. The pooled profile mixes the groups, and its gap opens in the early fifties, where the option of the least schooled arrives.

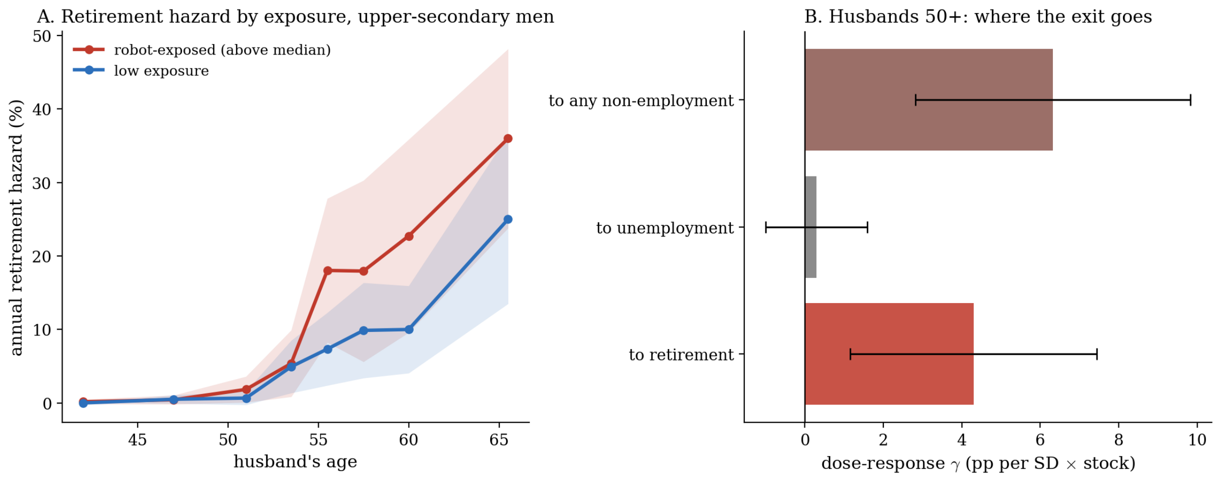

FIGURE 1. RETIREMENT HAZARDS BY AGE AND EXPOSURE FOR UPPER-SECONDARY MEN

Notes: The sample is the estimation sample: couples in which the husband is an employee in the first wave he is observed working, with his occupation score fixed at his first valid detailed occupation report. Panel A restricts to the 1,253 couples in which the husband has upper-secondary schooling (5,426 couple-waves at risk) and plots the annual hazard from work into retirement by the husband's age, splitting these couples at the median of their standardized exposure scores; shaded bands are pointwise 95 percent confidence intervals clustered by couple, and Appendix B draws the same panel for the other schooling groups. Panel B reports, for the 698 couples of the full estimation sample in which the husband is first observed at age 50 or older (2,572 couple-waves at risk), the coefficient on the exposure score interacted with the aggregate robot stock from a linear probability model with couple and year fixed effects; whiskers are 95 percent confidence intervals clustered by couple.

Equation (4) assumes that, absent the seniority pension, high- and low-exposure men would have retired in parallel. Figure 2 makes the comparison as a two-by-two on the couple panel: the raw retirement hazard of upper-secondary men before and after 55, split at the median exposure. The two groups retire at the same rate before the threshold, 0.0110 against 0.0106 per year, and diverge at it, for a difference-in-differences of +0.060, smaller than among the employees of Figure 1 because the couple panel adds the self-employed, whose response is absent. The flat pre-period is informative about one confound: the rule turns on contribution years, so a man who started work unusually early clears thirty-five of them before 55, and if exposure were correlated with early career starts within this group, exposed men would retire more before the threshold too. Appendix B reports the same comparison by tercile of exposure, the destination-specific coefficients showing the response is confined to retirement, and a randomization-inference test against the reassignment of scores across occupations.

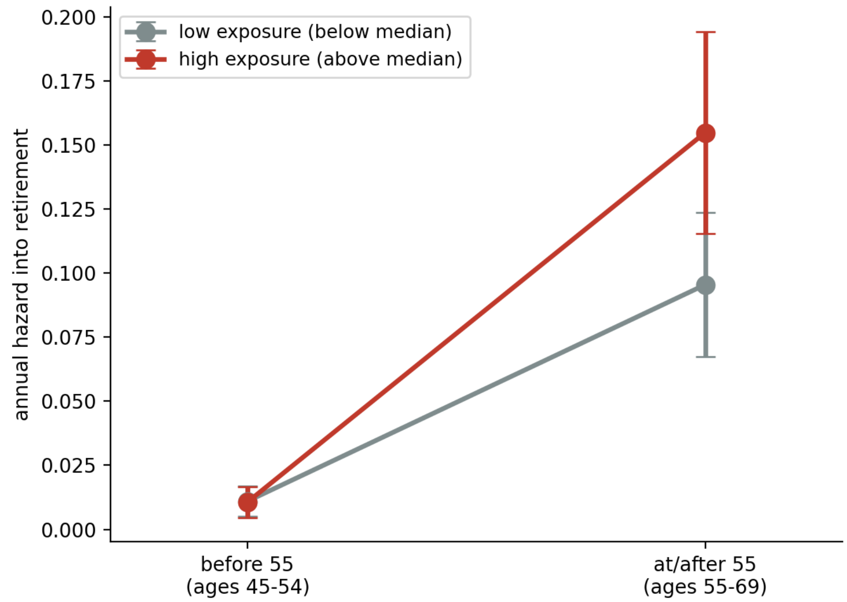

Figure 2. Retirement before and after the eligibility age, by exposure, for upper-secondary men

Notes: Annual hazard from work into retirement for husbands with upper-secondary schooling, before (ages 45–54) and at or after (ages 55–69) the mid-fifties eligibility ages, split at the median of the standardized exposure score. Whiskers are 95 percent confidence intervals. Upper-secondary men are the schooling group whose thirty-five contribution years are completed near age 55, as Section II.E discusses.

A. *Is the Exit a Push or a Pull?*

Whether automation makes a man choose to stop working or makes his employer stop him is central to the reform experiment: a chosen exit sits at an indifference margin, so delaying his eligibility keeps the man at work at second-order cost, while a dismissed man who cannot retire must land somewhere else. Italian institutions of the 1990s could in principle deliver either: seniority pensions on the one hand, employer-sponsored early retirement and long mobility bridges for dismissed older workers on the other. The two readings are close to observationally equivalent in the size of the exit response itself. Table 8 reports three kinds of evidence that separate them.

Table 8. The mechanism: how men leave work

Panel A. The monthly route out of work, 427 passages into retirement at 50–69

	Share	(s.e.)	Count	Exposure gradient	
Seamless: retirement follows the last month of work	0.9415	(0.0114)	402	−0.0339	(0.0123)
Via any non-employed month	0.0585	(0.0114)	25	+0.0346	(0.0122)
Via an unemployed month	0.0211	(0.0070)	9	+0.0149	(0.0085)
Via an inactive month	0.0398	(0.0095)	17	—	
Mean gap, months	0.7073	(0.2001)			

Panel B. The stated reason for stopping work

	Into retirement ($n = 342$)		Into unemployment or inactivity ($n = 79$)	
Reason	Share	Count	Share	Count
Retired at normal age	0.8801	301	0.0886	7
Obliged to stop by employer (closure, redundancy, early retirement, dismissal)	0.0439	15	0.5190	41
End of temporary contract	0.0000	0	0.2278	18
Closure or sale of own business	0.0000	0	0.0127	1
Own illness or disability	0.0175	6	0.0759	6
Wanted to retire or live off private means	0.0351	12	0.0506	4
Family, care, study, other	0.0234	8	0.0253	2

Panel C. The falsification: the same design at ages with no pension

Flow	Coefficient	(s.e.)	Couple-waves	Events
Into unemployment, ages 30–54	−0.0002	(0.0026)	11,433	145
Into inactivity, ages 30–54	+0.0018	(0.0025)	11,433	100
Into retirement, ages 30–54	+0.0065	(0.0019)	11,433	123
Into unemployment, ages 50–69	+0.0005	(0.0054)	4,336	73
Into inactivity, ages 50–69	+0.0095	(0.0061)	4,336	74
Into retirement, ages 50–69	+0.0470	(0.0118)	4,336	432

Notes: Estimation sample throughout. Panel A chains the survey's monthly calendar of activities across waves into a continuous monthly record of each man's status, and classifies each observed passage from work into retirement by whether any non-employed month intervenes; shares carry binomial standard errors and the gradient is the coefficient on the standardized exposure score with year and exit-age indicators, clustered on the couple. Panel B reports the stated reason for stopping work; all respondents answered the not-working version of the item, whose first category is retirement at the normal age and whose employer category includes employer-sponsored early retirement. Panel C is the paper's own within-couple design (exposure interacted with the national robot stock, couple and year fixed effects, standard errors clustered on the couple), applied to each flow at ages where no pension is available and at ages where one is.

First, the route. Chaining the survey's monthly calendar of activities across waves yields a continuous monthly record of each man's status on the estimation sample. Of 427 observed passages from work into retirement at ages 50 to 69, 94 percent are seamless, with the first month of retirement following the last month of work; 5.9 percent cross a month of non-employment, and 2.1 percent a month of unemployment.

Second, the stated reason. Among the men observed moving from work into retirement, 88 percent report that they retired at the normal age and 4 percent that their employer obliged them to stop, a category that explicitly includes employer-sponsored early retirement. Among the men moving from work into unemployment or inactivity the same question returns the opposite composition: 52 percent employer-initiated and 23 percent an expired temporary contract. The question can detect a dismissal; among retirements it finds almost none, and its employer-initiated share does not rise detectably with exposure.

Third, the falsification. If firms shed exposed older workers and the pension were merely the least costly exit, the shedding should be visible at ages where no pension is available, but it is not. In the same within-couple design that identifies the retirement response, automation moves the flow from work into unemployment at ages 30 to 54 by −0.0002 (standard error 0.0026) and the flow into inactivity by +0.0018 (0.0025), against +0.0470 (0.0118) on the flow into retirement at 50 to 69. Retirements do not cluster within sector and year beyond what independent exits would produce, and the gradient is as strong among public-sector employees, whom the collective-dismissal machinery did not touch, as elsewhere, at +0.0579 (0.0218).

The data do leave a residual. About 3 percent of matched annual retirement events carry an involuntary marker, and at monthly resolution 5.9 percent of passages do. That monthly share rises with exposure, by 3.5 percentage points per standard deviation (standard error 1.2), as does the receipt of an unemployment-type benefit in the year of separation. Exposure therefore shifts not only how many men leave but, at the margin, how a small minority of them leave. Section IV describes how the model treats this residual and Section VI reports the bound it implies for the reform.

FACT 2: *Exposure does not change what workers are paid over the diffusion period.*

Table 9 reports the price margin. Conditional on employment, exposure is associated with no change in pay as robots diffuse: on the estimation sample the within-couple coefficient on log earnings is +0.0008 (standard error 0.0233) at ages 50–59, where the retirement response is largest, and +0.0075 (0.0072) pooled across ages; the winsorized level and inverse-hyperbolic-sine versions are equally null. On the couple panel the pooled coefficient is +0.0268 (0.0103): the positive drift is contributed by the self-employed, who lie outside the model. Earnings inclusive of zeros do fall at ages 40–59: that is the employment margin of Fact 1 appearing in an earnings measure, and exposure does not move the price.

Table 9. Exposure and pay: the price margin

	Husband's baseline age				Pooled
Outcome	25–39	40–49	50–59	60–69	25–69
Panel A. Estimation sample: employee couples					
Log earnings | employed	+0.0114	−0.0030	+0.0008	+0.1239	+0.0075
	(0.0109)	(0.0108)	(0.0233)	(0.1559)	(0.0072)
	7,946	5,904	2,640	246	16,736
Winsorized level | employed, EUR	−143.3	−241.8	−482.0	−421.7	−134.2
	(184.4)	(191.2)	(346.0)	(1,837.6)	(122.5)
	8,069	5,977	2,713	257	17,016
Inverse hyperbolic sine | employed	+0.0210	−0.0178	−0.0349	−0.6700	+0.0023
	(0.0294)	(0.0335)	(0.0717)	(0.8140)	(0.0210)
	8,069	5,977	2,713	257	17,016
Log earnings, level design	−0.1034	−0.1450	−0.2722	−0.7226	−0.1415
	(0.0126)	(0.0199)	(0.0392)	(0.3598)	(0.0082)
	7,946	5,904	2,640	246	16,736
Panel B. Full couple panel					
Log earnings | employed	+0.0253	+0.0139	+0.0508	+0.1909	+0.0268
	(0.0142)	(0.0155)	(0.0316)	(0.2001)	(0.0103)
	10,657	8,159	4,234	650	23,700
Winsorized level | employed, EUR	−17.1	−92.0	+159.6	+965.7	+40.0
	(187.1)	(200.6)	(363.4)	(1,719.1)	(130.1)
	10,941	8,337	4,423	695	24,396
Inverse hyperbolic sine | employed	+0.0535	+0.0674	+0.1213	+0.3163	+0.0761
	(0.0362)	(0.0380)	(0.0738)	(0.4458)	(0.0251)
	10,941	8,337	4,423	695	24,396
Log earnings, level design	−0.1096	−0.1641	−0.2831	−0.5271	−0.1601
	(0.0137)	(0.0193)	(0.0349)	(0.1527)	(0.0093)
	10,657	8,159	4,234	650	23,700
Panel C. Earnings including zeros: this mixes in the employment margin					
Winsorized level including zeros, EUR	−108.3	−889.5	−1,234.2	+3,341.6	−311.3
	(201.1)	(268.2)	(476.8)	(1,191.2)	(202.9)

Table 9, continued

Outcome	Husband's baseline age 25–39	40–49	50–59	60–69	Pooled 25–69
	8,417	6,578	3,975	641	19,611
Inverse hyperbolic sine including zeros	+0.0886	−0.1991	−0.6234	+0.4316	−0.1044
	(0.0472)	(0.0800)	(0.1650)	(0.3557)	(0.0644)
	8,417	6,578	3,975	641	19,611

Notes: Each cell is a separate regression. In the trend design the outcome is regressed on the standardized exposure score interacted with the national robot stock, with couple and year fixed effects; the level design replaces the interaction with the score itself and the fixed effects with year and single-year-of-age indicators. Standard errors, in parentheses, are clustered on the couple; the third line of each block is the number of observations. Winsorized levels are in euros, reported to one decimal place. Panels A and B condition on the husband being employed in the wave; Panel C does not, so movements in its outcome come from the employment margin of Fact 1.

Exposed occupations do pay less, and by a great deal: in levels the differential is about 14 log points per standard deviation of exposure. This differential, however, predates the diffusion period: it is as large at ages 25–39, where no retirement response exists, as at 50–59, and it does not widen as robots arrive; the trend design tests exactly that widening and fails to reject zero. Exposure changes whether older men work. It does not change what workers are paid. That rules out modeling automation as a price and forces it into preferences.

Fact 3: *The two exits differ in how well they are insured.*

Among state-stable households, the husband's personal income when retired is 0.8156 of his income when working, and when unemployed 0.1436, a factor of nearly six. Table 10 reports the full budget. Retirement is close to fully insured and unemployment is essentially uninsured: with the wife out of work, household consumption falls by 15 percent when he retires and by 70 percent when he loses his job. The model's account of the household rests on this ratio, and it is measured directly from the income data.

Table 10. The household budget

Panel A. Personal income by state				
	Working	Unemployed	Inactive	Retired
Husband, y_h	1.0000	0.1436	0.4000	0.8156
	Employed	Non-employed	Retired	
Wife, y_w	0.6933	0.0343	0.5293	
Other household income, Y_0	0.1942			

Panel B. Implied household consumption			
	She employed	She non-employed	She retired
He works	1.8876	1.2285	1.7235
He unemployed	1.0311	0.3721	0.8671
He inactive	1.2875	0.6285	1.1235
He retired	1.7032	1.0441	1.5392

Notes: All quantities are normalized by the husband's net personal income when working, 15,940 EUR per year. Built from ECHP personal income components (total net personal income and its work, old-age pension, unemployment benefit and sickness or invalidity benefit components) on state-stable households, so that every euro is attributed to a person. None of these quantities is estimated.

Two channels outside the couple could in principle be doing the insuring, and neither does much. Within couple, the part of household income belonging to neither spouse rises by 0.0364 of his working earnings when he is unemployed and does not move when he retires, so co-resident earnings and family transfers close a small fraction of the gap. Public transfers move in the opposite direction from the replacement rate: the old-age pension component is worth 0.7100 of his working income to a retired man while the unemployment benefit component is worth 0.0273 to an unemployed one.

Fact 4: *Neither spouse increases labor supply when the other stops working.*

Table 11 reports within-couple regressions of each spouse's employment on the other's state, for couples in which both are aged 45–69. There is no added-worker effect in either direction. Both spouses work less when the partner stops, or do not move. The wife's employment falls by 3.5 percentage points when her husband retires and by 2.4 when he loses his job. The point estimates are ordered as an insurance account would predict, with the larger withdrawal where the loss is smaller, but her unemployment-state coefficient is estimated on a thin cell, its 95 percent interval runs from −0.0874 to +0.0390, and equality of the two responses cannot be rejected. The data do establish the sign: no coefficient in the table is positive. His employment falls by 13.2 points when she retires and by 13.1 when she is unemployed, but by only 2.5 points, insignificantly, when she is inactive, the margin on which most of her non-employment sits and the one least contaminated by shocks that hit both spouses at once.

Table 11. Within-couple cross-spouse labor supply responses, couples aged 45–69

	Partner's state (working is the omitted category)		
	Retired	Unemployed	Inactive
Dependent variable	Coef. (s.e.)	Coef. (s.e.)	Coef. (s.e.)
Her employment	−0.0353** (0.0167)	−0.0242 (0.0323)	−0.0303 (0.0262)
His employment	−0.1316*** (0.0427)	−0.1306** (0.0562)	−0.0247 (0.0321)
Her retirement on his	+0.0531*** (0.0152)	(joint retirement)	

Notes: Linear probability models with couple and year fixed effects; standard errors clustered on the couple. Sample: couples in which both spouses are aged 45–69, 7,577 couple-waves. Her unemployment cell is thin: transitions from her work into her unemployment are rare in the panel, and the 95 percent interval on that coefficient runs from −0.0874 to +0.0390. Pooling her unemployment and inactivity into a single non-employment state gives a coefficient on his employment of −0.0357 (s.e. 0.0309), which is the object the model of Section IV speaks to. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The income at stake is large. Conditional on employment the wife earns 0.6933 of her husband's income when working, 81 percent of what he loses when he becomes unemployed, and in the household's budget his job loss cuts consumption by 70 percent if she is out of work and by 45 percent if she is employed. The wife leaves this insurance unused, and Section V.F takes up the reason inside the model.

Fact 5: *Couples retire together.*

Within couple, the wife's retirement rises by 5.3 percentage points when her husband is retired. Figure 3 shows the stock and the flow. Her retirement rate is 0.2032 when he is retired against 0.0827 when he works, and 0.0558 when he is unemployed, which is lower than when he works. The flow behind the stock is her exit from employment at her ages 45–69: 0.0746 per year while he stays at work and 0.2389 in the year he moves into retirement, a cell of 113 transitions whose interval excludes the stay-at-work rate. In the year he moves into unemployment her exit is 0.2105, but that cell holds 19 transitions and its interval includes the stay-at-work rate: the joint exit is specific to retirement.

Conditioning on his state is a choice of presentation, and the reverse conditioning gives the same picture: his exit from work is 0.2857 in the year she moves from employment into retirement, on 56 transitions. Both spouses usually retire in the same year, so the annual data cannot say who moved first. Similar ages alone could make the spouses' retirement years close, with each on an age-driven schedule of their own; the timing says more. Conditional on single-year age dummies for both spouses and on year effects, her exit from employment remains 0.1235 (standard error 0.0410) higher in exactly the year he moves into retirement. The model of Section IV accordingly treats the couple's choice as joint and imposes no direction.

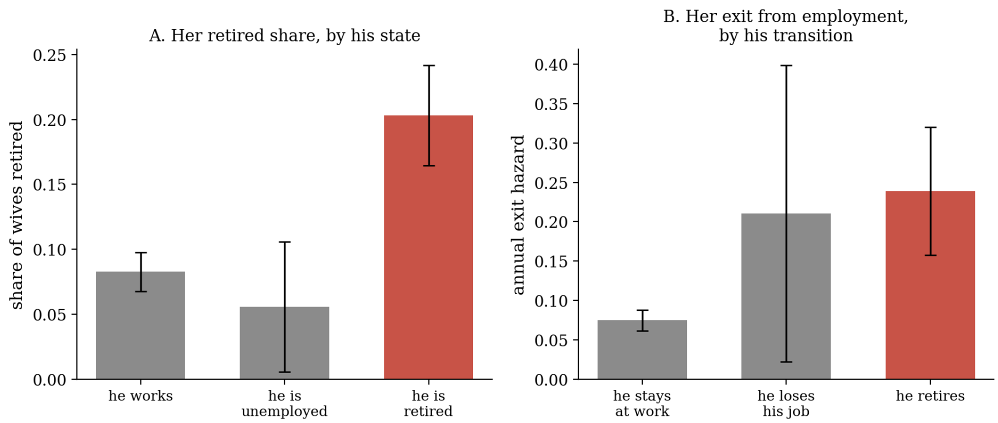

Figure 3. Joint retirement, by the husband's state

Notes: Estimation sample. Panel A: the share of wives retired by the husband's contemporaneous state, for couples with both spouses aged 45–69. Panel B: the wife's annual exit hazard from employment at her ages 45–69, by the husband's transition between the same two waves, conditional on his working at the first. Whiskers are 95 percent confidence intervals clustered by couple.

The five facts fit together. Exposure moves older men into retirement and into nothing else, at the ages the institution opens the exit, among the employees whose fund the rule governed, and without touching pay. The two exits open to an exposed man are insured at 82 and 14 percent. Neither spouse adds labor supply when the other stops working, and couples retire together, with retirement the only exit they share. The design identifies this pattern more sharply than the size of the gradient, which the inference schemes and the departures above place between zero and six percentage points. Section IV builds a life-cycle model of the couple around these facts, and Section V estimates it.

IV. A Life-Cycle Model of Household Exit

A. Environment

A couple is observed annually over the husband's age $a = 25, \ldots, 79$. The wife's age is $a_w = a - \Delta$, where the age gap Δ is drawn from its empirical distribution. There is no mortality: both spouses reach the terminal age, and every discounted calculation below runs over that certain horizon. The husband occupies a state $s \in \{W, U, O, R\}$ (working, unemployed, inactive, retired) and the wife a state $n \in \{E, N, R\}$ (employed, non-employed, retired). Her unemployment is pooled into non-employment: the panel's cell for transitions from her work into her unemployment is thin, and the coefficient it supports in Table 11 is too imprecise to support a separate state.

The household maximizes a single objective. This unitary assumption is restrictive in general: Chiappori (1988, 1992) show that spouses with distinct preferences who bargain efficiently generate labor supply functions a unitary model rejects, and Mazzocco (2007) rejects the full-commitment version of the collective model on intertemporal data. It is a safe simplification here for two reasons. First, the result to be explained is an absence of movement: neither spouse changes labor supply when the other stops working. Bargaining would show up as one spouse's labor supply moving against the other's, and the data show neither moving. Second, the wife's employment in this sample is limited by how rarely job offers arrive, so her bargaining weight has little to act on. A collective model would let the husband's pension entitlement shift that weight, but identifying the shift requires a variable that moves bargaining power without moving the budget, and the panel has none.

The couple is hand to mouth. Consumption is

$$c(s, n) \;=\; Y_0 + y_h(s) + y_w(n), \tag{5}$$

where $y_h(s)$ is the husband's personal income in state s, one when he works and the measured replacement rates otherwise, $y_w(n)$ is the wife's personal income in state n, and Y_0 is the household income belonging to neither spouse. All quantities are normalized by the husband's net personal income when working; every element of (5) is set from the data in Table 10, and none is estimated. The absence of assets is a restriction, imposed because the survey records neither wealth nor consumption, so nothing in it would identify an asset process. Two later sections measure what the restriction costs. For the welfare results it costs little: Section VI shows that granting perfect consumption smoothing, more than any asset could deliver, moves the headline cost by 4.3 percent of itself, because a delayed pension is expensive through the disutility of continued work, which no asset can insure. For the household mechanism, savings offer a competing explanation of the missing added-worker effect that the survey cannot test; Section V.F states it and what would distinguish it.

A last restriction concerns the husband: no firm can end his job in order to retire him. This restriction matters for the counterfactual, because a man pushed out instead of choosing to leave would, once his eligibility is delayed, land somewhere other than at work. It is imposed because the moments that would identify a firm-side channel are essentially zero. Among matched work-to-retirement transitions, 3.2 percent carry a firm-side marker, with an exposure gradient of +0.0082 (standard error 0.0103), both reported in Section V; 88 percent of men entering retirement report a normal-age exit against 4 percent employer-initiated; and in the within-couple design of Table 8 exposure does not raise the flow into unemployment at ages 30 to 54, where no pension can disguise a dismissal. The one qualification is the monthly record, in which 5.9 percent of passages into retirement cross a non-employed month, a share that rises with exposure.

Estimated freely, the channel returns zero; the reported model fixes it there, so the two involuntary moments are zero by construction, and the counterfactual of Section VI reports an imposed bound. With a firm-side release set at the involuntary share the data record and every released ineligible man routed to unemployment, 86 percent of the reform's employment gain survives, and the destination mix shifts toward unemployment (Section VI).

B. Preferences

Period utility is

$$\begin{aligned} u(s,n,a,z) \;=\; \log c(s,n) \;-\; \theta_h(a,z)\,\mathbf{1}[s=W] \;-\; \theta_w(a_w)\,\mathbf{1}[n=E] \\ -\; \kappa\,\mathbf{1}[n=E,\; n_{-1}\neq E] \;+\; \phi\,\mathbf{1}[s\neq W]\,\mathbf{1}[n\neq E], \end{aligned} \tag{6}$$

where

$$\theta_h(a,z)=\theta_0+\theta_1(a-50)^+ +\theta_2(a-60)^+ +\delta z, \qquad \theta_w(a_w)=\theta_{w0}+\theta_{w1}(a_w-45)^+. \tag{7}$$

Here z is the husband's standardized occupational exposure of Section II.D, $(x)^+ = \max(x,0)$, $\mathbf{1}[\cdot]$ is the indicator function, and n_{-1} is the wife's state in the previous period.

The costs of working θ_h and θ_w carry estimated intercepts, age slopes, and the exposure loading δ; κ is a fixed cost the wife pays in a period in which she enters employment; and ϕ is the utility of a period in which neither spouse works. Three features of (6) and (7) carry the paper's mechanisms.

The husband's cost of continuing to work rises with age through θ_1 and θ_2. Without that slope, no combination of pension generosity and eligibility timing produces a retirement hazard that rises tenfold between ages 50 and 69, because the value of the outside option is then the same at 52 and at 67.

Exposure enters as a level shift in that cost, δz, and nowhere else; Fact 2 rules out a price channel. The consequence is that exposure can act only on the retirement decision: separations into unemployment and inactivity are exogenous, so an exposed man who cannot retire has nowhere for the extra cost to send him. The model therefore predicts that exposure raises retirement, leaves unemployment and inactivity alone, and has no effect before eligibility, without any of those three being imposed.

The joint-leisure term ϕ is symmetric by construction across the husband's non-working states. It enters his retirement, his unemployment and his inactivity identically. Making it specific to retirement would build the asymmetry into the preferences by assumption. Here it emerges from the budget, and Section V.F shows how.

C. Institutions and Frictions

Pension eligibility is a persistent, probabilistic characteristic of the man. Each husband draws a latent eligibility age a_h^* from a logistic distribution with location a_{0h} and scale τ_h with probability $\bar{e}_h$, and never qualifies otherwise; he may retire at a if and only if $a \geq a_h^*$. This is the natural representation of a rule that turns on contribution years, and it lets the model generate a positive retirement hazard before any statutory age (4.4 percent at ages 50–54), which a hard threshold forces to zero. The representation has direct support: as Section II.E reports, men who had cleared thirty-five contribution years retired six to seven percentage points per year more often than men of the same age who had not, so the model's eligibility process is a contribution-years rule read off behavior.

The wife's eligibility follows the same distribution with its own parameters, with one addition: it accrues only in periods in which she is employed. A pension has to be contributed to. Without the restriction, every eligible non-employed wife claims immediately (her retirement income is 0.5293 of her husband's working income against 0.0343 out of work), all of her retirement arrives from non-employment, and the flow from employment into retirement, which the data put at 0.2838 at ages 60–69, is zero.

The wife's entry into employment is rationed. A non-employed wife receives an employment opportunity with probability

$$\mu(a_w) \;=\; \Lambda(m_0 + m_1(a_w - 45))\,, \tag{8}$$

where Λ is the logistic function and m_0 and m_1 are estimated parameters, and only then chooses whether to take it, paying the entry cost κ of (6). Her exit margin is unrationed. The restriction is read off the data: the entry hazard of non-employed wives falls from 0.0963 at ages 25–34 to 0.0289 at 55–59 and 0.0194 at 60–69, while her exit hazard from employment shows no comparable decline, and female employment in this sample falls from 47 percent at her ages 45–49 to 23 percent at 55–59.

The husband's employment risk is exogenous. From work he is separated into unemployment with probability π_U and into inactivity with probability π_O; from unemployment and inactivity a job opportunity arrives with probability λ_U and λ_O. Fact 1 supports treating the flow into non-employment as a shock: the hazard is flat in age and exposure does not move it. The flow out of inactivity is a choice, which is the model's one behavioral disability margin, and Section VI reports what the data say about it.

D. Choices and Taste Shocks

Within a period the couple observes its inherited states and the realized opportunities, and chooses jointly from the feasible sets they define. From work the husband chooses between continuing and retiring if eligible; if separated, between the state he was pushed into and retirement; from unemployment or inactivity, between accepting an arriving offer, remaining, and retiring. Retirement is absorbing. The wife chooses among employment (if she is already employed or has received an opportunity), non-employment, and retirement if eligible.

Taste shocks are nested extreme value in three levels, so the two spouses' choices can be differently noisy without imposed symmetry between them. Let $M(s', n')$ denote the deterministic value of the pair (s', n'),

$$M(s', n') \;=\; u(s', n', a, z) \;+\; \beta\,\mathbb{E}\,V_{a+1}(s', n', e_h', e_w')\,, \tag{9}$$

where β is the discount factor, e_h and e_w indicate the spouses' pension eligibility, whose processes the previous subsection describes, and V_{a+1} is next period's value function, built in (10)–(12) below. The expectation is over next period's eligibility, which evolves exogenously for the husband and, for the wife, only if she works this period. The husband's choice sits in the innermost nest with scale σ_h, the wife's choice between her two non-working states in the middle nest with scale σ_n, and her choice between working and not working in the outer nest with scale σ_w:

$$I(n') = \sigma_h \log \sum_{s' \in D_h} \exp(M(s', n')/\sigma_h)\,, \tag{10}$$

$$J = \sigma_n \log \sum_{n' \in \{N,R\} \cap D_w} \exp(I(n')/\sigma_n)\,, \tag{11}$$

$$V_a = \sigma_w \log[\exp(I(E)/\sigma_w)\,\mathbf{1}[E \in D_w] + \exp(J/\sigma_w)] + \sigma_w \gamma, \tag{12}$$

with $\sigma_n = \rho_n \sigma_w$ and $\sigma_h = \rho\, \sigma_n$, so that $\sigma_h \leq \sigma_n \leq \sigma_w$ holds by construction and the system is consistent with random utility maximization (McFadden, 1978). Here D_h and D_w are the feasible sets that the inherited states and realized opportunities define for the husband and for the wife, $I(n')$ is the inclusive value of his choice given her state n', J is the inclusive value of her two non-working options, and γ is Euler's constant. The middle nest is necessary. Her non-employment and her retirement are both states in which she does not work, while either is far from employment, so a single scale over her three options cannot make her exits from employment split as sharply as they do: at her ages 60–69 the data split them seven to one in favor of retirement, 0.2838 against 0.0405.

The state vector is $(a, s, n, e_h, e_w, z, \Delta)$ and the model is solved by backward induction. Appendix C gives the recursion and the numerical implementation.

V. Estimation

A. What Is Set Outside the Model

The discount factor is 0.96 per year, the conventional value in annual life-cycle models: it implies an annual discount rate of about four percent, the rate Gourinchas and Parker (2002) estimate from life-cycle consumption profiles. The budget of (5) is fixed at the values of Table 10: each spouse's personal income in each state, and the household income belonging to neither spouse, measured from the ECHP income components on the state-stable households of Section II, in units of the husband's income when working. Nothing in the budget is estimated. The distributions of exposure and of the spousal age gap are the empirical ones. Everything else is estimated: 23 parameters, all of them preferences, frictions, or institutions.

The model is simulated on the observed sampling design. Because the taste shocks and the mobility and eligibility draws are random, each of the 3,375 couples of the estimation sample is simulated twelve times, and every model moment averages over the twelve copies: the copies reduce simulation noise, twelve balances that noise against computing time, and the standard errors below carry the correction for the noise that remains. Each copy inherits the couple's own exposure score, age gap, birth cohort, and observation window, and the simulated panel then passes through the same sample rule as the data, that the husband be seen at work at least once inside the window. Data and model moments are therefore computed by the same estimator on panels of the same shape, and a gap between them can only come from the model's behavior. As a check, the same estimator code, run on the real panel, returns the data moments exactly.

B. Identification and the Search

The criterion is the sum of squared t-statistics of 54 moments,

$$Q(\Theta) \;=\; \sum_k \left(\frac{m_k(\Theta) - \hat{m}_k}{s_k}\right)^2, \tag{13}$$

where s_k is the sampling standard error of the data moment (clustered on the couple for regression coefficients, binomial for rates), floored at 0.006 so that a moment estimated on 6,727 observations cannot outweigh one estimated on 108. Table 12 lists the moments.

Table 12. Targeted moments: data and model

	Data	Model	t
Husband's exit from work			
Retirement hazard, 50–54	+0.0436	+0.0580	+2.4
Retirement hazard, 55–59	+0.1187	+0.1495	+3.5
Retirement hazard, 60–64	+0.1918	+0.2039	+0.8
Retirement hazard, 65–69	+0.4537	+0.4432	−0.2
Job-loss hazard, 50–69	+0.0168	+0.0152	−0.3
Inactivity hazard, 50–69	+0.0173	+0.0096	−1.3
Re-employment from unemployment	+0.3871	+0.3147	−2.0
Retirement from unemployment	+0.0323	+0.0750	+3.3
Re-employment from inactivity	+0.3154	+0.3038	−0.3

Table 12, continued

	Data	Model	t
Robot exposure			
Exposure $\times\mathbf{1}[\text{age} \geq 55]$ on retirement	+0.0599	+0.0522	−0.5
Exposure $\times\mathbf{1}[\text{age} \geq 55]$ on unemployment	−0.0005	+0.0096	+1.7
Husband's state distribution			
Working, 50–69	+0.6980	+0.7104	+2.1
Unemployed, 50–69	+0.0295	+0.0269	−0.4
Retired, 50–69	+0.2466	+0.2470	+0.1
Inactive, 50–69	+0.0259	+0.0157	−1.7
Retired, 40–44	+0.0022	+0.0022	−0.0
Retired, 45–49	+0.0159	+0.0177	+0.3
Working, 45–49	+0.9455	+0.9236	−3.7
Wife's employment and retirement, by her age			
Employed, 25–34	+0.4849	+0.5115	+3.8
Employed, 35–44	+0.5063	+0.4901	−2.7
Employed, 45–49	+0.4657	+0.4433	−2.5
Employed, 50–54	+0.3396	+0.3674	+3.0
Employed, 55–59	+0.2282	+0.2456	+1.5
Employed, 60–69	+0.0891	+0.1060	+1.6
Retired, 25–34	+0.0010	+0.0003	−0.1
Retired, 35–44	+0.0040	+0.0052	+0.2
Retired, 45–49	+0.0378	+0.0326	−0.9
Retired, 50–54	+0.0838	+0.0879	+0.7
Retired, 55–59	+0.1533	+0.1868	+3.5
Retired, 60–69	+0.4011	+0.3308	−3.8
Wife's entry hazard, by her age			
Entry, 25–34	+0.0963	+0.1080	+1.7
Entry, 35–44	+0.0633	+0.0689	+0.9
Entry, 45–49	+0.0492	+0.0412	−1.3
Entry, 50–54	+0.0274	+0.0267	−0.1
Entry, 55–59	+0.0289	+0.0152	−2.3
Entry, 60–69	+0.0194	+0.0057	−1.9
Wife's exit hazard, by her age and destination			
Exit to non-employment, 25–34	+0.0976	+0.1007	+0.4

Table 12, continued

	Data	Model	t
Exit to non-employment, 35–44	+0.0546	+0.0703	+2.6
Exit to non-employment, 45–49	+0.0523	+0.0591	+1.1
Exit to non-employment, 50–54	+0.0674	+0.0640	−0.4
Exit to non-employment, 55–59	+0.0836	+0.0686	−1.0
Exit to non-employment, 60–69	+0.0405	+0.0650	+1.1
Exit to retirement, 25–34	+0.0017	+0.0001	−0.3
Exit to retirement, 35–44	+0.0011	+0.0019	+0.1
Exit to retirement, 45–49	+0.0041	+0.0145	+1.7
Exit to retirement, 50–54	+0.0293	+0.0360	+1.1
Exit to retirement, 55–59	+0.0900	+0.0686	−1.3
Exit to retirement, 60–69	+0.2838	+0.1707	−2.2
Cross-spouse responses			
Her employment, he retired	−0.0353	−0.0514	−1.0
Her employment, he unemployed	−0.0242	+0.0308	+1.7
Her employment, he inactive	−0.0303	+0.0034	+1.3
His employment, she retired	−0.1352	−0.0906	+1.0
His employment, she non-employed	−0.0357	−0.0297	+0.2
Her retirement on his	+0.0531	+0.0269	−1.7

Notes: 54 moments; criterion 174.0, 39 moments within two standard errors. t = (model − data)/s_k, where s_k is the standard error of the data moment (clustered on the couple for regression coefficients, binomial for rates), floored at 0.006. Simulated on 12 replicates of the observed panel at the set of draws the final refinement minimized over; the criterion is a property of that draw, and Table D1 reports its distribution over 24 independent draws along with the Monte Carlo error in the moments themselves.

The level and the two slopes of the husband's cost of work are identified by the shape of the retirement hazard across the four age bands: the level sets its height at 50–54, θ_1 its growth to 60, and θ_2 the steep rise afterwards. The eligibility distribution is separately identified from the same profile because eligibility and the cost of work enter the hazard multiplicatively: eligibility bounds the hazard from above at every age, while the cost governs the fraction of eligible men who take the option. The retired share below age 50, which is 0.0022 at 40–44, further bounds how early the distribution can begin.

The exposure parameter δ is identified by the difference-in-differences coefficient (4) computed on simulated data, and it is identified separately from the cost of work because a level shift in the cost changes the hazard at every age while δ changes it only where eligibility is positive. The separation and arrival hazards are identified by the corresponding flows.

On the wife's side, the exit hazard from employment identifies the level of her cost of work, its age slope identifies θ_{w1}, and the split of her exits between non-employment and retirement identifies the middle nesting parameter ρ_n. The arrival rate $\mu(a_w)$ and the fixed cost κ are jointly identified by the level of her entry hazard across six age bands together with its responsiveness: rationing bounds entry from above whatever the incentive, while a fixed cost can be overcome by a large enough one. Finally ϕ is identified by the four cross-spouse coefficients of Table 11 and by joint retirement, and the scale parameters by the overall frequency of transitions.

The search starts by cycling over three blocks, each set of parameters against the moments that identify it, and then searches the full parameter space against all fifty-four moments. The block stage is only a starting point, because parameters also move moments outside their own block; the joint-leisure parameter is the extreme case, with nine-tenths of its leverage on moments outside the cross-spouse block. Appendix D gives the implementation.

C. Estimates

Table 13 reports the estimates with standard errors and Table 14 the implied exogenous processes. One parameter sits against a boundary: the share of women who ever become pension-eligible estimates at 0.9914, within one percent of its natural ceiling of one, so the data are best read as saying that essentially every woman in these cohorts eventually becomes eligible. No other parameter is within two percent of a bound, though two sit within five percent of theirs: the age slope of his disutility (0.0151, against a floor of zero) and the share of men who ever become eligible (0.9824).

Standard errors come from the simulated-method-of-moments sandwich. Writing $g(\Theta)$ for the vector of differences between the fifty-four simulated and data moments, the estimator minimizes $g(\Theta)'Wg(\Theta)$ with W the diagonal matrix of inverse squared data-moment standard errors; the variance of the estimator is $(1+1/S)(G'WG)^{-1}$, where G is the Jacobian of the simulated moments in the parameters computed by central differences at fixed simulation draws, $S = 12$ is the number of simulation replicates, and the collapse of the sandwich to this form reflects the same diagonal treatment of the moments that the criterion uses. Three parameters cannot be distinguished from zero: the age slope of his disutility of work before sixty (0.0151, standard error 0.0199), the nesting parameter on his choice margin (0.4205, standard error 0.2916), and the arrival rate of offers out of other non-employment (0.8422, standard error 0.5422). The exposure cost δ, at 0.8375 with a standard error of 0.1283, is not among them. For the location and scale parameters of the eligibility distributions a t-statistic against zero is mechanical, and the table reports standard errors without stars.

The precision of δ barely depends on how the two exposure gradients are weighted, because δ is identified by the whole exposure-by-age profile and loads little on either gradient alone. The check is this. The gradients enter the criterion weighted by couple-clustered standard errors, and treatment varies across occupation cells, where the retirement gradient's standard error is 0.0296 against 0.0170 and the unemployment gradient's 0.0052 against 0.0061. Re-weighting both at the occupation level requires the full sandwich formula, since the collapsed form above assumes $\Omega = W^{-1}$; evaluated at the same estimate it moves the standard error of δ from 0.1283 to 0.1305, a t of 6.42 against 6.53, and re-weighting the retirement gradient alone gives 0.1306. Table 13 reports the same calculation for every parameter.

The husband's cost of continuing to work has a level of 0.8528. It barely rises between 50 and 60, and its slope gains a further 0.1541 per year of age after 60. A one-standard-deviation increase in exposure adds 0.8375 to it at every age, which very nearly doubles the level: a large number, and the reason exposure moves behavior once retirement is available even though it is invisible in pay. The median latent eligibility age is 54.3 with a dispersion of 2.4 years, so that 14 percent of men are eligible at 50, 57 percent at 55 and 97 percent at 65. The estimated share who ever qualify is 0.9824. These estimates are the seniority rule as behavior reveals it.

On the wife's side the estimated arrival rate of employment opportunities falls from 0.2764 at age 30 to 0.0327 at 57. The fixed cost of entry is 3.4488, with a standard error of 0.4589, and its identification does not rest on that local standard error alone: fixing the cost at zero raises the criterion from 174.0 to 1,577 with all other parameters held at their estimates, and to 878 when the offer-arrival intercept, the other parameter operating only on the entry margin, is re-optimized at each value. On this sample the data distinguish a large fixed cost of entry from none. The entry friction the model needs is therefore two-sided (offers arrive rarely, and taking one is costly), although, as Section V.F shows, neither component makes her entry respond much to his job loss; the level of entry and its response are different objects.

The joint-leisure parameter is 0.1183 with a standard error of 0.0516, and the nesting parameters are $\rho_n = 0.1611$ and $\rho = 0.4205$, both inside the unit interval, so neither nest collapses. The second of the two is not distinguishable from zero, which is to say that the data do not require the husband's choice to be noisier than a degenerate nest would make it.

Table 13. Estimated parameters

		Estimate	(s.e.)
Husband: cost of continuing to work			
Level	θ_0	0.8528	(0.1976)
Slope per year of age above 50	θ_1	0.0151	(0.0199)
Additional slope per year above 60	θ_2	0.1541	(0.0449)
Robot exposure, per standard deviation	δ	0.8375	(0.1283)
Husband: institutions and employment risk			
Median latent eligibility age	a_{0h}	54.2826	(0.9227)
Dispersion of eligibility ages	τ_h	2.3503	(0.3515)
Share ever eligible	$\bar{e}_h$	0.9824	(0.0535)
Separation into unemployment	π_U	0.0212	(0.0027)
Separation into inactivity	π_O	0.0111	(0.0036)
Job arrival out of unemployment	λ_U	0.7081	(0.1292)
Job arrival out of inactivity	λ_O	0.8422	(0.5422)
Wife: cost of working and the entry margin			
Level	θ_{w0}	0.3799	(0.0105)
Slope per year of her age above 45	θ_{w1}	0.0220	(0.0037)
Fixed utility cost of entering employment	κ	3.4488	(0.4589)
Opportunity arrival, logit intercept at age 45	m_0	-2.3103	(0.1412)
Opportunity arrival, slope in her age	m_1	-0.0899	(0.0103)
Wife: institutions			
Median latent eligibility age	a_{0w}	55.9516	(0.5812)
Dispersion	τ_w	4.0777	(0.2797)
Share ever eligible	$\bar{e}_w$	0.9914	(0.0643)
Joint preferences and taste shocks			
Value of joint non-work	ϕ	0.1183	(0.0516)
Scale, her work / non-work decision	σ_w	1.3148	(0.1023)
Nesting, her non-work states: $\sigma_n = \rho_n \sigma_w$	ρ_n	0.1611	(0.0487)
Nesting, his decision: $\sigma_h = \rho\, \sigma_n$	ρ	0.4205	(0.2916)

Notes: 23 parameters estimated by simulated method of moments on the 54 moments of Table 12. No element of the budget in Table 10 is estimated. The discount factor is fixed at $\beta = 0.96$. The nesting parameters satisfy $\sigma_h \leq \sigma_n \leq \sigma_w$ by construction, which is the condition for consistency with random utility maximization. Standard errors are from the simulated-method-of-moments sandwich $(1 + 1/S)(G'WG)^{-1}$ with $S = 12$ simulation replicates, W the diagonal inverse-variance weighting of the criterion, and G a central-difference Jacobian at fixed simulation draws; cross-moment covariances of the data moments are not estimated, so the moments are treated as independent, as in the criterion. The two exposure gradients enter W with couple-clustered standard errors, and the couple is not the level at which the treatment varies. Re-clustering them on the occupation cell and evaluating the full sandwich $(1 + 1/S)(G'WG)^{-1}G'W\Omega WG(G'WG)^{-1}$ at the same estimate leaves δ at 0.1305 against 0.1283, a t of 6.42 against 6.53; the parameters that move are the age profile, θ_2 falling to $t = 2.56$ and τ_h to $t = 5.15$. The share of women eventually eligible sits at its natural ceiling of one, where an interior-optimum standard error is not the relevant object.

Table 14. The exogenous processes

	Data	Model
Panel A. Husband's realized flows, ages 50–69		
Work to unemployment	0.0168	0.0152
Work to inactivity	0.0173	0.0096
Unemployment to work	0.3871	0.3147
Inactivity to work	0.3154	0.3038
Panel B. Pension eligibility, $\Pr(a \geq a^*)$		
Husband at age 50	—	0.1367
Husband at age 55	—	0.5656
Husband at age 57	—	0.7472
Husband at age 60	—	0.9031
Husband at age 62	—	0.9469
Husband at age 65	—	0.9722
Wife at age 50	—	0.1869
Wife at age 55	—	0.4381
Wife at age 57	—	0.5591
Wife at age 60	—	0.7234
Wife at age 62	—	0.8081
Wife at age 65	—	0.8942
Panel C. Arrival of employment opportunities to a non-employed wife, $\mu(a_w)$		
At her age 30	0.0963	0.2764
At her age 40	0.0633	0.1346
At her age 47	0.0492	0.0766
At her age 52	0.0274	0.0502
At her age 57	0.0289	0.0327
At her age 62	0.0194	0.0211

Notes: Panel A reports realized flows, which the model is asked to match. They are not the underlying parameters: an arriving offer may be declined and a separated man may retire instead, so the flows differ from π_U, π_O, λ_U and λ_O, which are estimated jointly with the rest and reported in Table 13. Panel B reports the estimated cumulative distribution of the latent eligibility age; ECHP carries no contribution histories, so eligibility is identified from the age profile of exit. The wife's eligibility accrues only in periods in which she is employed. Panel C compares the estimated arrival rate with the raw entry hazard of non-employed wives in the corresponding age band; the arrival rate bounds the entry hazard from above because an arriving opportunity may be declined.

D. Fit

Table 12 reports the fit. Of the 54 targeted moments, 39 lie within two standard errors of the data and 22 within one, and the ones that fit best are the ones the paper rests on. The model reproduces Fact 1 in full, including the design that identifies it. The retirement hazard runs 0.0580, 0.1495, 0.2039, 0.4432 against 0.0436, 0.1187, 0.1918, 0.4537 in the data, with the top band matched to within a quarter of a standard error. The difference-in-differences coefficient of (4) computed on simulated data is 0.0522 against 0.0599, and the corresponding coefficient for unemployment is 0.0096 against −0.0005; neither is imposed. The retired share below 50, the pooled retired share at 50–69 (0.2470 against 0.2466), the job-loss hazard (0.0152 against 0.0168), and re-employment out of inactivity (0.3038 against 0.3154) all sit inside half a standard error. Re-employment out of unemployment is 0.3147 against 0.3871.

The cross-spouse responses that motivate the paper are matched in size and, for the larger ones, in sign. His employment when she is retired is −0.0906 against −0.1352 in the data, and his response to her inactivity, −0.0297 against −0.0357, is within a fifth of a standard error. Her employment when he is retired is −0.0514 against −0.0353, inside one standard error. The one response the model gets wrong in sign is the smallest: her employment when he is unemployed is +0.0308 in the model against −0.0242 in the data, a miss of 1.7 standard errors on a moment whose own interval runs from −0.0874 to +0.0390. Figure 4 shows the two life-cycle profiles at the center of the estimation.

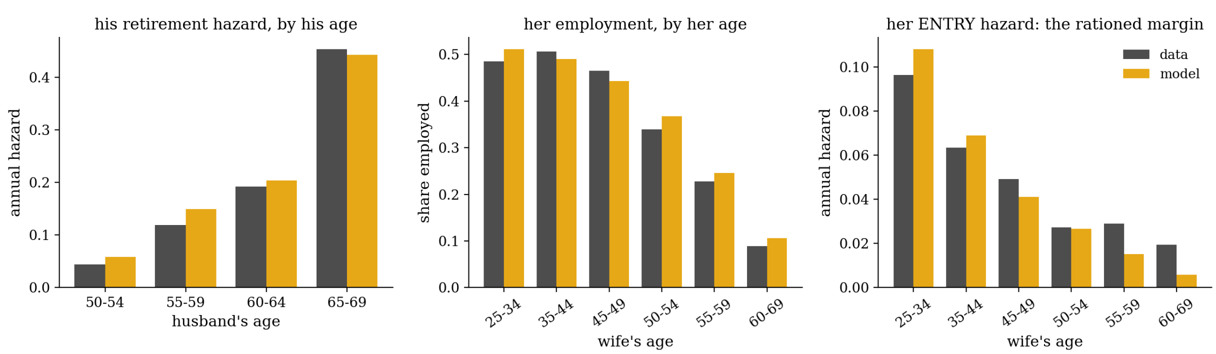

FIGURE 4. MODEL FIT: THE HUSBAND'S EXIT HAZARD AND THE WIFE'S PARTICIPATION MARGINS

Notes: Data and model, from Table 12. The right panel is the entry hazard of wives who were not employed in the previous wave, which the model treats as rationed.

The fifteen moments beyond two standard errors cluster on one margin, the timing of the wife's own retirement. The model retires her too early in her fifties and too late afterwards: her retired share at ages 55–59 is 0.1868 against 0.1533 and at 60–69 0.3308 against 0.4011, so the flow from her employment directly into retirement at the oldest ages comes out at 0.1707 against 0.2838 and joint retirement inherits the shortfall, 0.0269 against 0.0531. On the husband's side the retirement hazard is 3.1 percentage points too high at 55–59, retirement out of unemployment is 0.0750 against 0.0323, and the model keeps slightly too many men at work at 45–49, 0.9236 against 0.9455.

Three objects lie outside the model by construction. Return from retirement is 0.0296 in the data and zero in the model. The involuntary share of retirements and its exposure gradient are zero in the model because the firm-side channel is fixed at zero, as Section IV.A states. The model also has only two agents, so the 0.0364 of non-spousal household income that arrives when he is unemployed is absent from the budget; including it is one of the variants in Section V.F.

E. Moments Not Used in Estimation

Table 15 reports moments that never entered (13).

Table 15. Moments not used in estimation

	Data	Model
Exposure gradient on the flow into inactivity	+0.0032	+0.0035
Her entry hazard on his job loss, wives 25–44	+0.0944	+0.0445
Her entry hazard on his job loss, wives 45–69	+0.0321	+0.0171
Her exit hazard when he retires, wives 45–69	+0.2389	+0.1609
Her exit hazard when he stays at work, wives 45–69	+0.0746	+0.0796
Her participation, he retired	−0.0321	−0.0514
Her participation, he unemployed	+0.0042	+0.0308
His participation, she retired	−0.1320	−0.0922
Her employment rate, he retired	+0.2111	+0.2653
Return to work from retirement	+0.0296	+0.0000
Firm-initiated separations, from which the model abstracts		
Involuntary share among work-to-retirement events	+0.0317	+0.0000
Exposure gradient in that involuntary share	+0.0082	+0.0000

Notes: None of these moments enters the criterion. The first row is the exposure gradient on the flow from work into inactivity at the eligibility threshold, standard error 0.0064; the model's near-zero is a prediction, since exposure enters only the cost of continuing to work and separations into inactivity are exogenous. It stays untargeted because its squared t-statistic is 0.002, so adding it to the criterion would move the criterion of 174.0 by that much. The next two are the added-worker effect at the two ends of the age range, estimated on wives non-employed at $t-1$ with year effects; their standard errors are 0.0456 and 0.0338 respectively. The last two rows are reduced-form facts on 378 matched annual transitions from work into retirement: the share carrying a firm-side marker and its gradient in exposure, whose standard errors are 0.0090 and 0.0103. The model fixes the firm-side separation channel at zero, so its prediction for both is zero by construction; the counterfactual reports an imposed bound.

The criterion includes the level of the wife's entry hazard at six ages and the cross-spouse employment responses at ages 45–69. It does not include the response of her entry hazard to her husband's job loss. At wife ages 25–44 the model predicts +0.0445 against +0.0944 in the data, and at 45–69 +0.0171 against +0.0321. The prediction is positive at both ends, with a simulation standard deviation of 0.0104 at the younger ages, so its sign is a property of the estimated model and does not depend on the reporting draw; the model understates only the magnitude, by about half, and in both bands the miss is within one standard error of the data moment. The understatement has one source: the forces that cancel on the husband's unemployment cancel at every age, so nothing in the model switches on when the wife is young.

One untargeted prediction succeeds cleanly. Her exit hazard from employment when her husband stays at work is 0.0796 against 0.0746 in the data, and the model correctly puts it far below her exit hazard when he retires, the flow behind joint retirement that the criterion sees only as a stock, though it understates that second hazard, 0.1609 against 0.2389.

The state distribution across all nine age bands from 25–29 to 65–69 divides between the two spouses. On the wife's side, where none of the twenty-seven shares enters the criterion, twenty-two lie within two standard errors of the data and none beyond four. On the husband's side, of thirty-six shares, of which the criterion sees only three, sixteen lie within two standard errors and nine beyond four, and the model is too retired in every band from fifty onward, 0.3580 against 0.2485 at ages 55–59 and 0.5713 against 0.4496 at 60–64. The pooled retired share at 50–69, which the criterion does target, nevertheless matches, 0.2470 against 0.2466. The sample restriction reconciles the two. Requiring the husband to be seen at work inside the window removes more old couples from the simulated panel than from the data, and the composition shift offsets the over-retirement within each band. Appendix E reports the comparison, and it is why the fit here is assessed on the sample basis while Section VI reports the counterfactual on the population basis. Appendix F re-estimates the model on the full population, including the self-employed, and the fit deteriorates on the moments that identify the retirement margin.

F. What Generates the Absence of Insurance

Fact 4 left the question of why the wife does not insure her husband's job loss when her earnings would replace most of what he loses. Table 16 re-solves and re-simulates the model with one element changed at a time, holding every estimated parameter at its Table 13 value, and applies the estimators of Table 11 to the result.

Table 16. The sources of the missing added-worker effect

	Her empl. when he is		Entry resp.	Wives 50–54	
	Retired	Unemp.	to his job loss	Entry haz.	Employed
Data	−0.0353	−0.0242	+0.0321	0.0274	0.3396
Estimated model	−0.0514	+0.0308	+0.0171	0.0267	0.3674
No joint leisure ($\phi = 0$)	−0.0079	+0.0396	+0.0254	0.0292	0.3920
Her offer rate held at its age-30 value	−0.0914	+0.0331	+0.0124	0.0939	0.4079
No entry cost ($\kappa = 0$)	−0.0665	+0.0363	+0.0168	0.0439	0.3399
Near-deterministic ($\sigma \times 0.2$)	−0.0446	+0.0023	+0.0201	0.0018	0.4831
No frictions, no ϕ, near-determ.	+0.0094	+0.0649	+0.0541	0.2428	0.5540

Notes: Each row re-solves and re-simulates the model with one element changed and every estimated parameter held at its Table 13 value, then applies the estimators of Table 11 to the simulated panel. The third row freezes her offer arrival rate at its age-30 value; it does not remove the friction. The last column is the level of female employment each variant implies, against 0.3396 in the data.

The taste for joint leisure drives the wife's withdrawal, and it applies equally to both of her husband's exits. Setting $\phi = 0$ moves her employment response to his retirement from −0.0514 to −0.0079 and her response to his unemployment from +0.0308 to +0.0396. Because ϕ enters the two exits identically, it cannot generate the difference of more than 8 percentage points between the two responses in the estimated model. The difference comes from the budget. When he retires the household keeps 0.8156 of his earnings, the insurance motive is weak, and joint leisure dominates; when he is unemployed it keeps 0.1436, the insurance motive is strong, and the two roughly offset. The model's asymmetry therefore follows from Fact 3's replacement rates, which are measured, and this is the paper's central result on the household. The data are consistent with that ordering but do not establish it: the two data coefficients are −0.0353 and −0.0242, and the second is estimated on a cell too thin to sign the difference.

Figure 5 displays the decomposition against the confidence intervals of the two data coefficients.

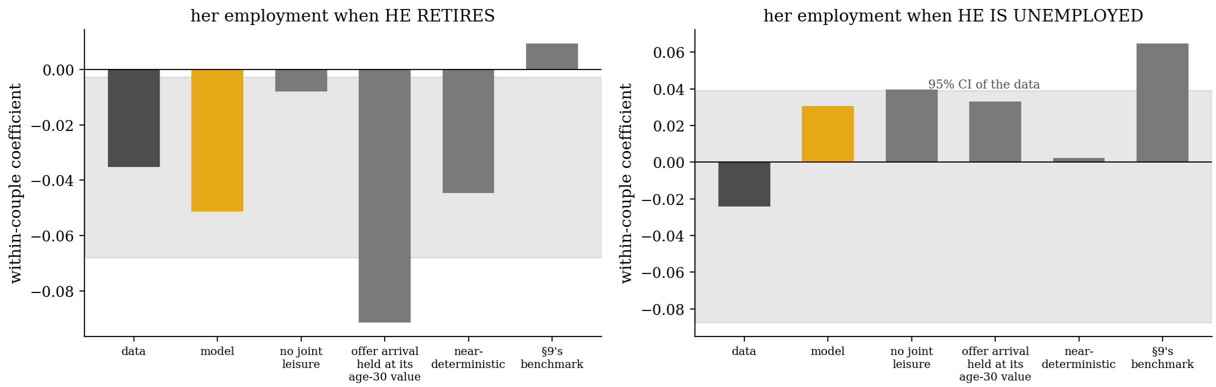

Figure 5. The wife's response to her husband's two exits, under alternative model variants

Notes: Bars are the within-couple coefficients of Table 11 computed on simulated panels. Shaded bands are the 95 percent confidence intervals of the corresponding data coefficients.

Rationing the wife's entry margin does little to suppress the added-worker effect. Freezing her offer arrival rate at its age-30 value raises her entry hazard at ages 50–54 from 0.0267 to 0.0939, more than threefold, but moves the response of that hazard to her husband's job loss only from +0.0171 to +0.0124 and her employment response to his unemployment from +0.0308 to +0.0331. Removing the fixed cost of entry raises the entry hazard to 0.0439 and leaves the entry response at +0.0168. Both components of the entry friction set the level of her entry hazard, which the model matches at every age, and neither affects its response to his shock. The reform in Section VI confirms this on the policy itself.

A household with no frictions is rejected by the same data. Removing the frictions, the joint-leisure term, and most of the choice noise together raises her employment response to his unemployment to +0.0649, outside the interval the data imply, and lifts the employment rate of wives aged 50–54 to 0.5540 against 0.3396 in the data. The variant fails on the level as well as on the response. This is why the paper states the size of the added-worker puzzle with the raw budget arithmetic of Fact 4 and the mechanism with the estimated model, and does not use the model to measure how large the response of a frictionless household would be.

Table 17 varies the measurement choices on which the household result rests.

Table 17. Sensitivity to the measurement of the budget

	Her empl. when he is		His empl. when	Joint
	Retired	Unemployed	she is retired	retirement
Data	−0.0353	−0.0242	−0.1352	+0.0531
Estimated model	−0.0514	+0.0308	−0.0906	+0.0269
b_R from residual measure[a]	−0.0489	+0.0379	−0.0423	+0.0103
b_U from residual measure	−0.0519	+0.0359	−0.0900	+0.0274
Non-spousal income when he is unemployed	−0.0510	+0.0288	−0.0916	+0.0269
ϕ on inactivity and retirement[a]	−0.0547	+0.0389	−0.0970	+0.0291

Notes: The residual measure computes replacement from the change in household income net of the couple's own earnings, and gives 1.0431 on retirement and 0.1107 on unemployment against 0.8156 and 0.1436 on the personal-income measure used throughout. The fourth row adds the within-couple increase in household income belonging to neither spouse when the husband is unemployed, +0.0364 of his earnings when working. Where a change would displace the block of parameters it identifies, that block is re-estimated on its own moments, which is the case for the 2 rows marked [a]; the others hold every parameter fixed. Every row is computed on the same simulated panel as the model fit reported in Table 12 (12 replicates of the observed panel at a single seed), so that variants differ from the estimated model only through the assumption each one changes.

Replacing the personal-income replacement rate with the residual measure raises b_R from 0.8156 to 1.0431, which makes the husband's retirement nearly costless and should therefore widen the asymmetry between his two exits. With his exit parameters re-estimated on his own moments it does so, but slightly: her response to his retirement moves from -0.0514 to -0.0489 and her response to his unemployment from $+0.0308$ to $+0.0379$. Taking the unemployment replacement rate from the same residual measure moves both by less. Adding the non-spousal household income that arrives on his unemployment, which the two-agent model omits, changes nothing material. Restricting joint leisure to his non-participation states, so that his unemployment carries no complementarity at all, requires the complementarity parameter to be re-estimated and leaves both responses close to the estimated model. Across all variants her response to his retirement lies between -0.0547 and -0.0489 and her response to his unemployment between $+0.0288$ and $+0.0389$, against a data estimate of -0.0242; every variant leaves her response to his unemployment inside the 95 percent interval of the data coefficient, in the case of the joint-leisure variant only marginally.

One competing explanation lies outside the model. A household with savings can meet a spell of his unemployment by running them down instead of sending the wife to work, the mechanism Blundell et al. (2016) place at the center of family self-insurance, and this model has no savings. The two explanations make different predictions: self-insurance says the wife's response should be muted where wealth is high and present where it is low, while the replacement-rate account says it should be muted wherever the exit is well insured, which is everywhere in retirement and nowhere in unemployment. The survey measures neither wealth nor consumption, so it cannot run that test, and the two accounts coexist on this evidence. The data do establish the pattern the model was built to explain: the wife responds to his retirement and barely to his job loss, in line with what the two exits pay. For the welfare results the distinction does not matter, as Section VI shows.

VI. Raising the Pension Eligibility Age

The counterfactual uses the estimated model to shift the distribution of the latent eligibility age to the right by one to five years, the shift Amato and Dini applied to the date at which a given man could stop working. Because eligibility is a persistent characteristic drawn before the reform, the experiment does not change who is more or less attached to work. It changes when the option arrives. The experiment compares steady states, holds prices and hazards fixed, and abstracts from the phased, partly grandfathered transition that was legislated.

The quantities reported here are population objects, computed over the full simulated life cycle instead of the estimation sample. The two are not the same. The sample requires the husband to be observed at work, so it over-represents late workers at exactly the ages the reform acts on: at ages 60–64 the estimation sample puts 45 percent of husbands in retirement, the model's counterpart of that sample 57 percent, and the model's population 90 percent. Appendix E reports the comparison. The section reports, in order, where the removed retirement goes, who bears its removal, how the reform interacts with the shock itself, the household's response, the fiscal saving and the welfare cost, what consumption smoothing could change, and an actuarial alternative.

A. Where the Men Go

Table 18 and Figure 6 report the reform. A five-year shift lowers the retired share at ages 60–64 from 0.8991 to 0.6977. Of each percentage point of retirement removed at ages 50–69, 0.52 returns to work, 0.29 becomes unemployment, and 0.19 becomes inactivity, a split that is stable across age bands.

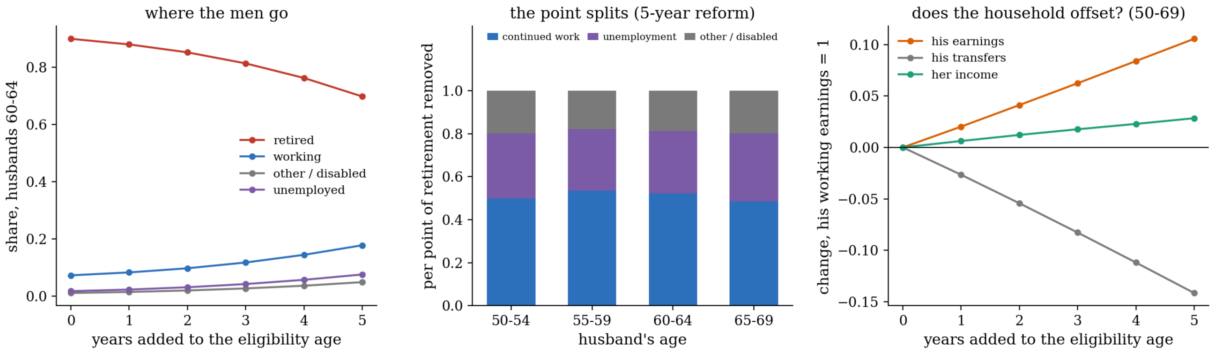

FIGURE 6. RAISING THE ELIGIBILITY AGE

Notes: Left panel, husbands aged 60–64. Middle panel, the destination of one percentage point of removed retirement under a five-year reform, by the husband's age. Right panel, the change in each component of household income at ages 50–69 relative to the baseline, in units of the husband's earnings when working.

Table 18. Raising the pension eligibility age

Panel A. Husbands aged 60–64, and the household at 50–69

Years added	His state at 60–64				Household, 50–69		
to eligibility	Work	Unemp.	Inact.	Retired	She empl.	Cons.	Transf.
0	0.0726	0.0174	0.0109	0.8991	0.2529	1.3028	0.6138
1	0.0829	0.0230	0.0146	0.8795	0.2591	1.3027	0.5906
2	0.0974	0.0310	0.0198	0.8518	0.2650	1.3018	0.5662
3	0.1176	0.0423	0.0269	0.8132	0.2703	1.3001	0.5410
4	0.1443	0.0570	0.0364	0.7624	0.2756	1.2977	0.5149
5	0.1777	0.0757	0.0489	0.6977	0.2813	1.2952	0.4883

Panel B. Where a point of retirement goes, 5-year reform

His age	Δ retired	To work	To unemp.	To inact.	Her empl.
50–54	−0.1804	0.4990	0.3015	0.1995	+0.2000
55–59	−0.3922	0.5347	0.2868	0.1785	+0.1131
60–64	−0.2014	0.5217	0.2894	0.1888	+0.1285
65–69	−0.0363	0.4850	0.3158	0.1991	+0.1987

Variants of the 5-year reform, ages 50–69

	Per point of retirement removed			Transfers	CEV, %	
	To work	To unemp.	To inact.	% change	From 25	From 50
Panel C. Bounds on the inactivity margin						
Inactivity a shock (lower bound)	0.5213	0.2920	0.1867	−20.4	−3.30	−8.90
Inactivity choosable (choice reading)	0.2095	0.0065	0.7840	−13.9	−1.50	−4.11
Panel D. An imposed firm-side separation channel						
No firm-side channel, as estimated	0.5213	0.2920	0.1867	−20.4	−3.30	−8.90
Firm-side channel imposed	0.4422	0.3947	0.1631	−20.3	−2.99	−8.08

Notes: The reform shifts the distribution of the latent eligibility age to the right by the stated number of years. Panel A: consumption and transfers are per couple-year in units of the husband's earnings when working; transfers are public transfers (old-age pension plus unemployment benefit plus sickness and invalidity benefit), whose state-specific values are taken from the data. Panel B: the three destination columns sum to one. Panel C: the estimated model treats entry into inactivity as an exogenous shock, the shock reading, which is the lower bound on diversion into it; the second row makes inactivity choosable at the same parameters, the choice reading, which is the upper bound. The inactivity-state transfer is measured on 48 person-years. Panel D imposes a firm-side release channel at the involuntary share observed in the data, 0.0125 per period with every released ineligible man routed to unemployment; the channel is imposed (freely estimated it returns zero), and 86 percent of the reform's employment gain at 50–69 survives it. All quantities are population objects, computed on the full simulated life cycle.

The split rests on one modeling choice: whether entry into long-term inactivity is a shock or a choice. Call these the shock and the choice readings; the estimated model is the shock reading. The data lean its way. Inactivity in this sample is largely a health state, with 0.121 of inactive husbands aged 50–69 receiving a sickness or invalidity benefit against 0.060 of the retired and 0.022 of those at work; exposure does not move the flow into it, a null of +0.0032 (standard error 0.0064) that the model predicts without being asked (Table 15); and the choice reading overshoots, placing 0.0728 of men aged 50–69 in inactivity against 0.0259 in the data. The sample still cannot decide the question outright, because a man who leaves work and never returns is largely invisible to a panel that requires him to be seen working. Panel C of Table 18 therefore reports the reform under both readings: under the choice reading the split becomes 0.21 work, 0.01 unemployment, and 0.78 inactivity, and every headline quantity below is reported as the interval the two readings define.

Panel D bounds the model's other restriction the same way. A firm-side release channel imposed at the involuntary share the data record, with every released ineligible man routed to unemployment, preserves 86 percent of the reform's employment gain and moves the destination mix toward unemployment, from 0.29 to 0.39 per point removed; freely estimated, the channel returns zero.

The actual reforms offer a quasi-experimental benchmark. Ardito (2021) follows the Amato arm of the episode, the old-age age rising from 60 to 65 for private-sector male employees, through administrative records: among treated cohorts pension claiming falls by 44 percentage points, employment rises by about 9, disability take-up rises by 4, and the largest destination is inactivity with no benefit at all. Her employment share of the removed claiming, roughly a fifth, lies at the choice-reading end of the interval reported here. The margins differ, since her threshold binds at 60–65 and men who stop without claiming are waiting out a delay in a state this model does not carry, but the reduced-form pattern, an employment response well below one for one with substitution into disability and unemployment benefits, is the counterfactual's. Her heterogeneity anticipates the next subsection: the delay, the employment rise, and the diversion into inactivity concentrate among workers with lower occupational grades and pay, with little response among the highly skilled. Bottazzi et al. (2006) show the reforms were understood as they happened: expected retirement ages rose across every employment group.

B. *Incidence by Exposure*

Panel A of Table 19 and Figure 7 report who bears the removal. Exposure raises the value of stopping, so exposed men are the ones using the early-retirement option, and withdrawing it takes most from them: at 55–59 the five-year reform removes 45 percentage points of retirement among above-median-exposure men against 33 among the rest, while by 65–69, where nearly everyone is eligible in both groups, the two are indistinguishable. This gap is the same to two decimals under either reading, -0.4455 and -0.3323 under the shock reading against -0.4453 and -0.3303 under the choice reading, which is why the paper leads with it: the readings govern only where the removed retirement goes.

Where the exposed men go does differ. At 55–59 their inactivity share rises by 10 points against 3 for the unexposed under the shock reading, because continued work in an automating occupation costs them most; under the choice reading, inactivity absorbs almost all of their removed retirement.

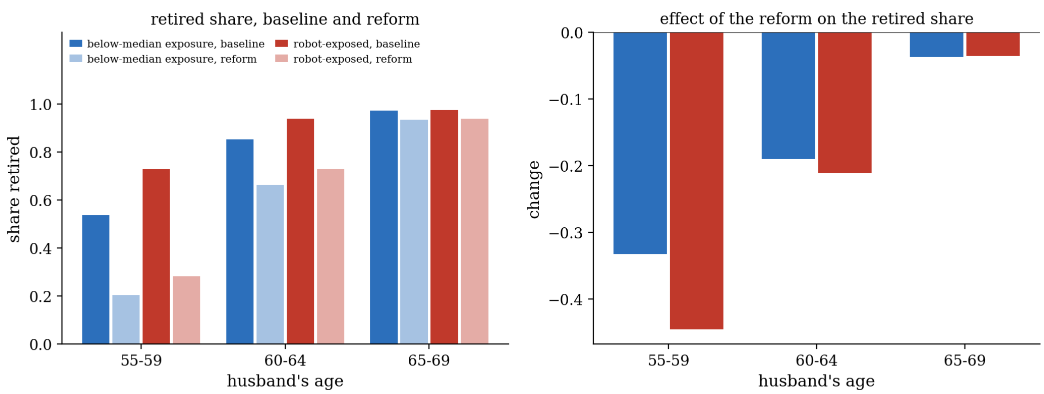

Figure 7. The reform's incidence by occupational robot exposure

Notes: Couples split at the median of the standardized exposure score. In the left panel the solid bar is the baseline retired share and the pale bar the share under the five-year reform.

C. *Automation and the Reform*

Table 20 turns the counterfactual on the shock itself. Setting $\delta = 0$ gives every occupation the mean exposure shift and re-solves. The exposure score is standardized to mean zero and a shift common to every occupation is absorbed in the level of the cost of work, so the design identifies exposure differences: the experiment equalizes the shock across occupations and cannot remove it. Equalization raises early retirement, moving the retired share at 55–59 from 0.6388 to 0.7293 and transfers at 50–69 from 0.6138 to

Table 19. Incidence, welfare and the budget

Panel A. Retired share by exposure, 5-year reform

	Below-median exposure		Above-median exposure	
His age	Baseline	Change	Baseline	Change
55–59	0.5378	−0.3323	0.7288	−0.4455
60–64	0.8531	−0.1902	0.9400	−0.2113
65–69	0.9722	−0.0370	0.9756	−0.0357

Panel B. Consumption-equivalent variation, percent

Years added	All couples		Below-median	Above-median	Public transfers
to eligibility	From 25	From 50	exposure	exposure	saved
0	+0.00	+0.00	+0.00	+0.00	0.0000
1	−0.70	−1.93	−0.23	−1.12	0.0232
2	−1.38	−3.77	−0.47	−2.17	0.0477
3	−2.03	−5.54	−0.73	−3.17	0.0729
4	−2.67	−7.24	−1.01	−4.13	0.0989
5	−3.30	−8.90	−1.32	−5.04	0.1255
the same reform under the choice reading of the inactivity margin					
5	−1.50	−4.11	−0.90	−2.03	

Panel C. Household and government on one discounted scale, PDV at age 25

Years added to eligibility	Transfers saved	Revenue on induced earnings	Government total gain	Household loss	Household loss per unit of government gain
1	0.1483	0.0610	0.2093	0.2484	1.1868
2	0.2928	0.1193	0.4120	0.4879	1.1841
3	0.4329	0.1739	0.6067	0.7203	1.1872
4	0.5698	0.2261	0.7958	0.9474	1.1905
5	0.7023	0.2760	0.9783	1.1713	1.1972

Notes: Panel A reports the retired share at baseline and its change under the five-year reform, with couples split at the median of the observed exposure scores. The quantity incidence in Panel A is, at these parameters, left alone by the inactivity margin: the five-year change in the retired share at ages 55–59 is −0.4455 for above-median and −0.3323 for below-median exposure under the estimated model, against −0.4453 and −0.3303 under the choice reading of Table 18 Panel C. Panel B: consumption-equivalent variation computed from the ex-ante value function at the husband's age 25, and again over the ages that bear the reform from 50; a negative number is the permanent proportional reduction in consumption that leaves the household as well off as the reform does. Transfers saved in Panel B are a flow, per couple-year at ages 50–69 in units of the husband's earnings when working. The second block of Panel B repeats the five-year row under the choice reading, the upper bound on inactivity diversion: the welfare incidence is not invariant to that bound, and the exposed-to-unexposed ratio falls from 3.83 to 2.25 across the two readings. Panel C puts every quantity on one present-value scale at age 25 at $\beta = 0.96$, including the terminal perpetuity the value function assumes; revenue applies the flat rate that finances the baseline transfer bill in present value, 0.3635, an accounting rate that never enters the household's problem. There is no general equilibrium, no behavioral response to the rate, and no mortality; the experiment compares steady states and abstracts from the phased transition that was legislated.

0.6601 per couple-year, because more occupations sit below the mean than above it, and taking away their relief raises the cost of work for the majority.

Run inside the equalized economy, the same five-year reform removes more retirement, saves more, and costs more: a transfer saving of 0.1364 against 0.1255, and a consumption-equivalent cost of 3.81 against 3.30 percent, the permanent consumption cut equivalent to the reform's welfare loss, whose construction Section VI.E reports. The aggregates therefore move modestly with the shock's dispersion. The interaction of automation with pension policy is distributional: exposure decides who uses the insured exit and who pays for its removal, and leaves the reform's totals close to unchanged.

The incidence result invites the design in the last row of Table 20: the five-year shift applied only to below-median-exposure men, leaving the exposed their original eligibility. The carve-out keeps 0.42 of the uniform reform's transfer saving at 0.19 of its welfare cost, and 0.81 of the retirement it removes returns to work, because the men it touches are the ones for whom continued work is cheap. The exempted half stays exactly at the baseline, so its cost is zero and the whole burden falls on the unexposed, at their uniform-reform cost of 1.32 percent. The actuarial design of Section VI.G keeps nearly all of the saving at two thirds of the cost; the carve-out buys less revenue at a far lower price per unit, and places the whole burden on the men the reform costs least.

Table 20. Automation and the reform

Panel A. The baseline economy, with exposure dispersed and equalized			
	Retired share, 55–59	Retired share, 50–69	Transfers, 50–69
Estimated economy	0.6388	0.6827	0.6138
Exposure equalized ($\delta = 0$)	0.7293	0.7326	0.6601
Panel B. The five-year reform, in each economy and designed around exposure			
	Transfer saving	CEV from 25, %	Share to work
Estimated economy	0.1255	−3.30	0.5213
Exposure equalized ($\delta = 0$)	0.1364	−3.81	0.4176
Exempting above-median exposure	0.0531	−0.62	0.8055

Notes: All rows re-solve or re-simulate at the estimated parameters. In Panel A, setting $\delta = 0$ gives every occupation the mean exposure shift: the score is standardized to mean zero and a shift common to all occupations is absorbed in the level of the cost of work, so the design identifies exposure differences, and the experiment equalizes the shock across occupations without removing it. Transfers are public transfers per couple-year at ages 50–69 in units of the husband's earnings when working. In Panel B the first two rows apply the five-year shift inside each economy, and the third applies it only to below-median-exposure men, leaving the exposed their original eligibility; the exempted men are left exactly at the baseline, so their consumption-equivalent cost is zero and the pooled cost is borne by the unexposed (−1.32 percent). The carve-out keeps 0.42 of the uniform reform's transfer saving at 0.19 of its welfare cost.

D. The Household's Response

The household absorbs the closure without offsetting it, as Fact 4 predicts. At ages 50–69 the five-year reform raises the husband's earnings by 0.1056 and lowers his transfers by 0.1416, while the wife's income rises by 0.0284 (Figure 6, right panel). Her income is a tenth of the total adjustment in household resources.

Her employment rises by 2.8 percentage points under the reform, and the rise is concentrated where the incidence is: 4.3 points in above-median-exposure households against 1.1 among the rest, so the wife's margin opens most where the household loses most. Table 21 decomposes the response. Two forces point the same way. The reform removes the states in which joint non-work is available, and switching that complementarity off accounts for 1.2 of the 2.8 points. The remainder is a small insurance response to a small resource loss: the pension replaces 82 percent of his income, but the men the reform diverts into unemployment and inactivity are replaced at 14 and 40 percent, so his net income falls slightly, by 0.0360 of his earnings when working. An income effect operating through a richer household would predict the opposite. The household ends the reform marginally poorer.

Table 21. The wife's employment response to the reform

	Offer rate	Her employment, 50–69 Baseline	Reform	Response
Estimated model (rationed entry)	—	0.2529	0.2813	+0.0284
Her offer rate held at its age-30 value	0.2764	0.3113	0.3342	+0.0229
An offer every period ($\mu = 1$)	1.0000	0.3474	0.3644	+0.0170
Joint leisure switched off ($\phi = 0$)	—	0.3009	0.3170	+0.0160

Notes: Each row re-solves and re-simulates the baseline and the 5-year reform with one element changed and every other estimated parameter held at its Table 13 value. The second row freezes her offer arrival rate at its age-30 value; the third removes the rationing altogether. The friction caps the level of her employment and leaves her response to him essentially unchanged. The reform changes his net income at 50–69 by −0.0360 of his earnings when working, so the household becomes slightly poorer and the insurance motive and the loss of joint non-work push her employment in the same direction.

Opening her entry margin does not overturn any of this. With her offer arrival rate held at its age-30 value her response is 2.3 points, and with an offer every period it is 1.7 points on a much higher base: the friction caps her level and has little effect on her response to him. This is a prediction of the model. The data contain no pension reform inside the observation window against which to check it.

E. Fiscal Saving and Welfare

Public transfers per couple-year at ages 50–69 fall from 0.6138 to 0.4883 under the five-year reform, a reduction of 20.4 percent under the shock reading and of 13.9 percent under the choice reading. The saving is close to linear in the size of the reform, because the eligibility distribution is close to uniform over the relevant range. Applying the reform to both spouses' eligibility raises the saving and the welfare cost together, with little change in the reallocation of his labor supply.

Panel B of Table 19 reports the household side as a consumption-equivalent variation: the permanent proportional cut in consumption that would leave the household as well off as the reform does, computed from the ex-ante value function at the husband's age 25. The five-year reform costs the average household 3.30 percent of permanent consumption under the shock reading and 1.50 percent under the choice reading; expressed over the ages that actually bear the reform, from 50, the same numbers are 8.90 and 4.11 percent, larger because ages 50 and above carry only a fraction of the age-25 present value. The incidence by exposure is the paper's central quantity: 5.04 percent for exposed households against 1.32 for unexposed ones under the shock reading, and 2.03 against 0.90 under the choice reading, a ratio between two and four whose ordering survives either reading. A uniform delay therefore places most of its burden on the workers whose occupations are being automated.

The fiscal and household sides can be put on one scale. Panel C of Table 19 discounts everything to the husband's age 25 at $\beta = 0.96$, including the terminal perpetuity the value function itself assumes, and adds a revenue side: continued earnings are taxed at the flat rate whose present value finances the baseline transfer bill, 0.3635. On that scale the five-year reform saves the government 0.7023 units of transfers and raises 0.2760 units of revenue, a total gain of 0.9783, against which the household would give up 1.1713 units of consumption to avoid it, or 1.20 units lost per unit the government gains. The comparison is bookkeeping. The rate is an accounting construction that never enters the household's problem; there is no general equilibrium, no behavioral response to the rate, and no mortality.

One caveat on this number remains unbounded. Prices are fixed while the reform is large: husbands' employment at ages 50–69 rises from 0.2141 to 0.3197, an increase of about a half in the stock of older male labor, and a general-equilibrium wage response to a supply increase of that size would reduce both the employment gain and the fiscal saving reported here. That pushes the welfare cost toward an upper bound.

F. *What Consumption Smoothing Could Change*

The household here cannot save, while the structural retirement literature gives households a savings margin for exactly this calculation (French, 2005; French and Jones, 2011; Blundell et al., 2016). Table 22 bounds what the restriction costs by granting *perfect* consumption smoothing, free borrowing and lending at the model's own discount rate, so the bound holds for any asset structure, any interest rate, any borrowing limit, and any initial wealth distribution. The answer is that little changes: the cost of the reform moves from 3.30 to 3.16 percent, 4.3 percent of itself, and the exposed-to-unexposed ratio from 3.83 to 3.82.

The bound is this tight because the reform barely changes the present value of consumption, by +0.046 percent: a man returned to work earns 1.0 where the pension paid

0.8156, so the household works longer without becoming much poorer over the life cycle. The cost is labor-supply disutility. Of the 3.30 percent, consumption accounts for 0.099 percentage points; his cost of continuing, $\theta_h(a)+\delta z$, accounts for 2.150, hers for 0.488, and the joint leisure the reform removes for 0.285, each part expressed as its own consumption equivalent. Her entry cost contributes 0.003 with the opposite sign, and the parts sum to 3.02 of the 3.30, the remainder being the surplus of the taste shocks, which the bound holds fixed along with the work costs. A bond smooths a consumption path. It does not make working less unpleasant, and the reform works by making men work.

Table 22. What consumption smoothing could change in the welfare cost

	Consumption-equivalent variation, %		Change, %
Five-year reform	As estimated	Under perfect smoothing	of itself
All couples	−3.30	−3.16	4.3
Below-median exposure	−1.32	−1.27	3.8
Above-median exposure	−5.04	−4.84	4.0

Notes: The estimated model is hand to mouth. This table bounds what that restriction costs the welfare number, for *any* asset structure, by granting perfect consumption smoothing – free borrowing and lending at the model's own discount rate, which no finite asset can improve on. The lifetime value change behind the consumption-equivalent variation is decomposed along the simulated paths into the flow terms of the payoff, and the consumption term is then replaced by the change in the present value of consumption, the object a household with free borrowing and lending cares about; the remaining terms are held at their estimated values, because a bond smooths a consumption path and does not make working less unpleasant. The present value of consumption changes by +0.046 percent under the reform, since a man returned to work earns 1.0 where the pension paid 0.8156. The exposed-to-unexposed ratio moves from 3.83 to 3.82. The bound leaves behavior uncovered: assets would also change when men retire and how much the wife works, a question for a re-estimation beyond this accounting exercise, and Section V.F states the consequence for the mechanism.

The bound does not cover behavior. A household that could save would also retire at a different date, and the wife would work a different amount. Establishing that requires a re-estimation, beyond this accounting exercise; Section V.F gives the consequence for the mechanism.

G. An Actuarial Alternative

The counterfactual so far holds the pension level fixed while moving its date, as Amato and Dini did to the seniority route. Dini's notional defined-contribution formula also raised benefits with the age at which they are claimed, and that is a different instrument. Table 23 reports it: early claiming remains available, and the public pension component is reduced by a fixed rate for each year it is claimed before the shifted eligibility age, with the transfer the household receives moving one for one with the pension so that nothing is fiscally free.

Table 23. An actuarial pension design

	His state at 60–64		CEV	Government, PDV at 25	
5-year reform	Retired	Working	%	Transfers saved	Total gain
Pure shift of the eligibility age	0.6977	0.1777	−3.30	0.7023	0.9783
Early door open, pension cut 4% per year	0.8231	0.1385	−2.32	0.8027	0.9678
Early door open, pension cut 8% per year	0.7081	0.1746	−3.29	0.7343	1.0062

Notes: The retired share at 60–64 without any reform is 0.8991. The actuarial variants leave the early door open and reduce the public pension component by the stated rate for each year it is claimed before the shifted eligibility age; the transfer the household receives moves one for one with the pension, so nothing is fiscally free. Implementation is a commitment menu: the model is solved for each waiting time from zero to 5 years and each couple takes the one that maximizes its ex-ante value. A man claiming later than his chosen floor still pays the floor's penalty, so the welfare cost is an upper bound and the fiscal saving is if anything overstated; the mean claim delay past the floor is 0.4551 years and 0.9152 of couples claim within one year of it. At 4 percent, 0.5289 of couples retire at the original age at a reduced pension and 0.3499 wait the full 5 years. Present values are at the husband's age 25 at $\beta = 0.96$, on the same scale as Table 19.

At four percent per year the design delivers the same fiscal gain at a lower welfare cost inside the model. Just over half of couples still retire at the original age and accept a reduced pension, and the retired share at ages 60–64 recovers to 0.8231 against 0.6977 under the pure shift. The government's present-value gain is essentially unchanged, 0.9678 against 0.9783, while the household's cost falls from 3.30 to 2.32 percent of permanent consumption. The actuarial design converts about a third of the welfare cost into a choice without giving up the budget. At eight percent per year the penalty is steep enough that 88 percent of couples wait the full five years, and the design collapses back onto the pure shift: the retired share is 0.7081 and the cost 3.29 percent. At four percent the choice also sorts on exposure: every above-median-exposure couple pays the penalty and keeps the original age, no below-median couple does, and the relief goes where the burden was, the exposed cost falling from 5.04 to 3.23 percent while the unexposed remain near their pure-shift cost, 1.29 against 1.32 percent.

The implementation is a commitment menu: the model is solved for each waiting time from zero to five years and each couple takes the one that maximizes its ex-ante value, so a man claiming later than his chosen floor still pays the floor's penalty. The welfare cost is therefore an upper bound and the fiscal saving is if anything overstated. The approximation is tight: the mean claim delay past the chosen floor is 0.46 years and 92 percent of couples claim within one year of it.

VII. Conclusion

This paper asks where an older worker goes when automation reaches his occupation, and what the answer costs his household, in a country whose pension system offered an unusually generous way out. Using Italian panel data that follow both spouses through the years in which robots diffused and the seniority pension was still open, I document five facts and build a life-cycle model of the couple around them. Exposure to robots raises the flow of older men into retirement by about six percentage points per standard deviation, and only there: it moves neither unemployment nor inactivity nor pay, it appears only once the seniority pension becomes available, and it is absent among the self-employed, whose fund the rule did not cover, and among wives, who rarely held the contribution years it required. The household absorbs these exits without offsetting them. Although the pension replaced most of a husband's income and unemployment benefit almost none of it, neither spouse goes to work when the other stops, and couples retire together. The estimated model reproduces this pattern with one taste for shared leisure set against the two measured replacement rates: when he retires the household loses little and the taste for joint leisure wins, so the wife withdraws; when he loses his job the household loses most of his income, the insurance motive offsets the taste, and she barely moves.

The findings bear on how pension reform interacts with automation. Because exposed men were the ones using the early exit, delaying eligibility by five years, as the Amato and Dini reforms did, takes more retirement from them, 45 percentage points at ages 55–59 against 33 from the rest, and costs their households two to four times more. The burden stays on the husband's side of the household: the wife's margin opens only a little, and most where the household loses most. Where the removed retirement goes is less certain, between one half and one fifth back to work depending on whether long-term inactivity is a shock or a choice, but the pattern by exposure is clear in either case, with the unexposed returning to work and the exposed diverting toward inactivity. The same model shows that this heterogeneity is where automation enters at all: giving every occupation the same exposure leaves the reform's totals nearly unchanged, so the shock decides who pays rather than how much is saved. Two designs act on that distribution. Allowing early claiming at a reduced pension raises nearly the same revenue at two thirds of the household cost, and the men who take it up are almost entirely the exposed; exempting the most exposed from the delay lowers the cost to a fifth of the uniform reform's while giving up more than half of its revenue. A quasi-experimental evaluation of the Amato reform finds the same pattern in the actual data: employment rising well below one for one with claiming, substitution into disability and unemployment benefits, and the response concentrated among manual and low-paid workers.

Reading the data through a model is what makes these statements possible. The reduced-form gradient identifies a sign and a route more firmly than a magnitude, and it says nothing about welfare; the model converts it into a cost of continuing work, carries that cost through the household's choices, and prices a reform the sample never observed. The price of that translation is the restrictions the model imposes, and the paper bounds the most consequential ones: the reform's cost moves by four percent of itself under perfect consumption smoothing, 86 percent of its employment gain survives an imposed firm-side layoff channel, and the only quantity that moves materially across readings of the inactivity margin is the destination of the removed retirement, which the data cannot pin down because the men who leave work for good are the ones a panel of workers rarely sees.

Three directions follow. The first is the destination question itself, which an administrative record linking disability receipt to employment histories would answer directly; the sickness and invalidity benefit already in these data is where such work would begin. The second is the household. The model explains why the wife does not insure her husband's retirement, but it understates by about half her response to his job loss when she is young, and it has no fertility and no child-care constraint, both of which shaped Italian women's participation in exactly these years (Del Boca, 2002; Attanasio et al., 2008). The third is the interaction itself. Every reform that raises retirement ages in an automating economy faces the distributional arithmetic this paper documents, and the Italian case, with its single generous exit and its sharp occupational gradient in exposure, is a setting in which that arithmetic can be measured with unusual clarity.

Data availability. The analysis combines four sources. (i) Household microdata: the Italian waves of the European Community Household Panel, waves 1–8, survey years 1994–2001, Eurostat users' database, documentation release PAN 166/2003-12. The Italian component is collected by ISTAT. The database is not public. (ii) Occupational robot exposure: the task-based index of Webb (2020), defined over 341 harmonized 1990 US Census occupations with 2010 employment weights, taken from the author's public replication material accompanying SSRN working paper 3482150 (first posted November 2019). (iii) Robot diffusion: operational stocks of industrial robots for Italy by industry from the International Federation of Robotics, *World Robotics* series. These data are licensed and are redistributed here only in the aggregated form reported in the text. (iv) Occupational crosswalks: three public concordances (international classification to the 2000 US Standard Occupational Classification, that to the 2000 US Census classification, and that to the harmonized 1990 Census codes), together with public US employment weights. All code that constructs the estimation samples, estimates the model and produces every table and figure in the paper is available from the author.

References

Acemoglu, Daron and Pascual Restrepo, "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy*, 2020, *128* (6), 2188–2244.

_ **and** _ , "Demographics and Automation," *Review of Economic Studies*, 2022, *89* (1), 1–44.

Ahituv, Avner and Joseph Zeira, "Technical Progress and Early Retirement," *Economic Journal*, 2011, *121* (551), 171–193.

Anelli, Massimo, Italo Colantone, and Piero Stanig, "Individual Vulnerability to Industrial Robot Adoption Increases Support for the Radical Right," *Proceedings of the National Academy of Sciences*, 2021, *118* (47), e2111611118.

_ **, Osea Giuntella, and Luca Stella**, "Robots, Marriageable Men, Family, and Fertility," *Journal of Human Resources*, 2024, *59* (2), 443–469.

Ardito, Chiara, "The Unequal Impact of Raising the Retirement Age: Employment Response and Program Substitution," *IZA Journal of Labor Economics*, 2021, *10* (1), 3.

Attanasio, Orazio, Hamish Low, and Virginia Sánchez-Marcos, "Explaining Changes in Female Labor Supply in a Life-Cycle Model," *American Economic Review*, 2008, *98* (4), 1517–1552.

Aubert, Patrick, Eve Caroli, and Muriel Roger, "New Technologies, Organisation and Age: Firm-Level Evidence," *Economic Journal*, 2006, *116* (509), F73–F93.

Bartel, Ann P. and Nachum Sicherman, "Technological Change and Retirement Decisions of Older Workers," *Journal of Labor Economics*, 1993, *11* (1), 162–183.

Battistin, Erich, Agar Brugiavini, Enrico Rettore, and Guglielmo Weber, "The Retirement Consumption Puzzle: Evidence from a Regression Discontinuity Approach," *American Economic Review*, 2009, *99* (5), 2209–2226.

Bessen, James, Maarten Goos, Anna Salomons, and Wiljan van den Berge, "What Happens to Workers at Firms that Automate?," *Review of Economics and Statistics*, 2025, *107* (1), 125–141.

Blanas, Sotiris, Gino Gancia, and Sang Yoon Lee, "Who Is Afraid of Machines?," *Economic Policy*, 2019, *34* (100), 627–690.

Blau, David M., "Labor Force Dynamics of Older Married Couples," *Journal of Labor Economics*, 1998, *16* (3), 595–629.

Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten, "Consumption Inequality and Family Labor Supply," *American Economic Review*, 2016, *106* (2), 387–435.

Bottazzi, Renata, Tullio Jappelli, and Mario Padula, "Retirement Expectations, Pension Reforms, and Their Impact on Private Wealth Accumulation," *Journal of Public Economics*, 2006, *90* (12), 2187–2212.

Brugiavini, Agar, "Social Security and Retirement in Italy," in Jonathan Gruber and David A. Wise, eds., *Social Security and Retirement around the World*, Chicago: University of Chicago Press, 1999, pp. 181–237.

_ **and Franco Peracchi**, "Micro-Modeling of Retirement Behavior in Italy," in Jonathan Gruber and David A. Wise, eds., *Social Security Programs and Retirement around the World: Micro-Estimation*, Chicago: University of Chicago Press, 2004, pp. 345–398.

Casanova, María, "Happy Together: A Structural Model of Couples' Joint Retirement Choices," 2010. Working paper, University of California, Los Angeles.

Casas, Pablo and Concepción Román, "Early Retired or Automatized? Evidence from the Survey of Health, Ageing and Retirement in Europe," *Journal of the Economics of Ageing*, 2023, *24*, 100443.

Caselli, Mauro, Andrea Fracasso, Sergio Scicchitano, Silvio Traverso, and Enrico Tundis, "What Workers and Robots Do: An Activity-Based Analysis of the Impact of Robotization on Changes in Local Employment," *Research Policy*, 2025, *54* (1), 105135.

Chiappori, Pierre-André, "Rational Household Labor Supply," *Econometrica*, 1988, *56* (1), 63–90.

_ , "Collective Labor Supply and Welfare," *Journal of Political Economy*, 1992, *100* (3), 437–467.

Cirillo, Valeria, Rinaldo Evangelista, Dario Guarascio, and Matteo Sostero, "Digitalization, Routineness and Employment: An Exploration on Italian Task-Based Data," *Research Policy*, 2021, *50* (7), 104079.

Coile, Courtney C., "Retirement Incentives and Couples' Retirement Decisions," *Topics in Economic Analysis and Policy*, 2004, *4* (1), 1–30. Article 17.

Costanzo, Claudio, "Robots, Jobs, and Optimal Fertility Timing," *Journal of Population Economics*, 2025, *38* (2), 1–33.

Cullen, Julie Berry and Jonathan Gruber, "Does Unemployment Insurance Crowd Out Spousal Labor Supply?," *Journal of Labor Economics*, 2000, *18* (3), 546–572.

Dauth, Wolfgang, Sebastian Findeisen, Jens Suedekum, and Nicole Woessner, "The Adjustment of Labor Markets to Robots," *Journal of the European Economic Association*, 2021, *19* (6), 3104–3153.

Del Boca, Daniela, "The Effect of Child Care and Part Time Opportunities on Participation and Fertility Decisions in Italy," *Journal of Population Economics*, 2002, *15* (3), 549–573.

Dottori, Davide, "Robots and Employment: Evidence from Italy," *Economia Politica*, 2021, *38* (2), 739–795.

Fadlon, Itzik and Torben Heien Nielsen, "Family Labor Supply Responses to Severe Health Shocks: Evidence from Danish Administrative Records," *American Economic Journal: Applied Economics*, 2021, *13* (3), 1–30.

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

_ **and John Bailey Jones**, "The Effects of Health Insurance and Self-Insurance on Retirement Behavior," *Econometrica*, 2011, *79* (3), 693–732.

Friedberg, Leora, "The Impact of Technological Change on Older Workers: Evidence from Data on Computer Use," *Industrial and Labor Relations Review*, 2003, *56* (3), 511–529.

Giuntella, Osea, Yi Lu, and Tianyi Wang, "How Do Workers and Households Adjust to Robots? Evidence from China," *Economic Journal*, 2025, *135* (666), 637–652.

Gourinchas, Pierre-Olivier and Jonathan A. Parker, "Consumption over the Life Cycle," *Econometrica*, 2002, *70* (1), 47–89.

Graetz, Georg and Guy Michaels, "Robots at Work," *Review of Economics and Statistics*, 2018, *100* (5), 753–768.

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

Halla, Martin, Julia Schmieder, and Andrea Weber, "Job Displacement, Family Dynamics, and Spousal Labor Supply," *American Economic Journal: Applied Economics*, 2020, *12* (4), 253–287.

Lu, Haiyang, Keya Zeng, and Weiliang Hu, "Does the Rise of Robotic Technology Reduce Fertility? Longitudinal Evidence from China," *Review of Economics of the Household*, 2025, *23* (4), 1371–1400.

Lundberg, Shelly, "The Added Worker Effect," *Journal of Labor Economics*, 1985, *3* (1), 11–37.

Matysiak, Anna, Daniela Bellani, and Honorata Bogusz, "Industrial Robots and Regional Fertility in European Countries," *European Journal of Population*, 2023, *39* (1), 11.

Mazzocco, Maurizio, "Household Intertemporal Behaviour: A Collective Characterization and a Test of Commitment," *Review of Economic Studies*, 2007, *74* (3), 857–895.

McFadden, Daniel, "Modelling the Choice of Residential Location," in Anders Karlqvist, Lars Lundqvist, Folke Snickars, and Jörgen Weibull, eds., *Spatial Interaction Theory and Planning Models*, Amsterdam: North-Holland, 1978, pp. 75–96.

Peracchi, Franco, "The European Community Household Panel: A Review," *Empirical Economics*, 2002, *27* (1), 63–90.

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

Stephens, Melvin, Jr., "Worker Displacement and the Added Worker Effect," *Journal of Labor Economics*, 2002, *20* (3), 504–537.

Stock, James H. and David A. Wise, "Pensions, the Option Value of Work, and Retirement," *Econometrica*, 1990, *58* (5), 1151–1180.

Webb, Michael, "The Impact of Artificial Intelligence on the Labor Market," 2020. Working paper, Stanford University. SSRN Working Paper 3482150.

Yashiro, Naomitsu, Tomi Kyyrä, Hyunjeong Hwang, and Juha Tuomala, "Technology, Labour Market Institutions and Early Retirement," *Economic Policy*, 2022, *37* (112), 811–849.

A. Data and Variable Construction

Activity states. The classification of Section II is validated against the panel's own employment, participation and hours variables, which are built from different underlying questions and are not used to construct the states. The validation confirms the category that matters: men who report self-employment are recorded as employed and as participating on those independent fields, men who report unemployment as participating but not employed, and men who report retirement as neither. The residual inactive category is recorded as neither employed nor participating, and it is kept apart from retirement because the counterfactual asks whether closing the pension door diverts men into it.

Income. The survey records total net personal income and its components: net income from work, old-age pension, unemployment benefit, and sickness and invalidity benefit. All income is net of tax and social contributions, refers to the calendar year before the interview, and is denominated in thousands of lire. The conversion is 1,936.27 lire to the euro and figures are not deflated. The component reading is verified against the activity states: the old-age pension component is worth 0.7100 of a working man's income to a retired man and 0.0086 to a working one, and the unemployment benefit component is small in every state, which is the source of the 0.14 replacement rate on unemployment. Replacement rates are computed on state-stable households and on personal income only, so that the wife's pension and any other household member's income are excluded by construction. Section II reports the comparison with the residual measure.

Occupational exposure. The exposure index of Webb (2020) is defined over 341 harmonized 1990 US Census occupations with 2010 employment weights. The crosswalk to the survey's occupational detail chains three public maps (international classification to the 2000 US Standard Occupational Classification, that to the 2000 Census classification, and that to the 1990 harmonized codes) with the digit-dropping and residual-category conventions those maps require. Scores are aggregated by 2010 US employment weights to the seventeen occupation groups the survey carries and, as a fallback, to the nine one-digit major groups. The survey's negative missing codes on the detailed occupation field are screened out before the husband's first observation is taken. Without that screening a missing code is treated as an occupation, 263 couples instead of 43 fall back on the one-digit classification, and the score takes 23 distinct values instead of 18. Panel C of Table B1 reports the estimate under both constructions and under a third built on the husband's modal occupation. One cell, teaching professionals, carries a legacy low value of 3.0; dropping it leaves every reported estimate unchanged.

The stated reason for stopping work. The reason-for-leaving item used in Panel B of Table 8 is administered in two versions. Respondents who are working at the interview are asked why they left their previous job, and the first response category is "obtained a better or more suitable job". Respondents who are not working are asked why they left their last job, and the first category is instead "retired at normal age". The two versions are pooled into a single variable in the user database, whose published code list reproduces only the first version. The distinction is restored using the respondent's activity status at the interview; every observation in the retirement-transition sample falls in the second version. The second response category, "obliged to stop by employer", covers business closure, redundancy, employer-initiated early retirement and dismissal in both versions, and the published label truncates that parenthetical.

The monthly route. Panel A of Table 8 chains the survey's monthly calendar of activities across consecutive waves into a continuous monthly record of each man's status, and classifies each observed passage from work into retirement by whether any non-employed month intervenes between the last month of work and the first month of retirement.

Aggregate robot diffusion. Operational robot stocks for Italy by industry are taken from the International Federation of Robotics and bridged to the European industrial classification. Regional exposure divides each sector's national robot stock by its base-year national employment and weights by the region's base-year employment share in that sector, so that all cross-regional variation comes from the fixed industry mix. Employment is aggregated from provinces to regions before shares are formed, which is exact. The denominator is the establishment-survey employment concept throughout, consistently in numerator and denominator. Using the labor-force concept in the denominator instead mixes statistical universes and inflates cross-regional dispersion, and is reported only as a robustness check in the source construction.

B. Validity of the Difference-in-Differences

This appendix collects the identification diagnostics for the threshold design of equation (4). Table B1 reports the effective sample behind the estimate, the coefficient under the departures discussed in Section III, and the exposure measure under alternative constructions. Figure B1 draws the hazard profiles of Figure 1 for the full sample and the remaining schooling groups. Three further checks support the upper-secondary difference-in-differences shown in Figure 2, and all three use the couple panel with the screened occupation score.

Table B1. Identification diagnostics for the exposure–retirement estimate

Panel A. The effective sample

	Couples	Couple-waves	Retirement events	Occupation cells	Variance share	β Coef.	(s.e.)
Straddle age 55	263	1,472	81	17	0.9901	+0.0480	(0.0163)
Do not straddle	989	2,865	351	18	0.0099	+0.0000	(0.0400)

Panel B. The coefficient under four departures from the specification

Specification	Coefficient (standard error) on: Exposure $\times \mathbf{1}[a \geq 55]$	$\mathbf{1}[a \geq 55]$	Tertiary $\times \mathbf{1}[a \geq 55]$	n
Equation (4)	+0.0599 (0.0170)	—	—	4,337
+ the age-55 main effect	+0.0587 (0.0183)	−0.0092 (0.0253)	—	4,337
+ a schooling threshold	+0.0443 (0.0218)	—	−0.0848 (0.0383)	4,327
Cell × year effects	+0.0235 (0.0203)	—	—	4,337
Unscreened exposure score	+0.0604 (0.0170)	—	—	4,337
Modal-occupation score	+0.0544 (0.0168)	—	—	4,337
All four together	+0.0092 (0.0246)	+0.0016 (0.0271)	−0.0958 (0.0458)	4,327

Panel C. The exposure measure

Score construction	Distinct cells	On the one-digit fallback	Couples whose score differs	β Coef.	(s.e.)
First observed, unscreened	23	263	—	+0.0604	(0.0170)
First observed, missing codes screened	18	43	184	+0.0599	(0.0170)
Modal over all valid waves	19	94	916	+0.0544	(0.0168)

Notes: Panel A splits the risk set of Table 5 by whether the couple is observed on both sides of age 55. The variance share is the share of the squared interaction after couple and year effects are removed; the non-straddlers' share is a projection residual, and estimating on them alone returns a coefficient of zero to machine precision. The straddling couples are 6 birth cohorts, with couple counts 1941 (59), 1942 (39), 1943 (38), 1944 (41), 1945 (43), 1946 (43). Panel B re-estimates equation (4) under each departure named in the row; the last row imposes all four at once. Cell × year effects give each exposure cell its own year effect. Schooling is the husband's at his first observation. Panel C rebuilds the exposure score end to end and re-estimates equation (4); the score used throughout the paper is the second row, and the modal score correlates 0.8589 with the unscreened one while the screened score correlates 0.9886. Standard errors clustered on the couple throughout.

Figure B1 repeats the Panel A construction of Figure 1 on the estimation sample for the full sample and the other schooling groups, each split at its own median exposure.

Among husbands with lower-secondary schooling or less, eligibility arrives before the early fifties for exposed and unexposed alike, both groups retire early, and the within-group gap is small. Among the tertiary-educated, retirement begins only in the late fifties and the cells are thin. In the pooled profile the gap opens in the early fifties because the exposed are disproportionately the least schooled, whose option arrives first.

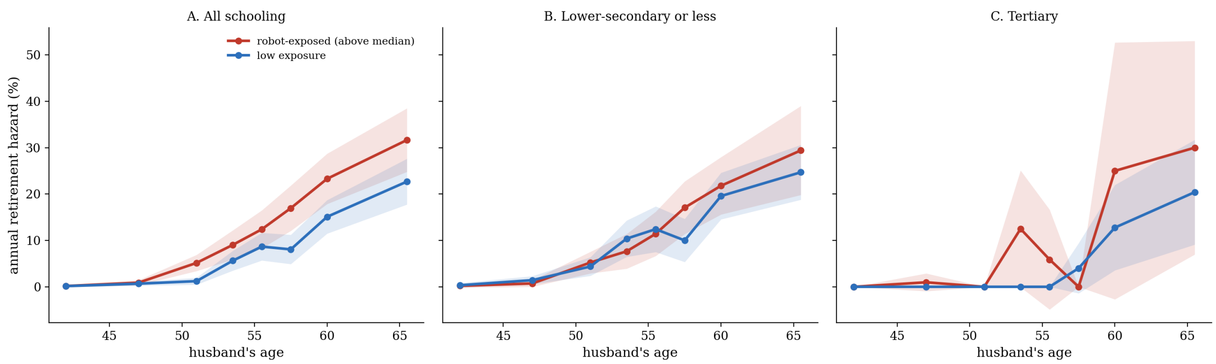

Figure B1. Retirement hazards by age and exposure, by schooling group

Notes: As Panel A of Figure 1, on the estimation sample: the annual hazard from work into retirement by the husband's age, split at the median of the standardized exposure score within each frame. Shaded bands are pointwise 95 percent confidence intervals clustered by couple.

The response is specific to the pension exit. Table B2 reports the exposure gradient at the threshold on each of the three destinations. It is +0.0365 on the flow into retirement and a precisely estimated zero on the flows into unemployment and into inactivity, whose standard errors are an order of magnitude smaller than the retirement coefficient under either clustering.

Table B2. The exposure response is specific to the pension exit

Flow from work into	Coefficient	Occupation s.e.	Couple s.e.	Events
Retirement	+0.0365	(0.0234)	(0.0115)	578
Unemployment	−0.0002	(0.0027)	(0.0042)	103
Other inactivity	+0.0033	(0.0048)	(0.0039)	83

Notes: The coefficient on the standardized exposure score interacted with $\mathbf{1}[\text{age} \geq 55]$, from linear probability models of the annual hazard from work into each destination, husbands 50–69 with couple and year fixed effects, on 6,817 couple-waves. Standard errors are clustered on the couple and, separately, on the occupation cell (20 cells). The response is confined to retirement; the two competing exits are precisely estimated zeros.

The response also scales with exposure. Figure B2 repeats the difference-in-differences of Figure 2 for upper-secondary men by tercile of the exposure score. The three groups leave work at the same rate before the eligibility ages and separate in order of exposure after them.

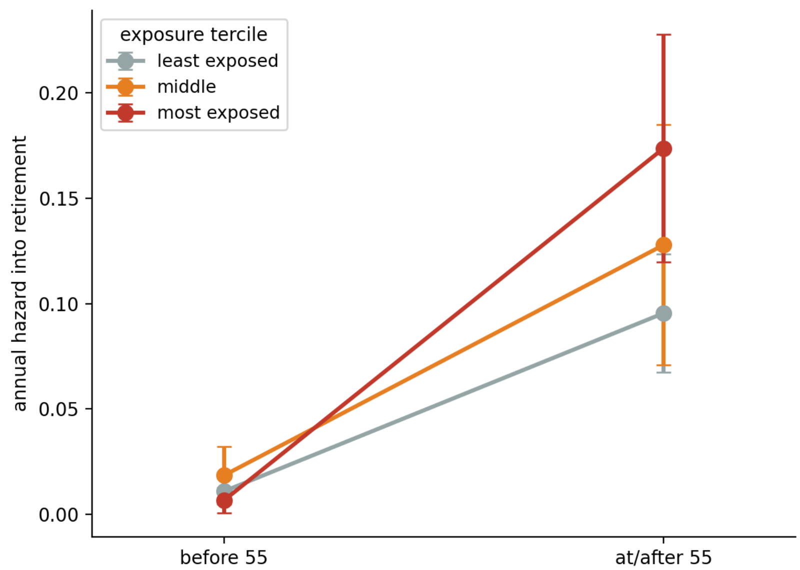

Figure B2. Retirement before and after the eligibility age, by exposure tercile, upper-secondary men

Notes: As Figure 2, with upper-secondary men grouped into terciles of the standardized exposure score. Whiskers are 95 percent confidence intervals.

Finally, because exposure varies across only twenty occupation cells, the sharpest threat is that the retirement result is an accident of which occupations received which scores. Figure B3 addresses it by reassigning the twenty cell scores at random across the cells 1,500 times and re-estimating equation (4) on each draw. The coefficient on the true assignment, +0.0365, lies in the tail of the resulting distribution: 8.3 percent of random reassignments produce a coefficient as large in absolute value.

The twenty cells take eighteen distinct scores, because two pairs of cells share a value, and the two counts support two versions of the same exercise. Permuting across the twenty cells, as here, gives $p = 0.083$. Permuting the eighteen distinct values, as Table 6 does, gives $p = 0.058$. The coarser permutation is the more conservative of the two, and neither reaches five percent. The design-based evidence on this coefficient falls just outside conventional significance however the permutation is organized.

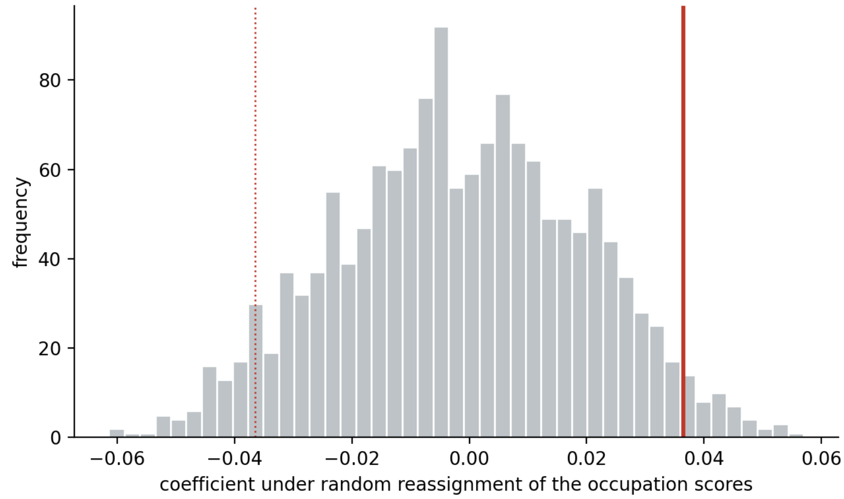

Figure B3. The equation (4) coefficient under random reassignment of the occupation scores

Notes: Distribution of the coefficient on exposure $\times\mathbf{1}[\text{age} \geq 55]$ over 1,500 random reassignments of the twenty occupation-cell scores across the cells, re-estimating equation (4) on husbands 50–69 each time. The solid line is the coefficient on the true assignment. The dotted line is its reflection.

C. Solution of the Model

The value function satisfies

$$V_a(s_{-1}, n_{-1}, e_h, e_w; z, \Delta) \;=\; \sum_b \Pr(b \mid s_{-1}) \sum_c \Pr(c \mid n_{-1}, a_w)\, \tilde{V}_a(b, c, e_h, e_w), \qquad (14)$$

where b indexes the husband's opportunity branch (no separation, separation into unemployment, separation into inactivity, arrival or non-arrival of an offer) and c the wife's, and $\tilde{V}_a$ is given by (10)–(12) evaluated on the feasible sets that branch pair defines. Eligibility enters the continuation in (9) through

$$\mathbb{E}\, V_{a+1} = \sum_{e'_h, e'_w} \Pr(e'_h \mid e_h, a)\; \Pr(e'_w \mid e_w, a_w, n')\, V_{a+1}, \qquad (15)$$

with the husband's transition hazard implied by the marginal distribution of a_h^* and the wife's equal to that hazard if $n' = E$ and zero otherwise. Retirement is absorbing for both spouses. At the terminal age the continuation value is a constant, which does not distort choices at any age.

The problem is solved by backward induction on a grid of exposure levels and spousal age gaps, with each simulated couple's policy obtained by bilinear interpolation at its own score and its own age gap. The exposure grid places a node at every distinct value the score takes, so the interpolation in that dimension is exact. The eligibility state is binary and persistent for both spouses. The husband's is simulated from a latent threshold drawn once, which reproduces the marginal distribution exactly, and the wife's from the period-by-period hazard, since hers depends on her employment path.

D. Numerical Implementation and Estimation

The criterion (13) is minimized in three stages. The search first cycles over three blocks: the husband's exit parameters against his hazards and state distribution, the wife's participation parameters against her two hazards, and the joint-leisure parameter against the cross-spouse responses. Each block is searched globally by differential evolution on the moments that identify it. A global search over all parameters against the full criterion follows, seeded from the block solution, and a Nelder–Mead pass completes it. Nothing is left block-optimized, and as Section V records, the block structure is only a warm start: the joint-leisure parameter's own-block share of precision-weighted Jacobian mass is 0.104, against a median of 0.96 across the twenty-three parameters, which is why full-criterion evaluations can deteriorate across block cycles and why the block stage is followed by a full-space search.

The block and global stages run at a reduced number of replicates of the observed panel, because they require tens of thousands of evaluations. The estimate is then refined directly against the twelve-replicate objective, and the fit reported in Table 12 is computed at that resolution.

The criterion is minimized on one set of simulation draws, and the value reported for the estimates is that draw's. Re-simulating the same parameter vector on twenty-four independent sets of draws at the same resolution gives a mean criterion of 211.2 with a standard deviation of 17.5 and a range of $[174.0, 235.9]$; the reported value is the smallest of the twenty-five. Two things follow. The level of the criterion is not an estimate of anything: it is quadratic in the simulated moments, so simulation error raises it in expectation, and minimizing over a fixed draw selects that draw's favorable side. Comparisons of the level across studies or across simulation designs are correspondingly uninformative and are not made here. It does not follow that the parameters are wrong. The draws are held fixed across parameter values, so this is the standard simulated method of moments.

The estimator remains consistent, its simulation variance falls as the number of replicates rises, and the standard errors reported in Table 13 already carry the simulation correction for the twelve replicates used. The remedy for the level of the criterion is more replicates.

Table D1 sizes the Monte Carlo error in the moments themselves. It is below the sampling standard error of the corresponding data moment for every reported moment, so the misfits recorded in Section V are due to the model, with simulation error too small to produce them. The moment for which simulation error is largest relative to the data standard error is the retirement hazard at 65–69, at 0.91. The moment for which it is largest relative to the model coefficient itself is the wife's employment response to her husband's inactivity, which is another way of saying that this coefficient is zero in the model as it is in the data.

Table D1. Monte Carlo error in the simulated moments and in the criterion

Panel A. The moments

Moment	Estimated model	Independent-draw mean	Monte Carlo s.d.	Data s.e.	Ratio
Her employment, he retired	−0.0514	−0.0395	0.0066	0.0167	0.40
Her employment, he unemployed	+0.0308	+0.0240	0.0102	0.0323	0.32
Her employment, he inactive	+0.0034	+0.0080	0.0178	0.0262	0.68
His employment, she retired	−0.0906	−0.0821	0.0164	0.0426	0.38
His employment, she non-employed	−0.0297	−0.0182	0.0095	0.0309	0.31
Her retirement on his	+0.0269	+0.0207	0.0043	0.0152	0.29
Her entry response, wives 25–44	+0.0445	+0.0463	0.0104	0.0456	0.23
Her entry response, wives 45–69	+0.0171	+0.0175	0.0117	0.0338	0.35
Retirement hazard, 55–59	+0.1495	+0.1514	0.0030	0.0088	0.34
Retirement hazard, 65–69	+0.4432	+0.4174	0.0436	0.0479	0.91
Exposure gradient in retirement	+0.0522	+0.0626	0.0054	0.0170	0.32
Retired share, 50–69	+0.2470	+0.2495	0.0025	0.0060	0.41

Panel B. The criterion

Canonical draw (the estimation seed)	174.0
Independent draws, mean	211.2
Independent draws, standard deviation	17.5
Independent draws, range	[174.0, 235.9]

Notes: The estimated model re-simulated on 24 independent sets of random numbers at the reporting resolution of 12 replicates of the observed panel. The estimated-model column is the draw on which the criterion was minimized and is the value reported throughout the paper; it is not among the 24 independent draws, and none of those falls below it. The last column of Panel A is the Monte Carlo standard deviation as a fraction of the sampling standard error of the corresponding data moment. Holding the seed fixed and varying the resolution instead gives 240.4 at 4 replicates, 174.0 at 12 replicates, 202.2 at 20 replicates.

E. The Estimation Sample and the Population

The estimation sample requires the husband to be observed at work at least once inside the observation window. The restriction is innocuous at younger ages and severe at older ones, because a man who retired before the panel opened is never seen working. Comparing the data, the model's counterpart of the estimation sample, and the model's population at ages 60–64: the working share is 0.4904, 0.4150 and 0.0726. The retired share is 0.4496, 0.5713 and 0.8991. The gap is widest on inactivity, which is why the diversion into it under the counterfactual carries the bound recorded in Section VI. Model fit throughout Section V is assessed on the sample basis. The counterfactual in Section VI is reported on the population basis.

The same restriction explains an apparent contradiction in Section V: the model is too retired in every age band from fifty onward, yet the pooled retired share at 50–69, which the criterion targets, matches at 0.2470 against 0.2466. The reconciliation is arithmetic. Within 50–69 the data put 0.438 of couple-waves at ages 50–54 and 0.056 at 65–69, while the simulated panel, after the same requirement that the husband be seen at work, puts 0.586 and 0.015 there: the model retires men earlier, so more of its couples fail the requirement at older ages and drop out. Applying the model's own band shares to the data's age composition gives a pooled retired share of 0.3200 against the data's 0.2466. The composition shift is therefore worth 0.0730 of the pooled share, and it allows a model that over-states retirement within every band to match the pooled moment. The offset is a property of the estimation sample itself, which is why the counterfactual is computed on the population.

F. The Model Estimated on the Full Population

Table F1 estimates the model on the full population, including self-employed couples, and compares it with the main specification. That sample is 28,290 couple-waves on 4,936 couples, against 19,611 on 3,375 in the main specification, and its retirement risk set contains 6,817 couple-waves and 578 transitions from work to retirement, against 4,337 and 432. The moment construction, the simulated panel design, the estimation stages and the search budgets are identical to the main specification. The exposure score is standardized on the sample being estimated, and the exposure grid again places a node at every distinct score.

The model fits the full population less well. The criterion is 261.5 against 174.0, and 33 of the 54 targeted moments fall within two standard errors against 39. The deterioration is concentrated on the moments that identify the retirement margin, because a single set of preferences must now reconcile the men who faced the seniority-pension threshold with the self-employed, who did not. Forced to fit a diluted gradient, the estimate of the disutility of continuing to work in an exposed occupation rises to 1.1223 against 0.8375. The two criterion levels are not comparable in level: the larger sample delivers tighter standard errors on the moments that identify the retirement margin, where the exposure gradient in the data is +0.0374 against +0.0599 among employees. The fit is nonetheless worse on the count of moments matched, and the restriction to employee households follows from the institution.

Table F1. The model estimated on the full population

		Main specification	Full population†
Level	θ_0	0.8528	1.0874
Slope per year of age above 50	θ_1	0.0151	0.0122
Additional slope per year above 60	θ_2	0.1541	0.0605
Robot exposure, per standard deviation	δ	0.8375	1.1223
Median latent eligibility age	a_{0h}	54.2826	54.6147
Dispersion of eligibility ages	τ_h	2.3503	2.6604
Share ever eligible	$\bar{e}_h$	0.9824	0.9956
Separation into unemployment	π_U	0.0212	0.0121
Separation into inactivity	π_O	0.0111	0.0134
Job arrival out of unemployment	λ_U	0.7081	0.6605
Job arrival out of inactivity	λ_O	0.8422	0.7103
Level	θ_{w0}	0.3799	0.3296
Slope per year of her age above 45	θ_{w1}	0.0220	0.0243
Fixed utility cost of entering employment	κ	3.4488	3.9149
Opportunity arrival, logit intercept at age 45	m_0	−2.3103	−0.8904
Opportunity arrival, slope in her age	m_1	−0.0899	−0.0924
Median latent eligibility age	a_{0w}	55.9516	57.0247
Dispersion	τ_w	4.0777	4.0558
Share ever eligible	$\bar{e}_w$	0.9914	0.9974
Value of joint non-work	ϕ	0.1183	0.2778
Scale, her work / non-work decision	σ_w	1.3148	1.0291
Nesting, her non-work states: $\sigma_n = \rho_n \sigma_w$	ρ_n	0.1611	0.1845
Nesting, his decision: $\sigma_h = \rho\, \sigma_n$	ρ	0.4205	0.5712
Criterion		174.0	261.5
Targeted moments within two s.e.		39 of 54	33 of 54
Couple-waves		19,611	28,290
Couples		3,375	4,936
Retirement risk set, couple-waves		4,337	6,817
Work-to-retirement transitions, 50–69		432	578
Exposure gradient in the data		+0.0599	+0.0374

The targeted moments whose fit differs most between the two estimations

	Main specification			Full population		
Moment	Data	Model	t	Data	Model	t
Working, 50–69	+0.6980	+0.7104	+2.1	+0.7128	+0.7453	+5.4
Exit to non-employment, 45–49	+0.0523	+0.0591	+1.1	+0.0582	+0.0830	+4.1
Retired, 55–59	+0.1533	+0.1868	+3.5	+0.1464	+0.1913	+6.1
Exposure $\times \mathbf{1}[\text{age} \geq 55]$ on retirement	+0.0599	+0.0522	−0.5	+0.0374	+0.0713	+2.7

Notes: †Full population (including self-employed couples). The comparison column estimates the same model on that sample. The moment construction, the simulated panel design, the estimation stages and the search budgets are identical to the main specification; the exposure score is standardized on the sample being estimated, and the exposure grid places a node at every distinct score in both. Each column's t-statistics use that column's own data standard errors. The full-population sample is larger and its standard errors are correspondingly tighter, while more of them reach the 0.006 floor, so the two criterion levels are not comparable in level.