COSTLY ATTENTION AND RETIREMENT

JAMIE HENTALL-MACCUISH
Department of Economics and Decision Science, HEC Paris

In UK data, I document the prevalence of misbeliefs regarding the State Pension eligibility age (SPA) and their predictivity for retirement. Exploiting policy variation, I estimate a life-cycle model of retirement in which, motivated by patterns in belief data, rationally inattentive households learning about uncertain pension policy endogenously generates misbeliefs. Misbeliefs explain 51% of the excessive (given financial incentives) drop in employment at SPA when constrained to replicate the belief data patterns and completely explain it when not. To achieve this, I develop a solution method for dynamic rational inattention models with persistent beliefs. Costly attention makes the SPA up to 15% less effective at increasing old-age employment. Hence, information letters improve welfare and increase employment.

KEYWORDS: Rational inattention, Retirement, Misbeliefs, Pensions, Behavioral Macro, Structural Econometrics.

1. INTRODUCTION

Understanding why households appear to deviate from rational behavior is crucial for policy design. If such deviations reflect fixed preferences or fixed features of household behavior, policy options to address them are limited. But mistaken beliefs about policy due to limited attention can produce similar departures, as emphasized by Gabaix (2019). In these cases, straightforward information provision might be effective. This paper argues that misbeliefs offer an alternative, or complementary, explanation for a well-known puzzle often attributed to fixed household preferences: the disproportionately large drop in employment at pension eligibility ages, despite weak economic incentives to stop working precisely then.[1] To achieve this, I develop a general-purpose solution method for dynamic rational inattention models with persistent beliefs and use it to estimate a model on UK data, targeting both observed beliefs and behavior.

Retirement is a compelling context to study the impact of misbeliefs due to their prevalence.[2] Many people are confused about pensions. In my data, 59% of women affected by pension age reform are mistaken about their pension age by over a year when within 2-4 years of eligibility. Initially, these misbeliefs seem strange, since the information is financially relevant and freely available. However, they become less surprising when we acknowledge that government policy is objectively uncertain (changing in unpredictable ways), *and* information is costly. Together, policy uncertainty and costly information can generate these misbeliefs as an optimal response. This paper asks whether these endogenously generated misbeliefs, in turn, help explain excess employment sensitivity to pension eligibility.

Jamie Hentall-MacCuish: hentall-maccuish@hec.fr

I thank Fabien Postel-Vinay and Eric French for discussions and encouragement. For helpful comments, I thank Richard Blundell, Mariacristina De Nardi, Gaetano Gaballo, Tomasz Michalski, Lars Nesheim, Cormac O'Dea, Morten Ravn, Morgane Richard, Victor Rios-Rull, Arthur Seibold, Johannes Spinnewijn, and participants at the NBER SI Behavioral Macro Session, IIPF Annual Congress, RES Junior Symposium, CESifo Public Area Conference, YES, NETSPAR Pensions Workshop, Econometric Society European Meeting, CERGE-EI, ENTER, Toulouse, Cambridge, UCL, UWO, UNSW, Tilburg, Edinburgh, Royal Holloway, HEC, and the IFS. For thoughtful discussion, I thank Andrew Caplin, Dylan Moore, Leanne Nam, Nicolas Fernandez-Arias, Jim Been, Patrizia Alexandra-Massner, and Sebastian Seitz. Funding from Grant Inequality and the insurance value of transfers across the life cycle (ES/P001831/1), ESRC studentship (ES/P000592/1), and HEC is gratefully acknowledged.

[1] This puzzle is documented in multiple countries as summarised in Gruber and Wise (2004).

[2] Documented, for example, in Gustman and Steinmeier (2005), Lusardi and Mitchell (2011), Ciani et al. (2023).

To investigate, I first document key facts on misbeliefs and excess employment sensitivity, then I separately and sequentially introduce policy uncertainty and information frictions (in the form of costly attention) into a model of retirement. Specifically, I estimate a dynamic lifecycle model of retirement (e.g. Rust and Phelan, 1997, French, 2005) with rationally inattentive households (e.g. Sims, 2003, Matějka and McKay, 2015, Caplin et al., 2019) deciding how much information about a changeable pension policy to acquire whilst incurring a disutility cost of information. The model endogenously generates observed misbeliefs about the pension eligibility age, but can it generate the otherwise puzzling sharp employment drop at this eligibility age (known as State Pension Age, or SPA)? The drop in employment at SPA is puzzling in the UK, as pension benefits are not tied to employment, so the SPA only incentivizes retirement for liquidity-constrained individuals unable to stop working earlier. Yet, employment also falls among those with substantial liquid wealth.

Counterintuitively, unawareness of the SPA is not only consistent with high employment sensitivity to the SPA but is essential to generating this sensitivity. The revelation of information upon reaching eligibility explains this. In the model, households pay a utility cost to learn about their eligibility age (SPA), modeled as stochastic to capture potential government reforms. Upon reaching the SPA, its value becomes fixed and is revealed, reflecting communication of eligibility and information disclosure during claiming. Thus, reaching the SPA is a positive information shock. It is also a positive wealth shock because, as households age past earlier alternative eligibility ages without receiving benefits, they rule those ages out, making now the earliest possible eligibility age. This information shock reduces precautionary labor supply, and since leisure is a normal good, the wealth shock also reduces labor supply. These mechanisms exist in a model with only policy uncertainty, but by introducing policy uncertainty and costly attention separately, this paper shows that the historically observed level of policy uncertainty is too low to generate meaningful changes. Hence, misbeliefs generated by costly attention are key to amplifying these positive shocks at the SPA.

These model mechanisms rely on the potential for the government to reform the SPA rather than on the occurrence of reforms, though the 1995 and 2011 reforms demonstrate that this potential is real. I do rely on the occurrence of reforms as identifying variation, firstly to estimate the probability of reform and secondly to causally identify the effect of the SPA on employment. Support for the external validity of misbeliefs shaping claiming and retirement patterns comes from van der Klaauw and Wolpin (2008) and Bairoliya and McKiernan (2023), both of which study the US, with the latter also focusing specifically on pension-policy misbeliefs.

I focus on costly attention to the SPA rather than any other burdens on people's attention for three reasons. One, pension policy uncertainty (unlike, for example, return uncertainty) resolves, or at least diminishes, upon eligibility, making it a potential explanation for employment responses at the SPA. Two, the SPA's simplicity (relative to other dimensions of pension-policy uncertainty, such as the benefit level) makes SPA misbeliefs easy to measure and, hence, study (additionally, this simplicity makes the misbeliefs we observe all the more surprising). Three, it is the dimension of pension policy corresponding to the most informative subjective-belief measure in my dataset.

In the data, misbeliefs about the SPA predict employment responses to it, motivating the joint study of misbeliefs and excess sensitivity. Women more mistaken about their SPA in their late 50s show a smaller response upon reaching it in their early 60s. The model replicates this pattern because varying returns to information lead to selection into attention. Women unconcerned by the SPA neither learn nor respond upon reaching it. So, whilst the information shock resulting from misbeliefs is an aggregate mechanism generating employment responses to the SPA, selection into SPA knowledge generates a cross-sectional correlation in which more mistaken individuals respond less. Thus, information endogeneity and return heterogeneity are

crucial for replicating the relationship between beliefs and employment. As well as being crucial to replicating the negative correlation between misbeliefs and employment responses to the SPA, endogenous information also replicates the fact that beliefs tend toward the truth as people age.

So, the endogeneity of beliefs drives the relationship between retirement and misbeliefs, but it complicates the model by introducing a high-dimensional state (prior beliefs) and choice (learning strategy). In static rational inattention models, prior beliefs represent ex-ante heterogeneity, but in dynamic models, today's learning affects tomorrow's beliefs, making beliefs a state variable. Many papers sidestep this by suppressing prior beliefs as a state variable.[3] While reducing the state space is beneficial and suppressing beliefs can be a good modeling assumption for specific situations, it limits the domain of application by implying that beliefs are irrelevant to choices. It cannot capture scenarios where data show that beliefs matter and vary across individuals, such as UK pension beliefs. Instead, I develop a solution method for dynamic rational inattention models that accommodates persistent beliefs by treating them as a state. The method is general-purpose, as it models beliefs nonparametrically without imposing restrictions on the data-generating process. It operationalizes the results of Steiner et al. (2017) for use with rich structural dynamic rational inattention models, while keeping the beliefs state-space representation, and addressing the computational challenges of high-dimensional state spaces by exploiting the sparsity property established in Caplin et al. (2019). Since beliefs are persistent, relevant, and imperfectly informed in many economically important environments, the scope for applications outside of retirement is considerable. For example, as the solution method makes tractable dynamic models where belief formation is an optimal response, it could be profitably applied to environments where beliefs are persistent, and their dynamics are important, such as job search (beliefs about job-finding rates) or portfolio choice (beliefs about returns).

The analysis uses data from the English Longitudinal Study of Ageing (ELSA), a micro panel survey. ELSA contains detailed subjective pension-belief data, in particular, self-reported and true SPAs, which allow the construction of a panel of SPA misbeliefs. Alongside these belief measures, it provides rich information on assets, labour-market status, and demographics. The survey is also linked to administrative records, in particular social-security contribution histories, enabling the estimation of individuals' State Pension entitlements. The recent presidential address to the Econometric Society (Almås et al., 2024) advocates for greater use of survey-based subjective-belief data, a call to which this paper responds. As noted in the address, survey data allow us to observe variables, such as beliefs, that are not visible in administrative records alone. Survey data, though, can be noisy, and I assess robustness to measurement error in beliefs data as carefully as possible.

I estimate the model using two-stage simulated method of moments, targeting asset and employment profiles, and, when present, identifying attention costs from changes in individual misbeliefs over time. Targeting changes in beliefs is possible thanks to my solution method, which, by retaining beliefs as a state variable, endogenously generates belief predictions that can be compared to the data. Thus, my solution method builds a bridge between the dynamic-rational-inattention literature and the subjective-belief-data literature. Policy uncertainty, combined with costly attention, increases the employment response to the SPA compared to a complete information baseline. Together they explaining 51% of the puzzling shortfall among richer household, when identifying the cost of attention from the beliefs data and completely explaining it when directly targeting the puzzle. The mean household is willing to pay £11.00-£31.00 to learn today's SPA, so estimated attention costs are low, consistent with other evidence that

[3]For example Miao and Xing (2024), Armenter et al. (2024), Turen (2023), Macaulay (2021), Porcher (2020).

apparently large deviations from optimizing behavior are explained by modest friction in dynamic environments (e.g., Chetty, 2012, Choukhmane, 2025). Large changes in the employment response at SPA stem from small attention costs because the concentrated response at SPA partly reflects an intertemporal shift in employment relative to the frictionless benchmark.

Pension eligibility ages are considered key to increasing old-age labor force participation, which is a common policy goal. Since costly attention increases employment response *at* the SPA compared to full information, one might assume it makes the SPA a better tool for this purpose. The opposite is generally true. Policy experiments comparing employment increases resulting from SPA changes in versions of the model with and without information frictions show that costly attention shifts part of the informed agent's response forward but can lower the overall response. Informed agents increase labor supply immediately, while less informed individuals, facing learning costs, respond closer to their SPA. Thus, informing individuals, for example, by sending letters, could raise old-age employment by up to 15%. In most policy experiments, the benefits to households and extra tax revenue from these letters, each separately, outweigh the costs: considered jointly, information letters are always welfare-enhancing.

Related Literature. Dynamic lifecycle models of retirement began with Gustman and Steinmeier (1986) and Burtless (1986). Key features introduced since then include uncertainty (Rust and Phelan, 1997), borrowing constraints (French, 2005), subjective life-expectancy and Medicare (van der Klaauw and Wolpin, 2008), and medical expenses (French and Jones, 2011). Much of this literature is US-focused, and some of its concerns, like medical insurance, are irrelevant to the UK. My model includes uncertainty, borrowing constraints, and individual heterogeneity. O'Dea (2018) models male UK retirees and is the closest paper in this literature.

Rational inattention began as a way to add costly attention to macroeconomic models (e.g., Sims, 2003, Maćkowiak and Wiederholt, 2009, 2015), but now touches most fields, e.g., industrial organization (Brown and Jeon, 2024), or labor economics (Bartoš et al., 2016). Matějka and McKay (2015) solve a general class of static discrete choice models with rationally inattentive agents, and Steiner et al. (2017) extends these results to dynamic discrete choice models. A key contribution of this paper is turning the theoretical solutions of Steiner et al. (2017) into a solution method for quantitative dynamic rational inattention models with history-dependent beliefs. Caplin et al. (2019) show rational inattention generically implies consideration sets, meaning solutions are sparse, which I leverage to reduce computational burden. Dynamic rational inattention typically avoids these computational issues by suppressing the belief distribution as a state variable (e.g. Miao and Xing, 2024, Armenter et al., 2024, Turen, 2023, Macaulay, 2021, Porcher, 2020). While reasonable for specific cases, this approach is not fully general and limits the range of questions that can be answered. Two recent papers (Miao et al., 2022, Afrouzi and Yang, 2021) also propose methods for dynamic rational inattention that incorporate beliefs as a state variable. Both use the linear-Gaussian-quadratic framework popular in macro rational inattention to speed up solutions, whereas my approach handles arbitrary noise and utility but lacks the performance gains from restricting the class of models. A closely related static rational inattention paper Boehm (2023) estimates a lifecycle model of older individuals, focusing on the one-shot annuity choice.

First highlighted in the US by Lumsdaine et al. (1996), a puzzlingly large drop in employment at pension-eligibility ages is observed across countries. In the US, the consensus was that liquidity constraints explained the drop at age 62, and Medicare eligibility the drop at age 65 (Rust and Phelan, 1997, French, 2005, French and Jones, 2011). Testing these explanations became possible after 2004, when the full retirement age increased. Part of the age 65 spike followed the full retirement age, despite Medicare eligibility staying at 65 (Behaghel and Blau, 2012), and Mastrobuoni (2009) found larger effects than standard models predicted. Pension

age increases around the world produced similar results: larger employment responses than financial incentives implied (summarised in Gruber and Wise, 2004). I document this in the UK, extending Cribb et al. (2016) by using richer data to rule out other potential explanations. Part of the literature has recently converged towards reference-dependence as the explanation of this puzzle (e.g. Seibold, 2021, Lalive et al., 2023, Gruber et al., 2022). I compare my results to this explanation in Section 8.2 and Appendix E.3.

The use of subjective belief data in structural microeconomic models is extensive but recent (Koşar and O'Dea (2022) provide a summary; some notable early examples include Pistaferri (2001) and Delavande (2008)). Most papers, however, do not model belief formation, limiting counterfactual analysis (e.g. de Bresser, 2023). Modeling belief formation as an optimal response to processing costs (enabled by my solution method) allows me to match model-generated beliefs to data rather than treating belief data solely as input. Early studies of pension beliefs (e.g. Bernheim, 1988, Manski, 2004) document misbeliefs about benefit levels. Caplin et al. (2022b) find substantial misbeliefs about eligibility ages in Denmark, similar to my findings in the UK. I use belief data to set initial conditions and identify a parameter from patterns in beliefs (patterns akin to Amin-Smith and Crawford (2018), prevalent misbeliefs predicting labor supply responses, and Rohwedder and Kleinjans (2006), errors decline as individuals age toward eligibility).

Structure of the paper. Section 2 provides background. Section 3 presents the data, and Section 4 descriptive and reduced-form analysis. Section 5 introduces the model, starting with a complete information baseline, then adding pension policy uncertainty and costly attention. Section 6 explains the solution method. Section 7 covers estimation. Section 8 discusses model fit and implications. Section 9 concludes.

2. BACKGROUND

The UK State Pension system has changed significantly since its introduction in 1948. I discuss the 2000-2016 system, especially post-2010, when the female SPA reform began.

State Pension benefit level. The UK State Pension comprises two parts: the Basic State Pension, based on contributing years, and a second tier, based on earnings, both calculated over working life. Working life is defined as spanning from the tax year an individual turns 16 to the year before they reach SPA (Bozio et al., 2010). So, benefit entitlement is frozen a year before SPA, meaning labor supply choices near SPA do not affect the pension amount.

The Basic State Pension began in 1948. By 2013, a full pension paid £107 per week ($203 in 2022 USD). Pro rata payments apply to those with fewer than 30 contributing years needed for the full pension. Contributing years include those in the labor force (earning above a minimum threshold) and spent caring for a child or disabled person post-1978. So, the timing of and reasons for labor market inactivity affect the pension amount.

The second tier of the State Pension began in 1978. Initially, it used an index-linked average of earnings between lower and upper limits over working life. Legislative changes resulted in varying accrual rates from 1978 to 2002, with a more progressive formula applied after April 2002. Thus, the timing of earnings affects second-tier entitlements. Private pension holders could also opt out in exchange for reduced payroll taxes.

Even in this simple outline, we see that due to protections for entitlements accrued under changing policies, the state pension benefit depends not only on total earnings and labor force participation but also on their timing and other factors (see Bozio et al., 2010, for details). Still, some general trends emerge. First, it is a relatively low benefit. It provides a 37% net

replacement rate for median earners, compared to 47%, 50%, and 58% in the USA, OECD, and EU, respectively. Second, it is a relatively flat-rate benefit. This is reflected in the larger drop in replacement rate between half and one-and-a-half times median earnings: 35 percentage points in the UK, versus 17, 21, and 14 in the USA, OECD, and EU (OECD, 2011).

State Pension Age and its reform. The State Pension Age (SPA) is the earliest age at which the State Pension can be claimed, serving as the UK's early retirement age. Deferring increased benefit generosity, but without a cap on deferral duration, hence implying no effective full retirement age.[4] So, the SPA is the sole focal age of the UK state pension system.

Unlike the State Pension amount, the SPA is a simple function of birth date and gender. The SPA was 65 for men and 60 for women until the Pensions Act 1995, which raised the female SPA from 60 to 65 incrementally, one month every two months, over ten years starting April 2010. The Pensions Act 2011 accelerated this change from April 2016, equalizing SPAs by November 2018, and legislated an increase for both genders to 66, phased in from December 2018. Figure 1a shows how these changes affected women by birth cohort. These reforms allow estimation of the risk that UK women face of SPA changes, a key model input. I also use variation from the 1995 reform (but not the 2011 reform, to avoid confounding from a change in benefit levels) to identify the SPA's impact on employment.

Communication and lack thereof. The government did not directly inform women affected by the reform, sending only the standard letter received by all pre-reform cohorts shortly before SPA. This lack of communication was controversial. From 2015, two campaign groups claimed the reforms discriminated against older women, with one unsuccessfully seeking to reverse the changes in the High Court. Their argument focused on the lack of communication. The government defended this by citing the absence of a national database in 1995, claiming direct notification was "essentially impossible". Reconciling this with letter-sending at SPA is beyond this paper's scope, but the absence of protests until 20 years after legislation supports the view that reported misbeliefs are genuine.

Private pensions. A large private pension market supplements the State Pension. Since private pension eligibility is not tied to SPA, it has little relevance to the employment response to SPA (more evidence in Appendix A.3.4).

Excess employment sensitivity and State Pension age. The UK SPA reform offers a unique opportunity to examine the excess employment sensitivity puzzle, as many common explanations for labor market exits at early retirement age are ruled out. First, UK law prohibits mandatory retirement based on age, banning it as age discrimination.[5] So, firm-mandated retirement cannot explain SPA employment sensitivity. Second, the state pension is not tied to employment status; individuals can claim it and continue working, and many do. Third, the UK pension system lacks tax incentives for labor market exits at SPA. Unlike the US system, there is no earnings test,[6] and while the state pension is taxable, a component of income tax, called National Insurance contributions, is removed at SPA.[7] Finally, it is worth restressing that

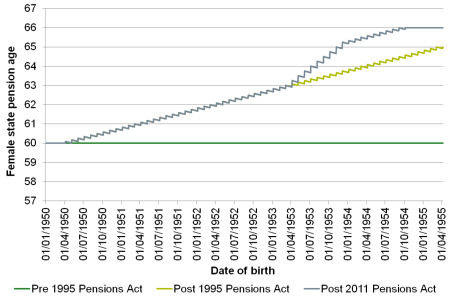
[4]Despite generous actuarial adjustments, deferral was rare, presenting a puzzle. Appendix F offers a model extension addressing this. Elsewhere, I abstract from the deferral puzzle, taking observed claiming as given.

[5]The Equality Act (2006) banned mandatory retirement below age 65, exceeding the highest SPA in this paper. The Equality Act (2010) extended the ban to all ages with exceptions in Appendix A.1.

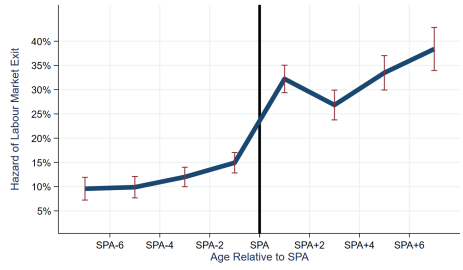
[6]Earnings tests penalize working while claiming retirement benefits, but they are *not* a feature of the UK system.

[7]Cribb et al. (2016) find changes to participation tax rates at SPA do not explain the employment response.

FIGURE 1.—Pension Legislation and Employment Response to the State Pension Age



(a) SPA by Date of Birth



(b) Fraction exiting labor employment

Note: Panel (a) shows State Pension Ages for women under the Pensions Act 1995, the Pensions Act 2007, and the Pensions Act 2011. Panel (b) plots the hazard of exiting employment at ages relative to SPA, with data plotted at two-yearly intervals to match ELSA's frequency.

benefit entitlement is frozen the year before SPA, making it unaffected by labor supply choices near SPA.

These facts show the State Pension acts as an anticipatable increase in non-labor income, with the SPA as the eligibility age. Announced in 1995 and starting in 2010, the reform provided at least 15 years of advance notice. The puzzle is not that employment responds to the reform, but the concentrated response at SPA despite the long notice period. In a standard life-cycle model with complete information and forward-looking agents, employment does not respond to anticipatable income changes unless liquidity constraints prevent intertemporal smoothing. Liquidity-constrained individuals cannot borrow against future pension income, forcing them to wait for this income to reduce labor supply.[8] Hence, what is puzzling is the concentrated employment response among women with substantial liquid wealth.[9]

3. DATA

Studying the employment response to the State Pension Age (SPA) requires a large sample of older individuals, and exploring its causes requires rich microdata. I use the English Longitudinal Study of Ageing (ELSA), as it is the UK[10] dataset best suited to these needs.

ELSA is a biennial panel dataset sampling the English population aged 50 and over, modeled on the US Health and Retirement Study (HRS). It provides rich microdata on labor-market conditions, earnings, and asset holdings. From wave three onward, ELSA collects data on SPA knowledge, crucial for studying misbeliefs. ELSA requests National Insurance numbers (equivalent to a US Social Security number) and consent to link administrative records, with 80% of respondents agreeing. These records improve pension-entitlement estimates, which are key for modeling SPA incentives. Survey data on health, education, and family further illuminate retirement motivations.

ELSA waves 1 (2002/03) through 7 (2014/15) cover those affected by the 1995 pension age reform and form the basis for analysis. The main sample includes women aged 55-75

[8]Loans using future pension benefits as collateral are not illegal but are not observed in practice.

[9]There was also a change in means-tested benefit eligibility at SPA. Individuals are no longer eligible for working-age unemployment benefits (Jobseeker's Allowance) but become eligible for the means-tested Pension Credit Guarantee. However, with the asset cut-off threshold used to study the excess employment sensitivity puzzle, the women would not qualify for Pension Credit Guarantee, and this richer subgroup are also not major users of job-search contingent Jobseekers allowance.

[10]ELSA (Banks et al., 2021) technically covers only England and Wales.

with 24,968 observations of 7,165 women. Different samples are used only when estimating particular model inputs, such as the spousal income process (dropping females, not males) or mortality process (including older ages). The female SPA reform began in 2010, making wave 5 the first post-reform wave. Earlier waves control for pre-trends and inform model inputs. The earliest affected cohort was born on 6 April 1950. Older cohorts serve as controls and also inform model inputs.

4. KEY MOTIVATING FACTS

4.1. *Excess Employment Sensitivity*

The sensitivity of employment to official retirement ages in excess of incentive is a puzzle observed in many countries (see Section 1). This section examines evidence of this puzzle for the UK SPA. The facts presented in this section are not novel, but the analysis places greater emphasis on demonstrating the puzzling nature of this employment response to SPA for standard complete information models, as this fact is central to the paper's thesis.[11] As liquidity constraints are the only standard complete-information mechanism for explaining SPA sensitivity (see Section 2), I focus on whether these constraints alone can account for the sensitivity of employment to the SPA. Readers happy to take this puzzling nature as established will lose little by skipping to the more novel facts on the relationship between misbeliefs and the employment response to the SPA in Section 4.2.

Figure 1b illustrates the excess employment sensitivity puzzle, showing the mean hazard rate of exiting employment by years from SPA. A sharp rise in exits at SPA is evident. While this is a correlation, the female SPA reform provides policy variation with which to estimate the SPA's causal effect on labor market exit.

To do this, I use a difference-in-difference approach, common in studies of employment responses to pension eligibility (e.g. Staubli and Zweimüller, 2013, Vestad, 2013, Atalay and Barrett, 2015, Cribb et al., 2016, Rabaté, 2019, Etgeton et al., 2023). The outcome is the hazard of exiting employment, which captures key transitions driving employment changes and is not contaminated by the level of employment, unlike using employment directly as the outcome. The main equation is:

$$y_{it} = \alpha \mathbb{1}\{\text{age}_{it} > SPA_{it}\} + \delta_{\text{age}_{it}} + \kappa_t + \gamma_{\text{cohort}_i} + \boldsymbol{\beta}' \mathbf{X_{it}} + \varepsilon_{it}, \tag{1}$$

where the hazard of exiting employment (y_{it}) of individual i, at time t depends on an indicator of being over the SPA ($\text{age}_{it} > SPA_{it}$); a set of quarterly age, period, and cohort dummies; and a vector of controls (X_{it}). The controls are marital status, years of education, highest qualification, self-reported health dummies, presence of a partner, partner's age, partner's age squared, partner's SPA eligibility, household non-housing non-business wealth (defined below), and a constant. I use this same set of controls throughout the regression in Section 4 and Appendix A. The hazard (y_{it}) is an indicator variable defined only if the individual was employed in the previous period; it takes the value 1 if they are no longer employed and 0 otherwise.

This form assumes cohort- and date-constant age effects, age- and date-constant cohort effects, and cohort- and age-constant date effects. Given these assumptions, which simply rephrase the parallel trends assumption, the parameter α is a difference-in-difference estimator

[11] Indeed, Cribb et al. (2016) conduct a similar analysis. They use the Labor Force Survey, which lacks sufficiently detailed asset information to rule out liquidity constraints as an explanation for the employment response to SPA. Since ruling out this explanation is key to the argument in this paper, I repeat some of their analysis in ELSA, which has richer asset information.

TABLE I

EFFECTS AND EFFECT HETEROGENEITY OF SPA ELIGIBILITY ON THE HAZARD

	(1)	(2)	(3)	(4)	(5)	(6)
Over SPA	**0.123**	**0.088**	**0.161**	**0.150**	**0.200**	**0.352**
s.e	(0.0235)	(0.0325)	(0.0348)	(0.0252)	(0.0459)	(0.0839)
Over SPA×Wealth Above Median	...	...	**-0.074**	...	...	...
s.e			(0.0487)			
Over SPA×Wealth in £100K	...	...	...	**-0.026**	...	...
s.e				(0.012)		
Over SPA× |SPA - SPA Self-report|	...	...	...	...	**-0.082**	...
s.e					(0.0378)	
Over SPA×SPA Above Self-report	...	...	...	...	...	**-0.207**
s.e						(0.0948)
Obs.	7,906	3,759	7,906	7,906	5,304	4,404

Note: The outcome variable is the hazard of exiting employment. Row 1 reports the treatment indicator for being over the State Pension Age (SPA). Rows 2–5 report interactions of this treatment indicator with: (2) an indicator for having above-median non-housing non-business wealth (NHNBW); (3) a continuous measure of NHNBW (in £100,000); (4) the absolute difference between the true and self-reported SPA; and (5) an indicator for the true SPA exceeding the self-reported SPA (i.e., under-estimation of the SPA). Self-reported SPAs are measured at age 58 or the closest available age, and wealth in the last interview before SPA. Each column reports estimates from a different specification of Equation 1: (1) the baseline regression on the full sample, (2) the baseline regression on an above-median NHNBW subsample, and specifications that fully interact the baseline equation with (3) being in the above-median NHNBW subsample, (4) the continuous NHNBW measure, (5) the absolute difference between the true and self-reported SPA, and (6) an indicator of under-estimating the SPA. As well as the age, period, and cohort dummies, all regressions control for marital status, years of education, highest qualification, self-reported health dummies, presence of a partner, partner's age, partner's age squared, partner's SPA eligibility, household non-housing non-business wealth, and a constant. Following Abadie et al. (2023), standard errors are clustered at the level of the treatment (i.e., birth cohort). Coefficients on controls and their interactions are not reported.

of the treatment of being above the SPA. The treatment is administered to all, but the reform induces variation in the duration of treatment. As a diagnostic for differential pre-trends I regress the hazard on age, period, and cohort dummies, the set of other controls, and interactions of the dummies, and I cannot reject cohort constant age effects or age constant cohort effects (p=0.376); age constant period effects or period constant age effects (p=0.144); or period constant cohort effect or cohort constant period effects (p=0.666).

Despite the well-known potential for bias of a staggered difference-in-difference (e.g., De Chaisemartin and d'Haultfoeuille, 2020), this simple difference-in-difference is preferred for the main text for ease of interpretation. Additionally, the final goal is to apply the same regression to simulated data as an auxiliary model during ex-post model validation, where bias is not a concern. As long as the same biased auxiliary model is used on both observations and simulated data, all that matters is the model's ability to replicate the results. However, Appendix A.3.3 addresses the potential for bias, allowing for heterogeneous treatment effects with the modern imputation method of Borusyak et al. (2024). Allowing for heterogeneity does not change the conclusion about SPA sensitivity in any important way.

Column (1) of Table I presents the results of estimating Equation 1. I find a 0.123 increase in the hazard of exiting work from being over SPA significant at the 0.1% level. This is in line with the findings in the literature. For example, studying the same reform using the Labor Force Survey Cribb et al. (2016) find an effect on employment of 6.3 percentage points with an average employment level in the range 41%-55%, implying an impact on the hazard rate of 11.5-15.4

percentage points.[12] To investigate if liquidity constraints explain the treatment effect, I restrict the sample to women from households with above-median non-housing non-business wealth (NHNBW)[13] in the wave before reaching SPA. The resulting threshold of £28,500 targets a group unlikely to face liquidity constraints affecting retirement choices. As the SPA was reformed in monthly increments and Equation 1 controls for quarterly age and cohort effects, the control group for estimating the treatment effect consists of individuals born in the same quarter but a few months younger, thus still below SPA. This narrow window strengthens the case against liquidity constraints: women with over £28,500 in NHNBW are unlikely to need to wait 1-3 months for the State Pension to stop working. Column (2) of Table I shows a treatment effect of 0.088 for this subgroup, similar to the full population and significant at 1%.

Column (3) of Table I encapsulates Columns (1) and (2) by estimating Equation 1 fully interacted with an indicator for being in the subpopulation used in Column (2) (i.e., allowing coefficients on all regressors in the specification in Equation 1 to differ for the subpopulation considered in Column (2)). Hence, the interaction with the treatment dummy tests whether there is a difference in treatment effects between those with above- and below-median assets. It is insignificant, showing no significant difference between groups. Dichotomizing assets into above and below median discards information, so Column (4) fully interacts Equation 1 with the continuous NHNBW variable in place of the subpopulation indicator. This interaction is significant but tiny: reducing the treatment effect by 1 percentage point requires an extra £38,460 in NHNBW. Indicating that, while wealth matters, its impact is too small for liquidity constraints to fully explain the SPA's effect on employment.

Table I captures the excess sensitivity puzzle in various ways, but a simple summary to test the model against is needed. While Column (4) provides finer-grained heterogeneity than Column (3), which consolidates Columns (1) and (2), Columns (1) and (2) more clearly embody the puzzle in two key findings: one, a significant employment response, which is, two, constant across a median asset split. So, I test the model against Columns (1) and (2).

Appendix A.3 provides robustness checks, including restricting to more liquid asset categories (Appendix A.3.1) and to alternative functional forms, such as dropping controls to address concerns about bad controls (Appendix A.3.2). These confirm that while assets influence the labor supply response to SPA, the effect is too weak for liquidity constraints to fully explain it. As mentioned, Appendix A.3.3 also relaxes the homogeneous treatment effects assumption using the modern imputation method of Borusyak et al. (2024).

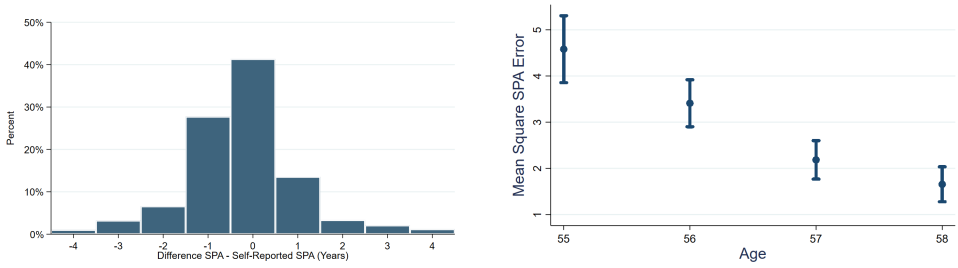
Appendix A.3.4 examines whether other factors like health, private pensions, or joint retirement explain the excess sensitivity and finds they do not. For example, in the case of joint retirement, Appendix A.3.4 examines whether the treatment effect of the SPA on the hazard of exiting employment varies with the presence of an out-of-work partner, and finds no significant heterogeneity along this dimension. Since partner employment does not change discretely at one's pension eligibility age and spillover effects of reaching pension age onto a partner's employment are an order of magnitude smaller than own effects[14], it is unsurprising that they play little role in explaining employment responses to one's own SPA. Similarly, the principal reason I find health and private pensions do not explain response to the SPA is that the SPA does

[12] Rabaté et al. (2024) calculate implied impacts on the hazard rate for various papers, including this one, and arrive at slightly different numbers because they start from observed drops in employment at pension age rather than the causal estimates of the impact on retirement.

[13] I use the designation NHNBW taken from the categorization in Carroll and Samwick (1996), but the variable may be better described as non-housing non-business non-pension wealth as it subtracts from ELSA's total non-pension wealth variable the value of primary residence and personal business assets. It includes current and savings accounts; shares; bonds; trusts; tax-privileged cash and equity saving instruments; other forms of land and property (excluding primary residences, farms, and personal businesses); and debt, whether credit card, personal, or other.

[14] See Lalive and Parrotta (2017), Johnsen et al. (2022), García-Miralles and Leganza (2024)

FIGURE 2.—SPA Beliefs



(a) Mistaken SPA Beliefs Age 58 (b) MSE in Self-reported SPA

Note: Panel (a) plots the frequency of errors in self-reported SPA at age 58 (binned to yearly accuracy). Panel (b) shows the mean squared error in Self-reported SPA plotted against respondents' age.

not significantly correlate with changes in either health or private pension eligibility status. The illegality of firm-mandated retirement has been used elsewhere as an empirical strategy to rule out its role in explaining employment response to pension eligibility (Rabaté, 2019). Of course, it is possible that firms illegally mandate retirement at SPA. Appendix A.3.4 provides some evidence against this, driving the result using self-declared reasons for employment termination.

The rest of this paper does not depend on the causal nature of the estimates presented in this section but uses them as an untargeted auxiliary model for a structural model. The key is the model's ability to replicate these results, not their causal nature. However, the analysis assumes that readers find these results puzzling under standard complete-information models. Placebo tests, in which I drop observations over SPA and replace the treatment in Equation 1 with indicators for aging past the dates one or two years before SPA, confirm with insignificant treatment effects that something specific is happening at SPA (full results in Appendix A.3.5), This is puzzling for those with substantial liquid wealth.

4.2. *Misbeliefs and Employment Sensitivity*

Compared to some other subjective belief data, such as inflation or survival expectations, pension beliefs have the distinctive feature that a currently correct answer exists. Hence, misbeliefs, in the sense of mistaken beliefs, are potentially observable. This section documents such misbeliefs about the SPA that are hard to reconcile with frictionless information, since people have clear incentives to know this information. It then investigates how these belief errors relate to observed retirement behavior. The core finding of this section is that discrepancies in self-reported SPA predict future retirement behavior in ways consistent with costly attention.

4.2.1. *Belief Data*

From wave three, ELSA asks respondents below SPA about their State Pension beliefs. This section focuses on the beliefs modelled in this paper, SPA beliefs, but ELSA records other State Pension beliefs. Beliefs about the State Pension amount were collected only in wave three; therefore, they lack statistical power and do not provide a panel dimension. The binary measure of basic awareness of the occurrence of the SPA reform shows that the vast majority know a reform has occurred, making this variable not very discriminating and indicating that a lack of awareness does not drive SPA belief errors. So, SPA beliefs are the most interesting of ELSA's belief data about the State Pension. A fact which motivated the decision to focus

on this aspect of pension policy uncertainty (more details about the other belief variables in appendix A.7).

ELSA has collected subjective SPA beliefs from individuals under SPA since Wave 3, resulting in five waves of data spanning almost a decade for the present purpose.[15] One feature that shapes the interpretation of these data is that ELSA elicits only point estimates for SPA beliefs, which, as noted by De Bruin et al. (2023), pose interpretation challenges. If individuals hold subjective priors, it is not clear which measure of central tendency the answer reflects or if it represents something else entirely. To operationalize the model, Section 7 provides a specific interpretation of these point estimates as a draw from the subjective belief distribution. The analysis in this section, however, does not rest on this particular mapping from individual belief distributions to responses, only that responses correlate with people's mean subjective SPA belief.

Thus, in this section, I treat SPA-belief interview responses ($SPA^r_{i,t}$) as a measure of individuals' expectations of their currently announced SPA ($SPA_{i,t}$) as implied by their subjective beliefs ($\tilde{E}_{i,t}[SPA_{i,t}]$), contaminated by measurement error ($\varepsilon_{i,t}$). In plainer terms, responses are a noisy measure of what individuals believe their currently announced SPA is. Measurement error is primarily included to account for the aforementioned uncertainty in interpreting responses.

The question asked in the ELSA survey is: "Do you know at what age in years and months you will reach the State Pension Age?" As well as being subject to measurement error, this question conflates the probability of an actual reform occurring between when the person is asked and when they reach SPA with their subjective uncertainty. If people respond precisely to this question, they would account for any expected increase in the SPA between the interview and their reaching SPA. SPA reforms, however, are observable events, and so we know the frequency with which they occur. Hence, we can correct for the probability of a reform occurring between the interview and the respondent reaching SPA.[16] If we assume that an individuals true $SPA_{i,t}$ follows a non-decreasing stochastic process with constant mean excepted increment ρ, then we can correct for the resulting drift to recover their subjective expected beliefs ($\tilde{E}_{i,t}[SPA_{i,t}]$) about the current value of their SPA with the following formula:[17]

$$(1-\rho)(SPA^r_{i,t} - age_{i,t}) + age_{i,t} = \tilde{E}_{i,t}[SPA_{i,t}] + \nu_{i,t} \tag{2}$$

where $\nu_{i,t} = (\rho - 1)\varepsilon_{i,t}$. The drift ($\rho$) is estimated in Section 7.1 to be 0.102. Accordingly, by accounting for the expected drift resulting from a potential reform between the time of interview and eligibility, we can use Equation 2 to infer a noisy measure of the individual's belief about their current SPA ($\tilde{E}_{i,t}[SPA_{i,t}]$) from their reported age of reaching eligibility ($SPA^r_{i,t}$).

As the SPA is an exact function of date of birth and gender, both recorded in ELSA, we can examine discrepancies between self-reported and an individual's currently announced SPA.

[15] Although this panel of belief data is substantial, it is a slight loss of sample size compared to Section 4.1 owing to missing responses in Waves 1-2 and among women over SPA.

[16] This assumes the individuals know the probability of reform but not the current value. It is equally possible that they know the current value but not the probability of it changing. The source of their confusion (current value or law of motion) is not separately identifiable; only the fact that they must be mistaken about one or both. Given this, I load uncertainty in this paper onto the current value, but I do not believe the implications are greatly changed if they are instead mistaken about the law of motion.

[17] The age at which someone reaches the SPA represents the stopping time $A_i = \inf_t\{age_{i,t} \geq SPA_{i,t}\}$. So if response are noisy measures of current subjective expectations of arriving at SPA, $SPA^r_{i,t} = \tilde{E}_{i,t}[A_i] + \varepsilon_{i,t}$ and $SPA_{i,t}$ a non-decreasing stochastic process with constant drift ρ, we can apply the additive drift theorem to get: $\tilde{E}_{i,t}[A_i] = age_{i,t} + \frac{\tilde{E}_{i,t}[SPA_{i,t}] - age_{i,t}}{1-\rho}$.

Figure 2a plots the difference between true and raw self-reported SPAs for reform-affected women at age 58, the last age when no cohort has received an SPA communication, or the closest age interviewed. Although the modal group reports their SPA correctly to within a year, this includes many mistakes by a margin of months, and the majority (58.7%) are off by a year or more. Can these discrepancies be explained by the objective probability of a SPA reform occurring? The answer is no. Firstly, since more people underestimate their SPA than overestimate it, correcting for the objective probability of an increase in the SPA makes the discrepancies on average larger, since $\rho > 0$. Secondly, even if we ignore the fact that this correction moves self-reports in the wrong direction, the size of the change is much smaller than required to account for these discrepancies, with the mean correction for individuals in the sample being just 0.12 years.

So, the objective probability of a reform cannot account for the observed discrepancies, which are instead best interpreted as a combination of belief errors and measurement noise. While the lack of a suitable instrumental variable makes it difficult to isolate measurement error fully, the discrepancies in self-reported SPA are predictive of future retirement behavior (Section 4.2.2 below), indicating that they reflect meaningful belief variation rather than purely noise. Additionally, Appendix A.6 shows individuals whose self-reported SPA increases save more. Hence, SPA belief data are also predictive of saving behavior in ways consistent with a face-value interpretation as reflective of individuals' subjective beliefs. Accordingly, I interpret drift-corrected self-reports as a noisy measure of individuals' mean belief about their current SPA, and the discrepancy between these and the true SPA as a noisy measure of misbelief, whilst accounting for measurement error as robustly as possible given the lack of an instrument.

These measured misbeliefs are not only prevalent but also show traits consistent with costly information, such as learning. Learning over time is likely under costly information acquisition as knowledge is retained, and the instrumental value of knowing your SPA rises with age. Figure 2b supports this, showing a decline in mean squared errors of self-reported SPAs as women age toward their SPA. The model uses these declining errors to identify the attention cost.[18]

4.2.2. *Relation to Employment Sensitivity*

This paper's model of endogenous SPA knowledge makes two distinct predictions about the relationship between SPA misbeliefs and the employment response at the SPA. Firstly, as SPA knowledge is endogenous, selection implies that larger absolute belief errors correlate with smaller employment responses to the SPA. This is because if the SPA is irrelevant to an individual's actions, she will choose not to learn it or respond upon reaching it. Secondly, because overestimating the SPA implies a positive wealth shock upon learning the actual value, it predicts that an individual who overestimates their SPA will, all else equal, have a larger employment response at the SPA than a comparable individual who underestimates. Other mechanisms are also at play at the SPA,[19] so the model does not necessarily predict a null or negative response among under-estimators, but it does predict a smaller one.

These two predictions are driven by distinct groups: those who care about their SPA (prediction two) and those who do not (prediction one), which are not readily observable. Since these

[18] Additionally, Appendix A.4 shows that these discrepancies are not driven by people just reporting the pre-reform age or some focal age as the self-reports cluster around each cohort's true SPA, consistent with a costly attention model. The appendix also details self-report errors at their natural monthly frequency and belief heterogeneity by years of education.

[19] Information revelation at SPA also implies a positive information shock reducing precautionary labor supply.

predictions sometimes conflict and are driven by groups whose membership is unobserved, each prediction may obscure attempts to detect the other, potentially leading to aggregation bias. For example, if we observe an individual who massively over-predicts their SPA but does not react upon reaching it, is this evidence against the second prediction, or is this confirmation of the first?

The data robustly support the first prediction. Column (5) of Table I shows the results of Equation 1 fully interacted with the absolute error in self-reported SPA at age 58 or the nearest age observed (i.e., interacting all regressors in Equation 1 with the absolute error). The significant negative interaction suggests that for each additional year of error in SPA self-reporting, the employment response drops by 8.2 percentage points. So, those least informed about the SPA before age 60 have the smallest employment response upon reaching SPA after 60. This pattern aligns with a model of endogenous costly information acquisition: individuals who place lower value on knowing their SPA acquire less information and exhibit smaller behavioral responses. In contrast, such selection would not arise if self-reports were purely measurement error or if information arrived exogenously.

Since in this specification we are taking a non-linear function (the absolute value) of the measurement-error-contaminated variable, the direction of the induced bias is not guaranteed (as discussed in Schennach, 2016). Hence, it is particularly important to consider the role of measurement error in this context. Appendix A.5 repeats this analysis, using an alternative measure of the size of the mistake, the square of misbeliefs, with which the measurement error always attenuates results. Appendix A.5 also reverses this regression using the observed employment response to the SPA as a proxy for incentives to learn about your SPA to study selection more directly. These robustness exercises support selection into SPA knowledge driving a correlation between the employment response upon reaching SPA and the size of the SPA misbelief earlier in life.

The data provide stronger support for the second prediction when accounting for the aggregation bias introduced by the mixture of individuals who select out of SPA knowledge by restricting the sample. Specifically, I exclude individuals with an absolute error above 5 quarters, as those who select out of SPA knowledge are more likely to be found among individuals with larger absolute errors. Column (6) of Table I shows treatment effect heterogeneity according to whether individuals over- or under-predict their SPA at 58 or the closest age observed, resulting from fully interacting Equation 1 with an indicator for under-estimating. In agreement with prediction two, it shows a significantly larger response to the SPA among those who overestimate their SPA. Without the sample restriction, the point estimate still aligns with prediction two but does not reach statistical significance (details in Appendix A.5). With the threshold variable used here, measurement error will attenuate the results toward zero (Aigner, 1973), perhaps helping to explain why significant support for the second prediction is sensitive to specification decisions.

Appendix A.5 replicates the analyses from columns (5) and (6) without applying the drift correction (Equation 2) to SPA self-reports, and additionally estimates a regression that accounts for treatment heterogeneity along both the size and direction of belief error dimensions simultaneously. The results of these robustness exercises are broadly consistent with the findings here: strong support for prediction one, and noisier, though directionally consistent, support for prediction two.

Recent work (e.g., Seibold (2021), Lalive et al. (2023)) has made important progress on the puzzle of excess employment sensitivity by modeling reference-dependent preferences. These models abstract from the role of beliefs and assume full information about policy. By contrast, this paper emphasizes that individuals may be uncertain about their eligibility due to attention frictions and that this uncertainty contributes to the retirement response. The results in this section, which document misbeliefs and heterogeneity in employment responses based on belief

accuracy, provide evidence consistent with this new channel. Although reference dependence as an explanation for employment responses to pension ages does not require perfect information, it does not account for misbeliefs, making information frictions a more unifying explanation. Additionally, the fact that misbeliefs are predictive of employment responses to pension ages is not naturally generated by a framework in which employment responses stem from fixed preference parameters. Furthermore, the simple fact that people are mistaken about their eligibility age, as found here and by Caplin et al. (2022b), may cast some doubt on the idea that they view it as a key reference point. That said, there is interesting evidence in favor of reference-dependence (discussed in Appendix E.3), and whether these explanations are complementary or whether one can account for all observations is a question left for future research. This paper presents evidence that endogenously generated misbeliefs are also part of the story, as they can explain patterns in the belief data and how these relate to employment responses to the SPA. Section 8.2 provides further comparison between these two explanations.

Regarding the external validity of the mechanism, it is worth emphasizing that although I use the occurrence of the reform to identify variation, the proposed mechanisms rely solely on pension misbeliefs and the potential for reform. Appendix A.8 documents similar employment and misbelief patterns for men, who were not subject to a reform, consistent with the view that this misbelief channel exists in the absence of a reform.

5. MODEL

Section 5.1 presents the baseline standard complete information model. Section 5.2 introduces two additions: objective uncertainty about government pension policy and costly information acquisition about this uncertain policy.

5.1. *Complete Information Baseline*

Key features are summarized before diving into details. The model's decision-making unit is a household containing a couple or a single woman, but when a husband is present, his labor supply is inelastic. The household maximizes lifetime utility from bequests, leisure, and equivalized consumption by choosing consumption, labor supply, and savings. Households face risk over i) whether they get an employment offer, ii) the wage associated with any offer, and iii) mortality. The households receive non-labor income from state and private pensions after the relevant eligibility age for each.

In more detail, households are divided into four types indexed by k, based on the high or low education status of the female and the presence or absence of a partner. Periods are indexed by the age of the female (t). Each period, households choose how much to consume (c_t), how much to invest in a risk-free asset (a_t) with return r, and, if not involuntarily unemployed, how much of the women's time endowment (normalized to 1) to devote to wage labor ($1 - l_t$) (40, 20 or 0 hours per week) at a wage offer (w_t) that evolves stochastically. Unemployment (ue_t), where $ue_t = 0$ indicates employment (presence of a wage offer) and $ue_t = 1$ unemployment (the absence), also evolves stochastically. The partner's labor supply is inelastic, and so his behavior is treated as deterministic. Although doing so abstracts from leisure complementarities and joint retirement, Appendix A.3.4 shows that the partner's labor supply is not significant in explaining labor supply responses to one's own SPA, which is the focus of this model. The wife receives the state pension once she reaches the SPA (SPA), a parameter varied to mimic the UK reform, and a private pension once she reaches the type-specific eligibility age ($PPA^{(k)}$), which is a type-specific parameter that does not vary with SPA. Both pensions, $S^{(k)}(.)$ the state pension and $P^{(k)}(.)$ the private pension, are treated as type-specific functions of average

lifetime earning ($AIME_{t+1} = \frac{(1-l_t)w_t + AIME_t t}{t+1}$) [20]. From age 60, the women face a probability of surviving the period (s_t^k). Finally, households value bequests through a warm glow bequest function (De Nardi, 2004). The full vector of model state is $X_t = (a_t, w_t, AIME_t, ue_t, t)$.

Utility. The warm glow bequest motive creates a terminal condition ($T(a_t)$) that occurs in a period with probability $1 - s_{t-1}^{(k)}$:

$$T(a_t) = \theta \frac{(a_t + K)^{\nu(1-\gamma)}}{1-\gamma}$$

where θ determines the intensity of the bequest motive, and K determines the curvature of the bequest function and hence the extent to which bequests are luxury goods. The functional form around $a_t + K$ is the household's utility from consumption (see below), so the warm-glow bequest approximately captures the utility a descendant gains from these assets, and hence altruism as a motive, whilst keeping parameters to a minimum.

Whilst alive, a household of type k has the following homothetic flow utility:

$$u^{(k)}(c_t, l_t) = n^{(k)} \frac{((c_t/n^{(k)})^\nu l_t^{1-\nu})^{1-\gamma}}{1-\gamma}$$

where $n^{(k)}$ is a consumption equivalence scale that takes the value 2 if the household is a couple and 1 otherwise. In other words, utility takes an isoelastic form, with curvature γ, over a Cobb-Douglas aggregator of consumption and leisure, with consumption weight, ν.

Initial and terminal conditions. ELSA interviews people from 50, but the model starts with women aged 55 because this is the youngest age with significant numbers of SPA self-reports for multiple SPA-cohorts, thus allowing me to initialize state variables (a_t and $AIME_t$ but in the model with incomplete information, also beliefs) from the empirical distributions for different SPA-cohorts. At age 100, the woman dies with certainty.

Labor market. The female log wage (w_t) is the sum of a type-specific deterministic component, quadratic in age, and a stochastic component:

$$\log(w_t) = \delta_{k0} + \delta_{k1} t + \delta_{k2} t^2 + \epsilon_t \tag{3}$$

where ϵ_t follows an AR1 process with persistence ρ_w and normal innovation term with standard error σ_ϵ, and has an initial distribution $\epsilon_{55} \sim N(0, \sigma^2_{\epsilon,55})$. The quadratic form of the deterministic component of wages captures the observed hump-shaped profile and is common in the literature.

The unemployment status of the woman (ue_t) evolves according to a type-specific conditional Markov process. From 80, the woman can no longer choose to work; this is to model some of the limitations imposed by declining health. As spousal income results from the confluence of wages, mortality, and pension income, it follows a flexible polynomial in age:

$$\log(y^{(k)}(t)) = \mu_{k0} + \mu_{k1} t + \mu_{k2} t^2 + \mu_{k3} t^3 + \mu_{k4} t^4$$

This specification averages out and abstracts away from both idiosyncratic spousal income and mortality risk. In effect, the household dies when the woman dies, and the husband's mortality

[20] This is average yearly earnings, to keep notation in line with the literature I use the abbreviation Average Indexed Monthly Earnings, which is the variable US Social Security depends on.

risk only turns up insofar as it affects average income, as if husbands were a pooled resource amongst married women of type k. This allows me to ignore transitions between married and single, which, while important for the wider labor supply behavior of older individuals (e.g. Casanova, 2010), are of secondary importance to employment responses to the SPA. The function $y^{(k)}(t)$ amalgamates spousal labor and non-labor income, including pensions. Both female wage and spousal income are post-tax.

Social insurance. Unemployment status is considered verifiable, so only unemployed women ($ue_t = 1$) can claim the unemployment benefit (b).

The wife receives the state pension as soon as she reaches the SPA, which abstracts away from the benefit-claiming decision. This is done for two reasons, both touched upon earlier. Firstly, over 85% of people claim the State Pension at the SPA, so, in terms of accuracy, little is lost by this simplification. Secondly, this small fraction deferring receipt occurs despite deferral having been actuarially advantageous during the period studied. This presents another puzzle for standard models of complete information, which generally imply acceptance of actuarially advantageous offers. In Appendix F, I provide preliminary evidence that costly attention to pension policy may also help explain this claiming puzzle, but in the main text, I abstract from it. Abstracting from it here gives the baseline model a chance to solve the excess sensitivity puzzle.

Lifetime average earnings ($AIME_t$) evolve until the woman reaches the age she starts to receive her private pension ($PPA^{(k)}$), at which point it is frozen. Both state and private pensions are quadratic in $AIME_t$ until they reach their maximum, at which point they are capped. Until being capped, the pension functions have the following forms

$$S^{(k)}(AIME_t) = sp_{k0} + sp_{k1}AIME_t - sp_{k2}AIME_t^2$$

$$P^{(k)}(AIME_t) = pp_{k0} + pp_{k1}AIME_t - pp_{k2}AIME_t^2$$

These pension functions abstract from the details of state and private pension systems but capture some key incentives in a tractable form. The state pension is a complex, path-dependent function shaped by past and current regulations (see Bozio et al., 2010). This functional form captures the dependence of the state pension on working history without getting into these difficulties. Being type-specific allows $S^{(k)}(.)$ to capture indirect influences of education and marital status on the state pension; for example, being a stay-at-home mum counted towards State Pension entitlement (after the enactment of a reform). Every private pension scheme is different, but the dependence of $P^{(k)}(.)$ on $AIME_t$ reflects the dependence of most defined benefit schemes on lifetime earnings. This functional form less accurately reflects the structure of defined contribution systems, which are essentially savings accounts. However, the model captures retirement savings through a risk-free asset and starts after the statutory defined contribution eligibility age, beyond which they can be accessed without penalty.

Total deterministic income. Combining spousal income, benefits, and private and state pension benefits into a single deterministic income function yields:

$$\begin{aligned} Y^{(k)}(t, ue_t, AIME_t) = y^{(k)}(t) + b\mathbb{1}[ue_t = 1] + \mathbb{1}[t \geq SPA]S^{(k)}(AIME_t) \\ + \mathbb{1}[t \geq PPA^{(k)}]P^{(k)}(AIME_t) \end{aligned}$$

Household maximization problem. The Bellman equation for a household of type k is:

$$V_t^{(k)}(X_t) = \max_{c_t, l_t, a_{t+1}} \left[u^{(k)}(c_t, l_t) + \beta \left(s_t^{(k)} E[V_{t+1}^{(k)}(X_{t+1})|X_t] + (1 - s_t^{(k)}) T(a_{t+1}) \right) \right]$$

subject to the following budget, borrowing, and labor supply constraints:

$$c_t + (1+r)^{-1} a_{t+1} = a_t + w_t(1 - l_t) + Y^{(k)}(t, ue_t, AIME_t), \tag{4}$$

$$a_{t+1} \geq 0, \tag{5}$$ &

$$ue_t(1 - l_t) = 0. \tag{6}$$

5.2. *Two Additions: Policy Uncertainty and Costly Attention*

This section adds two features to the complete information model. Section 5.2.1 introduces objective policy uncertainty in the form of a stochastic SPA, reflecting potential SPA variation over the lifecycle caused by pension reform. Section 5.2.2 adds costly attention to the stochastic SPA, in the form of disutility for more precise information, allowing the model to capture misbeliefs. These additions are introduced independently, resulting in three model versions: the baseline from Section 5.1, a version with policy uncertainty and informed households, and the full model with rationally inattentive households.

5.2.1. *Policy Uncertainty: the Stochastic SPA*

To capture the objective policy uncertainty resulting from the fact that governments can and, sometimes, do change pension policy, I make the SPA stochastic.

Although the SPA does change, introducing an important dimension of uncertainty, changes are not sufficiently frequent to estimate a flexible stochastic SPA process. For this reason, I impose a parsimonious functional form on the stochastic SPA:

$$SPA_{t+1} = \min(SPA_t + e_t, \overline{SPA}) \tag{7}$$

where $e_t \in \{0, 1\}$ and $e_t \sim Bern(\rho)$. So each period, the SPA may stay the same or increase by one year, as the shock is Bernoulli, up to an upper limit of $\overline{SPA} = 67$. This captures a key aspect of pension uncertainty, that in recent years governments have reformed pension ages upward but generally not downward, whilst maintaining a simple tractable form. The lowest SPA, I consider possible is the pre-reform age of 60. Hence, as the law-of-motion only allows for increases, SPA_t is bounded below by $\underline{SPA} = 60$ and above by $\overline{SPA} = 67$.

In the model, the variable SPA_t represents the current best available information about the age the woman will reach her SPA, and as such, the data analog is the SPA the government is currently announcing for the woman's cohort. Only one SPA cohort is modeled at a time. So there is no conflict in having a single variable SPA_t whilst, in reality, at a given point in time, different birth cohorts have different government-announced SPAs.

5.2.2. *Costly Attention (Rational Inattention)*

The second addition is a cost of information acquisition about the stochastic SPA. This addition allows the model to reflect misbeliefs, and the modeling choices made in modeling costly information acquisition are motivated by the patterns seen in the belief data discussed in Section 4.2 and Appendix A. Firstly, we observed that although individuals are very mistaken about their own SPA, the vast majority are aware that a reform has taken place (Appendix A.7). Hence, to capture the pattern of misbeliefs, it is important to model the intensive margin of

information acquisition rather than the extensive margin of awareness or unawareness alone. Secondly, Section 4.2 showed that selection into SPA knowledge was important for explaining the relationship between misbeliefs and the employment response to the SPA. Hence, the process of information acquisition must, at least in part, be a choice that allows beliefs to respond to the incentives to be informed. The rational inattention paradigm fits these desiderata very well. It allows highly flexible information acquisition, capturing both intensive and extensive margin choices, including the extremes from zero to exhaustive information acquisition. It is a paradigm in which the chosen information is an optimal response to incentives and constraints, but which also naturally captures the unpredictable, random nature of information arrival, since the information acquisition strategy is stochastic. This section lays out the technical details of how introducing consistent, incentive-driven, Bayesian learning, in the form of rational inattention, changes the model.

Directly observed vs learnable states. As the goal of this paper is to explain the employment response at SPA, the stochastic SPA (SPA_t) is the only state variable subject to a cost of information acquisition. The SPA_t is not directly observed by the household. Instead, learning about its value requires paying a utility cost (defined below). In contrast, other stochastic states such as wages (w_t) and unemployment status (ue_t) are observed directly. This can be understood as these variables being more salient. To aid the exposition of information acquisition, these directly observable states are combined into a single vector of salient states $X_t = (a_t, w_t, AIME_t, ue_t, t)$.[21] Additionally decisions are grouped into a single variable: $d_t = (c_t, l_t, a_{t+1})$.

Within period timing of learning. Since the household does not directly observe SPA_t, it is a hidden state. It is still a state as it is payoff-relevant, but since the household does not observe it, it cannot enter the decision rule. Instead their decision will depend on what they believe the SPA to be, introducing a new state variable: the household's belief distribution over possible SPA values, denoted $\underline{\pi_t} = (\pi(spa))_{spa=\underline{SPA}}^{\overline{SPA}} \in \Delta(8) \subseteq \mathbb{R}^8$.

Each period, the household chooses an information strategy that specifies how they will acquire information about SPA_t, then, using the information they receive, they make their labor supply and savings decisions. That is, first, the household selects a signal distribution or information strategy ($\underline{f_t}[X_t, \underline{\pi_t}](z \mid SPA_t)$), then, conditional on the signal drawn from this distribution ($z_t \sim Z_t$), it chooses actions ($d_t[X_t, \underline{\pi_t}](z_t)$). The household can chooses as signal any valid probability distribution without additional constraint,[22] allowing the model to flexibly capture relationships between SPA_t and beliefs. This flexibility can capture the partial informedness that we see in the data, but also allows for (in the guise of a degenerate signal) the idea of looking up and remembering your SPA. Hence, it captures the extensive margin (e.g., looking up and remembering your SPA) as well as the intensive margin learning (e.g., learning from friends) that leads to the partial informedness we see in the data.

Although this choice is unconstrained, households suffer a utility cost for choosing more informative signals (discussed below). Due to this information-acquisition cost, households will generally choose noisy signals, not because being perfectly informed is ruled out, but because it is rarely optimal. Hence, the modeling device naturally reflects both the responsiveness of information acquisition to incentives (one of the key desiderata) and the random component of learning that is outside our direct control of any individual.

[21] This is the same collection of variables in X_t as when it was defined in the baseline model. I highlight this as a change because X_t has not absorbed the new state SPA_t and hence no longer contains all states.

[22] The household makes this choice having observed the other state variable which is why the function choose depends on X_t and $\underline{\pi_t}$

Since the information strategy ($\underline{f_t}$) specifies how the household acquires information not what information they acquire (z), the choice of information strategy can be thought of as the active choice of conditions under which to pay attention to the randomly arriving information from diverse source outside their control like older siblings, colleagues, or the media. For example, the appearance of news stories about pension reforms is beyond an individual's control, but whether to continue reading past the headline is an active choice. Hence, this modeling device reflects many features of the messy real-world learning process.

Bayesian learning. Beliefs are updated using Bayes' rule (starting at an initial belief distribution ($\underline{\pi_{55}}$)). The posterior after observing signal z_t is:

$$Pr_t(spa \mid z_t) = \frac{f_t(z_t \mid spa)\pi_t(spa)}{Pr_t(z_t)} = \frac{f_t(z_t \mid spa)\pi_t(spa)}{\sum_{spa'=60}^{\overline{SPA}} f_t(z_t \mid spa')\pi_t(spa')} \tag{8}$$

Next period's prior is then formed by applying the law of motion for SPA_t (Equation 7) to the posterior:

$$\pi_{t+1}(spa) = \begin{cases} (1-\rho)Pr_t(spa \mid z_t) + \rho Pr_t(spa-1 \mid z_t) & spa < \overline{SPA} \\ Pr_t(spa \mid z_t) + \rho Pr_t(spa-1 \mid z_t) & spa = \overline{SPA} \end{cases}. \tag{9}$$

This assumes households know the law of motion of SPA_t, but not its current value. As mentioned in Section 4.2, we cannot separately identify mistakes about the law of motion from mistakes about the current value, and this paper loads mistakes onto the current value. This assumption is not viewed as central, and I conjecture that, qualitatively, the conclusions would be robust to loading misbeliefs into the law of motion. What is key is that information frictions generate SPA misbeliefs, not which part of the data-generating process of SPA_t is subject to those frictions.

Entropy and mutual information. The cost of attention is modeled using tools from information theory. Entropy (in the information-theoretic sense) measures uncertainty: it captures the minimal space needed to store or transmit the information in a random variable. Mutual information measures how much this uncertainty is reduced by learning another variable. These concepts will define the utility cost of information acquisition (below).

DEFINITION—Entropy: Let $X \sim P_X(x)$. The entropy of X is the expected value of $-\log(P_X(x))$:

$$H(X) = \mathbb{E}_X[-\log(P_X(x))]$$

DEFINITION—Conditional entropy: The conditional entropy of X given Y is defined as:

$$H(X \mid Y) = \mathbb{E}_Y[H(X \mid Y = y)]$$

DEFINITION—Mutual Information: Let X and Y be random variables. The mutual information between them is the expected reduction in uncertainty about X from learning Y:

$$I(X,Y) = H(X) - H(X \mid Y)$$

In this model, the utility cost of attention is proportional to the mutual information between the belief over SPA and the chosen signal capturing information acquired during that period. At the end of this section, I discuss reasons for this choice of functional form.

Utility with attention costs. The household faces a trade-off between more informed decision-making and the cost of acquiring that information. On the one hand, learning about the SPA improves savings and labor supply choices. On the other, acquiring more precise information comes at a utility cost. So, the value of information is the instrumental value of making better saving and labor supply choices, while its cost is a direct utility cost. This trade-off is captured in the period utility function:

$$u^{(k)}(d_t, \underline{f_t}, \underline{\pi_t}) = n^{(k)} \frac{((c_t/n^{(k)})^{\nu} l_t^{1-\nu})^{1-\gamma}}{1-\gamma} - \lambda I(\underline{f_t}; \underline{\pi_t}) \tag{10}$$

Here, $I(\underline{f_t}; \underline{\pi_t})$ is the mutual information between the belief $\underline{\pi_t}$ and the chosen signal $\underline{f_t}$, and λ is the cost of attention.[23] Expanding the mutual information and combining terms gives:

$$I(\underline{f_t}; \underline{\pi_t}) = \sum_{spa} \sum_{z} \pi_t(spa) f_t(z|spa) \log \frac{f_t(z|spa)}{\sum_{spa'} \pi_t(spa') f_t(z|spa')}$$

What does this cost of information acquisition represent? While your SPA is a single number freely available online, looking it up does not capture the full costs of learning it. These should include information processing, storage, and recall costs, as well as straightforward hassle or time costs. For illustration, the author has paid the hassle cost of looking up his SPA but not the cognitive cost of remembering it. Hence, I would show up in survey data as having SPA misbeliefs, and I cannot use my SPA in decision-making. Thus, the minimum data- and model-consistent conceptualization includes both cognitive and hassle costs. Given that the time cost of looking up one's SPA can be measured in seconds and yet misbeliefs are prevalent in the data, it seems likely that the cognitive cost of information processing is more important than the time cost, consistent with the modeling assumption of a utility rather than a time cost.

Of course, the pension system is multifaceted, and people find many facets confusing, whilst the model concentrates all costs of information acquisition on tracking the SPA. Thus, it is possible that the cost of information may also capture learning and the resolution of uncertainty about the pension policy more broadly. An extension in Appendix F explores household learning about actuarial adjustment for deferred claiming.

Revelation of uncertainty. In practice, SPA eligibility was communicated to individuals in the UK via letter shortly before SPA, and the claiming process involved a phone call that explicitly clarified the implications of claiming. The model reflects these institutional features by resolving uncertainty upon pension receipt. Upon reaching SPA_t, the woman learns her true SPA_t and starts receiving the state pension. So, the household knows that if they do not receive the woman's state pension benefits, she is below her SPA. This also ensures households never unknowingly exceed their budget.

Dynamic programming problem. The household's state now includes both the belief over SPAs and the latent true SPA

$$(X_t, SPA_t, \underline{\pi_t}) = (a_t, w_t, AIME_t, ue_t, t, SPA_t, \underline{\pi_t}),$$

[23] This paper abstracts from preference parameters heterogeneity, and so in particular misses heterogeneity in information acquisition cost. Although more educated people almost certainly have lower costs of information acquisition, as shown in Appendix A.5, they are actually more mistaken about their SPA. This indicates that heterogeneity in incentives to learn, which are captured by the model and are larger for the less educated, is more important in this context.

The Bellman equation becomes:

$$V_t^{(k)}(X_t, SPA_t, \underline{\pi_t}) =$$
$$\max_{d_t, \underline{f_t}} E\Big[u^{(k)}(d_t, \underline{f_t}, \underline{\pi_t}) + \beta\big(s_t^{(k)} V_{t+1}^{(k)}(X_{t+1}, SPA_{t+1}, \underline{\pi_{t+1}}) + (1 - s_t^{(k)})T(a_{t+1})\big)\Big] \quad (11)$$

subject to the same constraints in Equations 4 - 6 as the baseline model, and where now the utility function includes an information cost as per Equation 10. As previously stated, the choice objects ($\underline{f_t}[X_t, \underline{\pi_t}](z \mid SPA)$ and $d_t[X_t, \underline{\pi_t}](z_t)$), are restricted to be functions of the observed state, and the hidden state (SPA_t) appears in the value function only because expectations and continuation value are conditioned on it.

A challenge buried in this Bellman equation is the formation of next-period beliefs, which, due to Bayesian updating, depend upon the full distribution of the signal. Hence, we need the solution to form the continuation value. This problem is taken up in Section 6.

Functional form of attention cost. The information acquisition cost is key to the model mechanisms. I assume it is proportional to the expected entropy reduction for three reasons.

Firstly, a cost of information acquisition that is directly proportional to mutual information is among the most common in the costly information literature, leading to two important advantages. It is tractable as many useful results are available for this functional form,[24] and it follows a convention. Tractability is important in models of costly information, which can become too complex to solve, and following a convention has merit because it restricts the degrees of freedom available to fit the data.

Secondly, as argued by Mackowiak et al. (2018), this functional form offers a disciplined behavioral model by replicating numerous empirically supported departures from classical models. It endogenously generates behaviors that look like heuristics, or rules-of-thumb, observed often enough to be christened biases in the behavioral literature.[25]

Thirdly, reasons exist to believe that the cost of cognition depends on entropy. The information-theoretic concept of entropy sets a lower bound on efficient transmission and storage of information. Thus, if the brain processes information efficiently, mutual information should factor into the ideal cost of attention function. This is not to say that an ideal cost of attention function would be linear in mutual information, and recent works such as Caplin et al. (2022a) generalize the traditional entropy penalty in multiple ways. Laboratory evidence (e.g. Dean and Neligh, 2023, Bronchetti et al., 2023) indicates that the entropy-based cost of attention omits features of human attention that other cost functions better capture, although other experiments find it adequate (e.g. Fuster et al., 2022). Outside of such a controlled setting, however, it is not always clear which departures from the entropy-based costs are most relevant or whether sufficient data variation exists to identify their extra parameters. As it seems that entropy enters an ideal cost function, my cost function can be considered a first-order approximation over this dimension.

6. MODEL SOLUTION

By introducing a high-dimensional state $\underline{\pi_t}$ (beliefs) and a high-dimensional choice $\underline{f_t}$ (signal), rational inattention has complicated the model. Indeed, solving a dynamic rational inatten-

[24] Until Miao and Xing (2024) extended results from Steiner et al. (2017) to universally posterior separable function, we only knew how to solve the dynamic rational inattention model with entropy-based cost of attention.

[25] For example, Kõszegi and Matějka (2020) show this attention cost generates mental budgeting (quantity allocated to a category being fixed and composition changing) and naive diversification (composition being fixed and quantity allocated changing) in different situations. Caplin et al. (2019) show it leads to consideration sets.

tion model without suppressing the belief distribution as a state variable represents an important contribution of the paper. To achieve this, I combine theoretical results into a general-purpose solution method for dynamic rational inattention models with history-dependent beliefs, such as the one in this paper. This solution method addresses computational challenges by first filtering out actions that will never be taken. In my application, this filtering reduces the number of potential choices by two orders of magnitude and often finds a solution without solving the optimization problem.

The solution method can be considered general-purpose because, one, it stores the belief distribution non-parametrically, and two, it does not rely on any specifics of the data-generating process. The only substantive restriction it imposes on the class of dynamic rational inattention models with this entropy-based cost of attention is that the problems must be discrete choice so that it falls within the class solved by Steiner et al. (2017). Since any computational method requires some degree of discretization, discretizing a problem can be seen as a computational approximation. Due to this restriction, I discretize the assets and labor supply choices. Section 6.1 explains the general-purpose method, and Section 6.2 details specific to solving the model of this paper.

6.1. *Solving Dynamic Costly Attention Models with History-dependent Beliefs*

This section presents a solution method for dynamic rational inattention models with history-dependent beliefs. I use the model of retirement decision from this paper to explain the method, but it applies to any dynamic, discrete-choice, rational-inattention model. Section 6.1.1 outlines key results from Steiner et al. (2017). Section 6.1.2 uses these results and presents the method.

6.1.1. *Analytic Foundations of Solution Method*

Steiner et al. (2017) shows that a wide class of models has logit-like solutions. Specifically, the class of dynamic discrete choice models subject to an additively separable entropy penalty for information acquisition about the state variable, of which the model presented in this paper is an example. The key results from their paper that are needed to understand the solution method are explained below. The results are general, but I use my model to explain them to avoid extra notation and concepts.

If we define the effective conditional continuation value (i.e., the expected continuation value conditional on taking a given action) as:

$$\overline{V}^{(k)}_{t+1}(d_t, X_t, SPA_t, \underline{\pi_t}) = \\ E\left[s^{(k)}_t V^{(k)}_{t+1}(X_{t+1}, SPA_{t+1}, \underline{\pi_{t+1}}(d_t)) + (1 - s^{(k)}_t)T(a_{t+1}) \middle| d_t, X_t, SPA_t, \underline{\pi_t}\right], \tag{12}$$

where expectations are over X_{t+1} and SPA_{t+1} (Section 6.1.2 belows describes finding $\underline{\pi_{t+1}}(d_t)$), then the Bellman equation 11 becomes:

$$V^{(k)}_t(X_t, SPA_t, \underline{\pi_t}) = \max_{d_t, \underline{f_t}} E\left[u^{(k)}(d_t, \underline{f_t}, \underline{\pi_t}) + \beta \overline{V}^{(k)}_{t+1}(d_t, X_t, SPA_t, \underline{\pi_t})\right].$$

Steiner et al. (2017) show the optimal information acquisition strategy is to receive an action recommendation as signal, which results in a one-to-one mapping from signals to actions. The intuition for this is that, since information is costly and the only benefit of information is making better choices, gathering any information not directly reflected in a choice would be wasteful. The linearity of information costs in mutual information is key to this one-to-one mapping;

without it, it is possible that delaying or advancing information acquisition might decrease the cost of information. Using this mapping, we can substitute actions for signals and the conditional choice probabilities ($d_t|SPA_t \sim \underline{p_t}(.|SPA_t)$) for the signal function ($\underline{f_t}$) throughout the problem, since one is simply a relabeling of the other. Thus, we can combine the choice of a stochastic signal function ($\underline{f_t}$) and a deterministic decision conditional on the signal ($d_t(z_t)$) into a single choice of a stochastic decision ($d_t|SPA_t \sim \underline{p_t}(.|SPA_t)$).

With the problem now re-expressed in terms of choice probabilities, it is a simple calculus exercise to solve for the optimal distribution of action which they show has actions that are distributed with conditional choice probabilities $d_t|SPA_t \sim \underline{p_t}(.|SPA_t)$ and associated unconditional probabilities $d_t \sim \underline{q_t}(.)$ (i.e., $q_t(d) = \sum_{spa=\underline{SPA}}^{\overline{SPA}} \pi(spa) p_t(d|spa)$) that satisfy:

$$p_t(d|spa) = \frac{\exp\left(n^{(k)} \frac{\left(\left(\frac{c}{n^{(k)}}\right)^{\nu} l^{1-\nu}\right)^{1-\gamma}}{\lambda(1-\gamma)} + \log(q_t(d)) + \frac{\beta}{\lambda} \overline{V}_{t+1}^{(k)}(d, X_t, SPA_t, \underline{\pi_t}) \right)}{\sum_{d' \in \mathcal{C}} \exp\left(n^{(k)} \frac{\left(\left(\frac{c'}{n^{(k)}}\right)^{\nu} l'^{1-\nu}\right)^{1-\gamma}}{\lambda(1-\gamma)} + \log(q_t(d')) + \frac{\beta}{\lambda} \overline{V}_{t+1}^{(k)}(d', X_t, SPA_t, \underline{\pi_t}) \right)}, \tag{13}$$

$$\underline{q_t} = \arg\max_{\underline{q}} \sum_{spa} \pi_t(spa) \log\left(\sum_{d \in \mathcal{C}} q(d) \exp\left(n^{(k)} \frac{((c/n^{(k)})^{\nu} l^{1-\nu})^{1-\gamma}}{\lambda(1-\gamma)} + \frac{\beta}{\lambda} \overline{V}_{t+1}^{(k)}(d, X_t, SPA_t, \underline{\pi_t}) \right) \right). \tag{14}$$

The logit-like structure captures how choice probabilities reflect not only utility differences but also information costs. In contrast to the better-known random-utility multinomial logit, the source of stochasticity is noise in the information, which is not pushed to zero because it would be costly to do so. This distinction is reflected by the presence of the endogenous objects $q_t(d)$ on the right-hand side of Equation 13 and leads to substantive predictive differences between the two logit foundations discussed in Matějka and McKay (2015). The solutions to the discrete choice rational inattention model are so similar to solutions to a model with EV1 preference shocks because the EV1 distribution maximizes entropy within the class of distribution that achieves a given expected utility. So, it is the least costly distribution for the agent to achieve a utility level.

6.1.2. *General-Purpose Solution Method*

At its core, the solution method is to solve Equation 14 for $\underline{q_t}$ and substitute the solution into 13 to get $\underline{p_t}$. This basic description corresponds to an infeasible brute-force version of my solution method. The problem with this brute-force version is that solving the high-dimensional optimization problem in Equation 14 at every point in the state space, which contains the high-dimensional state $\underline{\pi_t}$, is computationally infeasible. My solution method leverages the fact that there are actions that a rationally inattentive agent will never take for which there are conditions we can check (such as being strictly dominated) that are less computationally costly than the optimization in Equation 14. My algorithm uses these conditions to filter out actions that will never be taken before solving Equation 14. High-level pseudocode summarizing the algorithm, which may help follow the discussion that follows, is in Appendix C. To get to a complete description of the algorithm, two hurdles must be passed.

The first hurdle is that knowing which belief next period will result from an action this period requires knowing the full probability distribution of actions. This follows because we do not know how strong a signal an action is for a given SPA unless we know how likely households are to take that action, given other possible SPAs. It follows that the conditional effective continuation value ($\overline{V}_{t+1}$) is not known, even though next period's value function (V_{t+1}) is known, because we do not know the beliefs tomorrow that will result from an action

today ($\underline{\pi_{t+1}}(d_t)$), which, as a state, enters V_{t+1}. To see this, substitute the distributions of actions for the distribution of signals in the Bayesian updating formula 8 and apply the results from Equations 13 and 14 to get:

$$Pr(spa|d_t) = \frac{\pi_t(spa)\exp\left(n^{(k)}\frac{((c/n^{(k)})^{\nu}l^{1-\nu})^{1-\gamma}}{\lambda(1-\gamma)} + \beta\overline{V}^{(k)}_{t+1}(d, X_t, spa, \underline{\pi_t})\right)}{\sum_{d'\in\mathcal{C}} q_t(d')\exp\left(n^{(k)}\frac{((c'/n^{(k)})^{\nu}l'^{1-\nu})^{1-\gamma}}{\lambda(1-\gamma)} + \beta\overline{V}^{(k)}_{t+1}(d', X_t, spa, \underline{\pi_t})\right)}.$$

Then the prior at the start of next period ($\underline{\pi_{t+1}}$) is formed by applying the law of motion of SPA_t (Equation 7) to this posterior as per 9. That is:

$$\pi_{t+1}(spa) = (1-\rho)Pr_t(spa|d_t) + \rho Pr_t(spa-1|d_t).$$

Thus, beliefs given choices ($\underline{\pi_{t+1}}(d_t)$) are a function of the posterior, which depends not only on the exponentiated payoff but also on $\underline{q_t}$. So, we need a solution ($\underline{q_t}$) to know $\underline{\pi_{t+1}}(d_t)$ and hence to form the effective conditional continuation values (Equation 12).

Steiner et al. (2017) evade this difficulty by removing the beliefs from the state space and replacing them with the full history of actions. They can do this because, given initial beliefs, the full history of signals, or equivalently actions, perfectly predicts the beliefs in period t. This is an inspired step in their proof that extends Matějka and McKay (2015) to the dynamic case, as it allows them to show we can ignore the dependence of continuation values on beliefs. For applied structural modeling, it is often a non-starter because it introduces redundant information into the state space. If two action histories lead to the same beliefs, they do not truly represent different states.[26] Redundant information in the state space is problematic, as the curse of dimensionality often makes this the binding constraint to producing richer models. That the redundant information grows exponentially with the number of periods moves this from problematic to a non-starter for many applications.

Hence, I rely on the theoretical results of Steiner et al. (2017), which used the history of action state-space representation, but in practice, I use the more compact belief state-space representation for the actual computational work. To get around the issue that I need $\underline{q_t}$ to know $\overline{V}_{t+1}$ to work with the belief state-space representation, I use a simple guess-and-verify fixed-point strategy. First, I guess a value $\underline{\tilde{q}_t}$ and use this to form next period beliefs conditional on having taking an action ($\underline{\pi_{t+1}}(d_t)$) using Equations 8 and 9. With these, I form the effective conditional continuation value ($\overline{V}_{t+1}$) using Equation 12. Then given $\overline{V}_{t+1}$ I solve 14 for $\underline{q_t}$. If the resulting $\underline{q_t}$ is sufficiently close to $\underline{\tilde{q}_t}$, I accept this solution otherwise I replace $\underline{\tilde{q}_t}$ with $\underline{q_t}$ and repeat.[27]

By increasing the computation required at each state, this solution to the first hurdle, however, exacerbates the second: the high computational demands resulting from the high-dimensional state $\underline{\pi_t}$. Previously, models of dynamic rational inattention have generally avoided this problem by suppressing the belief distribution as a state variable (Miao and Xing, 2024, Armenter et al., 2024, Turen, 2023, Macaulay, 2021, Porcher, 2020).[28] Although potentially reasonable

[26] In Steiner et al. (2017), past actions can affect beliefs and current utility. Hence, two histories leading to the same belief might represent truly different states. This is not the case here.

[27] Although I have not proved this is a contraction mapping, the fixed point iteration always converges and generally in relatively few iterations.

[28] Sometimes this is justified as an explicit information sharing assumption in the model. Often, it is justified by noting that local posterior invariance (Caplin et al., 2022a) extends to global posterior invariance if all actions are taken with positive probability. However, Caplin et al. (2019) show that solutions are rarely strictly interior as rational inattention often implies consideration sets. Hence, extending local posterior invariance to a global property is restrictive.

in specific applications, suppressing beliefs prevents dynamic rational inattention from modeling situations in which beliefs matter and vary across individuals, as is the case for pension beliefs in the UK. Heterogeneous subjective beliefs that impact choices have been documented in portfolio selection (Giglio et al., 2021), school choice (Kapor et al., 2020), and compliance with health guidelines (Bhalotra et al., 2025). Thus, suppressing beliefs as a state variable imposes real limitations on the domain of applicability of rational inattention. Moreover, these applications could potentially benefit from the dynamically consistent Bayesian model of incentive-driven belief formation offered by rational inattention.

My solution method keeps the belief distribution as a state whilst leveraging results of Caplin et al. (2019) to lighten the computational burden. They show that often rational inattention implies consideration sets. Hence, the solving conditional choice probabilities (CCPs) $\underline{p_t}$ are sparse. That is, households take various actions with zero probability. This sparsity reflects decision-makers' ability to rule out some actions as clearly not beneficial without further information. To illustrate this idea, a household with total cash-on-hand of £10,000 can immediately rule out saving £9,999 without gathering any information about their SPA.

I propose two criteria that ex ante identify actions that will be taken with zero probability, without solving the optimization problem. I then remove these from the decision problem. This filtering step always reduces the dimensionality of the optimization in Equation 14. Moreover, if a single action remains after filtering, we have solved the problem without further calculation. For my model, filtering leaves a single action in over 50% of cases.

The first and simplest criterion for culling actions is removing strictly dominated alternatives. The agent is rationally inattentive and so will never select an action strictly dominated in all possible realizations of the SPA.[29] Hence, all actions strictly dominated across all realizations of SPA_t can be removed. Checking this first criterion is helpful at two points in the procedure. Firstly, before making an initial guess for $\underline{\tilde{q}_t}$, by removing any actions strictly dominated across all possible *joint* combinations of SPA_t realization and belief on the belief grid $\underline{\pi_{t+1}}$. Doing this before entering the loop that solves for $\overline{V}_{t+1}$ reduces unnecessary computational burden in that fixed point iteration for $\underline{q_t}$. However, it imposes a much stricter condition, dominant across all joint realizations of $\overline{SPA_t}$ and $\underline{\pi_{t+1}}$, than needed to drop an action, dominant across all realizations of SPA_t. Therefore, having made an initial guess for $\underline{\tilde{q}_t}$, and so having prediction for next period beliefs given any action ($\underline{\pi_{t+1}(d_t)}$) and hence the conditional continuation value, I secondly remove actions strictly dominated across all realizations of SPA_t for each belief in the belief grid as I loop over beliefs to find the policy rule at those beliefs. I do this for each belief during each iteration of the loop that solves for $\overline{V}_{t+1}$.

In my model, the dimension reduction achieved by dropping strictly dominated actions is large, often two orders of magnitude. Abstracting from borrowing constraints, the household faces 1,500 options, 500 saving levels, and 3 labor supply choices. A household will never assign positive probability to more actions than the random variable they are learning about (SPA_t) has points of support. SPA_t has two points of support at the age of 65, increasing to 8 at age 59. Hence, filtering often reduces the initial choice set in the high hundreds (once

[29]Thinking about the wider applications of this solution method, rational inattention does not rule out choices that appear ex post strictly dominated, only those strictly dominated across all possible states of the world. For example, if a firm offers a multi-dimensional product like a mortgage that depends on the state of the economy, the rationally inattentive agent will never choose a product that is worse on all dimensions in all states of the world. They may choose a product that is ex post strictly dominated across all dimensions in the realized state of the economy because they assign some positive probability to the state being other than it actually is and believe that the company offers optimal products in that state. Hence, although choices appear strictly dominated ex post, they are not strictly dominated ex ante.

we impose borrowing constraints) to single digits. The runtime required to perform a single filtering is negligible compared to the runtime required to solve Equation 14.

Removing strictly dominated actions only uses ordinal preferences. The second criterion used to filter also uses the cardinal preferences encoded in expected utility. It exploits the necessary and sufficient condition from Caplin et al. (2019). Using these, it is easily shown (see Appendix B.1) that if there exists a decision $d^\star = (c^\star, l^\star)$ which satisfies:

$$\sum_{spa} \pi_t(spa) \frac{\exp\left(n^{(k)} \frac{((c/n^{(k)})^\nu l^{1-\nu})^{1-\gamma}}{\lambda(1-\gamma)} + \frac{\beta}{\lambda} \overline{V}_{t+1}^{(k)}(d, X_t, spa, \underline{\pi_t})\right)}{\exp\left(n^{(k)} \frac{((c^\star/n^{(k)})^\nu l^{\star 1-\nu})^{1-\gamma}}{\lambda(1-\gamma)} + \frac{\beta}{\lambda} \overline{V}_{t+1}^{(k)}(d^\star, X_t, spa, \underline{\pi_t})\right)} < 1, \tag{15}$$

for all other decisions $d = (c, l)$ then it is the only action taken ($q(d^\star) = 1$). Unlike dropping strictly dominated alternative, which reduces the dimensionality, making solving Equation 14 easier, checking Equation 15 is only beneficial when the optimal behavior is to take the same action in all realizations of SPA_t. So, the benefits of checking condition 15 depend on how frequently, in the problem faced, it reveals the optimal choice without needing to solve an optimization. When filtering does not leave a single action, I employ sequential quadratic programming to solve Equation 14, an algorithmic choice suggested by Armenter et al. (2024).

Appendix C details two other computational difficulties. Firstly, the large state space significantly increases the storage requirements for the solutions. With this issue, the sparsity proved by Caplin et al. (2019) is again helpful as I can use sparse matrix storage techniques. Secondly, when λ is small, Equation 14 can lead to underflow.

6.2. *Computational Details Specific to this Model*

All versions of the model (the baseline, with policy uncertainty but informed households, and with rationally inattentive households) are solved by dynamic programming, specifically backward induction. Beliefs ($\underline{\pi_t}$) and learning ($\underline{f_t}$) alter the nature of the within-period problem in the version with rationally inattentive households in some periods. Only in some periods because $\underline{\pi_t}$ and $\underline{f_t}$ are only relevant before the SPA. After the SPA, the true value is known, and so beliefs ($\underline{\pi_t}$) and learning ($\underline{f_t}$) about the SPA are irrelevant. Periods after the SPA are solved, like periods in the other two versions, by simple search techniques to find the optimal choice from the discrete set of assets and labor supply choices.

In the version with rationally inattentive households, we proceed by backward induction from terminal age $t = 100$ using standard techniques for the within-period problem until age $t = 66$. We can proceed back as far as age $t = 67$ because SPA_t is bounded above by 67, so the woman receives her state pension with certainty from that age onward. Standard methods can also solve the period $t = 66$ because, at this age, the household is perfectly informed. Either she has reached her SPA and policy uncertainty has been resolved, or she infers $SPA_t = 67$ with certainty, as she knows the data-generating process. In this period, $\underline{\pi_t}$ is not a state variable, but SPA_t is, as receipt of the state pension affects available resources.

At all earlier ages (t <66), if $SPA_t \leq t$, then uncertainty has been resolved, meaning the model can still be solved using standard techniques, and, moreover, the exact value of SPA_t is irrelevant. All that matters to the household is that they receive the benefit, so we can solve for a single representative SPA with $SPA_t \leq t$. Conversely, when the SPA is in the future ($SPA_t > t$), the agent cannot infer the true value of the SPA, and so both the agent's beliefs ($\underline{\pi_t}$) and the true value of the SPA (SPA_t) are states and the agents needs to choose a learning strategy ($\underline{f_t}$). Each year we proceed backward, the list of future potential SPAs ($SPA_t > t$) grows by one, increasing the combinations of $\underline{\pi_t}$ and SPA_t for which we need to solve a

problem with uniformed learning agents that is not solvable by simple search techniques. As $\underline{\pi_t}$ is a distribution over all future SPAs, its points of support also grow by one with each step in the backward induction. For example, at age $t = 65$, there are two potential future SPAs (66 and 67), and if SPA_t takes on either of these values, the agent can no longer infer its true value, and so beliefs ($\underline{\pi_t}$) become a state and the choice of signal function relevant. This growth of problem complexity along two related dimensions, rational-inattention-relevant potential future SPAs and the size of the belief distribution over them, continues until we reach $t = 59$. At this point, all SPAs 60-67 are future, and rational inattention is relevant regardless of the value of SPA_t and the support of $\underline{\pi_t}$ is fixed.

7. ESTIMATION

The model is estimated by two-stage simulated method of moments. The first stage estimates, outside the model, parameters of the exogenous driving processes and the initial distribution of state variables (a small number of parameters are also set, drawing on the literature). Using the results from the first stage, the second stage estimates the remaining preference parameters $(\beta, \gamma, \nu, \theta, \lambda)$ using the simulated method of moments.

7.1. *First Stage*

The parameters of the wage process, the state and private pension system, and the unemployment transition matrix are estimated outside the model. The curvature of the warm-glow bequest and the interest rate are taken from the literature.

Initial conditions. To set the initial conditions of the model, I need values for a_t, w_t, $AIME_t$, ue_t, and in the version with rationally inattentive households $\underline{\pi_t}$. Initial wages w_t are drawn from the estimated initial wage distribution (see below), and all agents start as employed ($ue_t = 0$). Beliefs ($\underline{\pi_t}$) are initialized from the type- and SPA-cohort specific empirical distribution, and assets (a_t) and average earnings ($AIME_t$) from their joint type- and SPA-cohort specific empirical distribution. The empirical counterpart used for assets is household non-housing non-business wealth. Using the full work histories in the administrative data linked to wave 5 of ELSA, I construct a measure of $AIME_t$. As this is only possible for a subsample, to estimate the joint distribution of $AIME_t$ and a_t, I impute missing $AIME_t$ values with a quintic in wealth and a rich set of observed characteristics (details in Appendix D). To initialize beliefs from the point-estimate belief data, I assume that responses represent a draw from an individual's subjective beliefs distribution.[30] Given this assumption and imposing a homogeneous prior among individuals of a given type and cohort, the cross sectional distribution will be the same as an individuals first period prior. Thus when simulating an individual of a given type and SPA cohort I set $\underline{\pi_{55}}$ equal to the empirical discrete distribution of self-reported SPAs at age 55.

Wage equation. I assume wage data is contaminated with serially uncorrelated measurement error ($\mu_{j,t}$), leading to the following variant of Equation 3 as data generation process:

$$\log(w_{j,t}) = \delta_{k0} + \delta_{k1} t + \delta_{k2} t^2 + \epsilon_{j,t} + \mu_{j,t}$$

[30] This assumption is consistent with evidence from psychology that averaging multiple responses elicited from an individual improves accuracy (Vul and Pashler, 2008). It also enables the construction of an individual's subjective belief distribution from point estimates.

for women j, of type k, and at age t. The parameters of the age-dependent deterministic component of the wage process $(\delta_{k0}, \delta_{k1}, \delta_{k2})$ are estimated by type-specific regression. The parameters of the stochastic component of the wage equation $(\rho_w, \sigma_\epsilon, \sigma_{\epsilon,55}, \sigma_\mu,)$ are found minimizing the distance between the empirical covariance matrix of estimated residuals and the theoretical variance-covariance matrix of $\epsilon_t + \mu_{j,t}$ (similar to Low et al., 2010). I correct for selection into employment following French (2005) by repeating the wage regression in the model, adjusting the model's wage equation parameters until the regression on the model and the data agree.

Pension systems. Both pensions are type-specific functions of average lifetime earnings. These are estimated on the $AIME_t$ measures constructed from the administrative data described above. As the state pension is relatively insensitive to education and the private pension to marital status, I simplify the state pension to be marital-status-specific and the private pension to be education-specific. I estimate the private pension claiming age $(PPA^{(k)})$ as the type-specific mean earliest age women are observed with private pension income.

Unemployment transition matrix. I classify a woman as unemployed if she claims an unemployment benefit and estimate type-specific transition probabilities in and out of unemployment.

Stochastic State Pension age. I estimate the probability of an increase in the SPA, ρ, on the cumulative changes to the original female SPA of 60 experienced by reform-affected cohorts. That is, I select the ρ that minimizes the mean error in SPAs, given that the data-generating process is Equation 7.

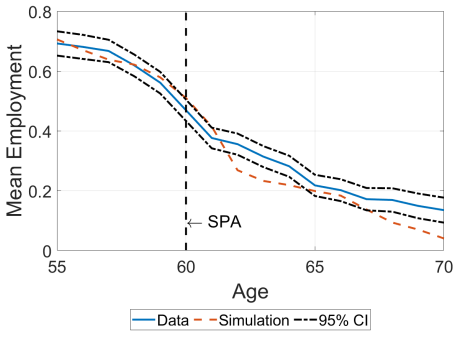
Parameters set outside the model. The curvature of the warm-glow bequest is taken from De Nardi et al. (2010) and the interest rate from O'Dea (2018). Prices are deflated to 2013 values using the RPI. Survival probabilities are taken from the UK Office for National Statistics life tables and combined with ELSA data to estimate type-specific survival probabilities following French (2005), details in Appendix D.

7.2. *Second Stage*

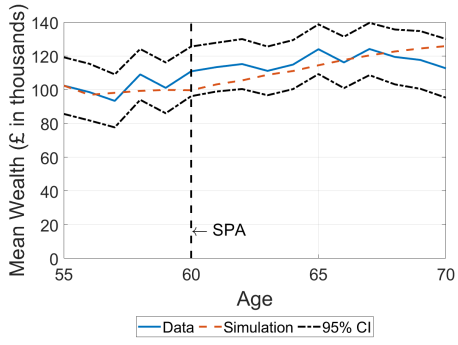
In the second step, moments are matched to estimate the preference parameters: the isoelastic curvature (γ), the consumption weight (ν), the discount factor (β), and the bequest weight (θ), as well as the cost of attention (λ) in the version with costly attention.

The 42 pre-reform moments of mean labor market participation and asset holdings from ages 55 to 70 were used to estimate $(\beta, \gamma, \nu, \theta,)$. To avoid cohort effects or macroeconomic influences, a fixed-effect age regression was estimated, including birth-year effects, SPA-cohort-specific age effects, aggregate unemployment (to half a percentage point), and an indicator for being below the SPA. Target profiles were then generated using these regressions with average pre-reform cohort values (details in Appendix D). In the model version with rationally inattentive households, λ is identified from an additional moment: the reduction in the mean squared error of self-reported SPAs between ages 55 and 58 for the same pre-reform SPA-cohort as other targeted moments $(SPA = 60)$. Since beliefs at 55 are initialized from the data (see Section 7.1), the fit in that period is mechanical (a slight undershooting results from discretizing beliefs). Hence, only beliefs at age 58 are targeted to identify λ, with beliefs at the two intervening ages (56 and 57) being untargeted. As self-reported SPAs are assumed to reflect draws from a subjective belief distribution, I generate a random draw from each simulated individual's beliefs at age 58 $(\underline{\pi}_{58})$ and calculate the MSE of simulated self-reported SPAs using the distribution of these draws.

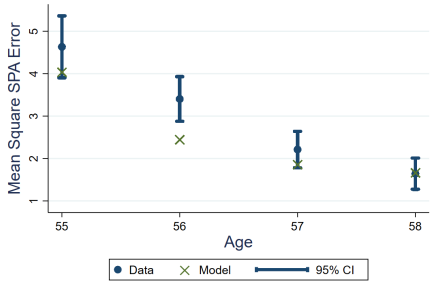
FIGURE 3.—Model Fit and Parameter Estimates



(a) Employment Profile



(b) Asset Profile



(c) Belief Profile

TABLE

ν: Consumption Weight	0.439 (0.0000038)
β: Discount Factor	0.985 (0.0000004)
γ: Relative Risk Aversion	3.291 (0.0000170)
θ: Warm Glow bequest Weight	100 (2.626)
λ: Cost of Attention	6×10^{-8} (2×10^{-11})

(d) Parameter Estimates

Note: Panels (a)-(c) show model fit to targeted profiles, the empirical profile is for the pre-reform SPA cohort with a SPA of 60. Panel (d) shows estimated parameters (analytic standard errors in brackets calculated following Newey (1985)).

8. RESULTS

Section 8.1 evaluates model fit and ability to replicate key facts on excess employment sensitivity, misbeliefs, and their relationship. Section 8.2 compares these results to those of the reference-dependence preference retirement literature Section 8.3 explores the model implications.

8.1. *Model Evaluation*

This section presents the model fit and each version's ability to replicate the employment responses to the SPA and their relation to beliefs. The objective probability of a pension reform occurring in any given year is estimated to be $\hat{\rho} = 0.102$. Other, less novel, first-stage parameter estimates are in Appendix E.1.

Figures 3a and 3b show the fit of the version with policy uncertainty to the pre-reform employment and asset profiles when simulated with the pre-reform SPA of 60. Since the baseline version and the version combining policy uncertainty with rational inattention produce similar fits to these static profiles, their profiles are placed in Appendix E.2. The mean square error of model-predicted and data beliefs used to identify λ are presented in Figure 3c. The value estimated of this parameter is $\hat{\lambda} = 6 \times 10^{-8}$, and Table 3d reports all second-stage parameter

TABLE II

UNTARGETTED MODEL FIT TO REGRESSION RESULTS

	(1) Baseline	(2) Policy Uncert.	(3) $\hat{\lambda} = 6 \times 10^{-8}$	(4) $\lambda = 5.0 \times 10^{-7}$	(5) Data (95% C.I)
Treatment Effect being above SPA on employment					
Whole Population	0.019	0.014	0.041	0.068	0.123 *(0.0766,0.1696)*
Assets >Median(£28,500)	0.018	0.014	0.054	0.088	0.088 *(0.0230,0.152)*
Treatment Effect Heterogeneity by Absolute SPA Error					
Interaction	—	—	-0.047	-0.046	-0.082 *(-0.1577,-0.0061)*
Treatment Effect Heterogeneity by SPA Error Positivity					
Interaction	—	—	-0.047	-0.046	-0.207 *(-0.3973,-0.0168)*

Note: The top panel shows employment response across the wealth distribution (Table II). The second panel shows heterogeneity in SPA labor supply response by the absolute size of self-reported SPA error at 58. The second panel shows heterogeneity in SPA labor supply responses by direction of self-reported SPA error at 58, and the third panel by the absolute size of the error. Some results are identical to three decimal places but differ to four decimal places.

estimates.[31] Although the three model versions achieve very similar fits to the static employment profiles, the three versions predict very distinct responses to SPA changes.

To analyze this response to the SPA, I simulate the model with the SPAs observed in ELSA waves 1-7 ($SPA = 60$, $SPA = 61$, $SPA = 62$) and repeat the regression from Section 4.1 on the simulated data. I adapt Equation 1 to the model's simpler environment, estimating the treatment effect of being above the SPA on the hazard of exiting employment using a two-way fixed-effects difference-in-differences approach. This regression includes the treatment indicator, full age, and cohort fixed effects (excluding period effects, which align with age in the model), and model counterparts to empirical controls (assets, marital status, and education). As in Section 4.1, I repeat this on the subsample with above-median empirical assets (£28,500) before SPA. Results are in Table II 's top panel. Column (5) repeats the empirical treatment effects from Columns (1) and (2) of Table I. The baseline model fails to match either.

This baseline's failure reflects the excess employment sensitivity puzzle, which prompted this investigation into policy uncertainty and costly attention. To assess their impacts separately, I introduce them sequentially. Column (2) shows that policy uncertainty alone has no effect. This is because objective uncertainty is low (SPA changes are rare). Both this version and the baseline fail to match the treatment effects for the whole population and for those with above-median assets at SPA, but are closer to the lower response of the richer subgroup.[32]

Column (3) of Table II shows that the model version with rationally inattentive households matches the employment response to the SPA significantly better than the baseline or the policy uncertainty versions, but still falls short of the data. Costly attention closes 21% of the gap for the whole population and 51% for the richer subgroup, with only the richer subgroup's estimate falling within the 95% confidence interval. So, the introduction of rational inattention

[31] These are the estimated parameters obtained from estimating the version with rational inattentive households. To separate the impact of changes in other parameters from the introduction of rational inattention, I have also run an exercise in which I hold other parameters constant and estimate λ from the additional moment. The results are available on request but change very little, as misbeliefs mostly strongly affect λ anyway.

[32] Section 4.1 highlights the ex-ante puzzling response of the wealthy, and targeting the two treatment effects directly allows the baseline to match the overall population response but not the wealthy subgroup's (results available on request). Thus, I consider the wealthy group's response puzzling, though the baseline struggles most with the aggregate at the estimated parameters.

improves the model's ability to account for employment responses to the SPA in an economically significant way. This improvement is achieved without directly targeting these responses or introducing free parameters, as data-driven restrictions are imposed on the added model features: policy uncertainty is estimated from past SPA reforms and information friction from misbeliefs.

Is the increase in employment response from introducing rational inattention reasonable? The introduction of rational inattention increases the impact of the SPA on the probability of exiting employment by 2.7 percentage points, and this increase is driven by two potential mechanisms: precautionary labor supply and wealth effects triggered by the resolution of pension uncertainty. Upon reaching SPA, households in the model overestimate their own SPA by 0.7 years, on average, implying that the wealth shock they receive is roughly equivalent to learning they have an extra year of pension income. More precisely since average yearly state pension benefits are £4,691, the implied wealth shock is £3,284. Then assuming an otherwise constant hazard, the 2.7 percentage point increase in hazard implies an expected decrease in employment of 0.18 years of employment. If this increase in hazard is driven purely by wealth effect, this implies, at the mean, a marginal propensity to earn (MPE) of -0.64. Using employment responses to pension-wealth increases among women in Germany, Artmann et al. (2023) estimate an MPE of -0.3. Exploiting lottery winnings, Imbens et al. (2001) estimate an MPE of -0.11 in Massachusetts; Cesarini et al. (2017) find MPEs in Sweden, in the range -0.17 to -0.15; and Golosov et al. (2024) estimate an MPE of -0.5 in the United States. Using changes in survivors' benefits, Coyne et al. (2024), Rabaté and Tréguier (2024), and Giupponi (2019) estimate MPEs of -0.3, -0.2, and -0.8, respectively. Hence, the implied wealth effect would fall within the range estimated in the literature.

In addition to the wealth effect, the increase in employment could also be driven by precautionary motives. Precautionary labor supply has proven more difficult to estimate than wealth effects, and, to the best of my knowledge, no estimates specifically related to pension uncertainty exist. Pistaferri (2003) estimates precautionary labor supply and finds it to be small but significant, and Jessen et al. (2018) quantify precautionary labor supply motives as accounting for 2.8% of hours worked among married men. These estimates are less comparable to the model output than those of wealth effects, but perhaps a 2.7 percentage-point increase in the probability of exiting work can be considered of the same order as 2.8% of hours.

It is important, however, to test model mechanisms as well as estimate them, as emphasized by Fang et al. (2007), among others. Section 4.2 discussed key patterns in the belief data, including the declining pattern of misbeliefs used to identify λ, and, crucially here, two patterns related to the relationship between misbeliefs earlier in life and employment responses later. These two patterns provide natural, untargeted moments for testing the model's mechanisms.

The first pattern was that individuals who are worse informed about their SPA in their late 50s exhibit smaller labor supply responses at SPA in their 60s as captured by the negative interaction term in Column (5) of Table I. Two opposing forces in the model link the accuracy of earlier SPA knowledge to labor supply responses to it. Endogenous SPA knowledge implies that those less dependent on the SPA acquire less information about it. Conversely, households that are less well informed due to unfavorable signal realizations rather than their choice of signal distribution face a larger shock upon learning their SPA, prompting a greater reaction. Which of these two mechanisms dominates, and hence whether the model generates a positive or negative relationship between misbeliefs and employment response to the SPA, is not predetermined. The middle panel of Table II shows that the model generates a negative relationship, indicating the model reproduces the observed direction of this relationship, with the empirical counterpart repeated in the final column.

The second pattern was that those who over-estimate their SPA, as opposed to those who under-estimate, have a larger employment response upon reaching SPA as captured by the

negative interaction term in Column (6) of Table I. This is because, all else equal, they receive a larger shock upon reaching their SPA. The bottom panel also shows that the model replicates the direction of the dependence between SPA employment responses and SPA misbeliefs.

8.2. *Comparison to Reference point retirement*

A prominent explanation for the sharp employment response at pension eligibility is reference-dependent preferences or norms that shift in utility from leisure at a pension age. This mechanism provides a compelling account of behavioral discontinuities and has been influential in explaining responses to statutory retirement ages (e.g. Seibold, 2021). However, reference dependence typically abstracts from widespread misbeliefs about the eligibility age. Although reference dependence does not require complete information, the widespread misbeliefs documented here and in Caplin et al. (2022a) about the age it posits as a salient reference point are somewhat in tension with it as an explanation. More importantly, there is an advantage to having a more parsimonious explanation that accounts for both misbeliefs and excess sensitivity simultaneously. Moreover, the predictive value of misbeliefs for employment responses suggests they are related phenomena, something that follows naturally from the model of costly attention but fits uneasily with explanations that abstract from misbeliefs, such as reference-dependence.

Previous studies from this literature generally introduce a reference point or utility kink to match the response at pension age directly. Instead, I estimate an attention cost from belief data without directly targeting the puzzle. When I also choose a parameter value to directly target the puzzle, in Column (4) of Table II, costly attention accounts for 47% of the gap for the whole population and completely accounts for the richer subgroup, with both estimates falling within the 95% confidence intervals. So, on this metric, costly attention can be seen as doing a comparable job to reference dependence in explaining excess employment sensitivity.

Directly targeting the employment response at SPA sacrifices some of the most appealing features of costly attention. Namely, it provides a unified framework to explain both the observed employment response and belief inaccuracy, and it allows us to bring in additional data on beliefs to identify the underlying mechanism. When the model is thus constrained to match observed misbeliefs, however, it only partially explains the excess employment sensitivity.

Therefore, it is worth considering why the model underestimates employment responses to the SPA rather than just pivoting to targeting the puzzle itself. I consider two potential explanations. Firstly, the model attributes all policy learning to the SPA, while real-world pension systems involve multiple complex features—some of which may also become salient around eligibility. This could understate the degree of learning at SPA and, hence, the size of the shocks received. Appendix F explores this possibility by extending the model to include uncertainty over deferral rules, though a lack of data on deferral rate beliefs makes this exercise more speculative. Secondly, belief-driven mechanisms may operate in parallel with behavioral ones, such as reference dependence. There is suggestive evidence that framing effects, such as labeling eligibility as a retirement norm, can influence labor supply responses—consistent with reference-point models. I discuss this further in Appendix E.3. Appendix E.4 also presents results from an estimated model version that includes, in addition to costly attention, a share of passive decision-makers (as in Chetty et al., 2014), who retire mechanically at their SPA. This extension provides a simple way to capture potential reference dependence, allowing for the possibility that both mechanisms are working in parallel. This exercise indicates that reference dependence can explain the portion of the employment response not explained by misbeliefs.

Thus, the costly attention model is perhaps best understood not as a rejection of reference dependence, but as a complementary explanation that addresses dimensions (such as belief heterogeneity) not captured by reference-based preferences. In doing so, it might help explain why

responses vary across individuals, cohorts, and institutional contexts. For example, Deshpande et al. (2024) find smaller employment responses to the U.S. full retirement age during reform periods. If behavior were driven entirely by fixed preferences, such variation would be hard to reconcile. With costly attention, however, policy salience and information campaigns during reform periods can reduce misbeliefs, attenuating responses. Further research could shed light on the relative importance of these two explanations.

8.3. *Model Implications and Predictions*

Attention cost size. λ is hard to interpret, having natural units of utils per bit. While utils are known to be non-interpretable, denoting in bits is also problematic, as it exaggerates costs, since models contain far fewer learnable bits than reality. Most models contain only single or double-digit bits of information, fewer than in an average sentence. Reality holds vastly more information, making per-bit information cost a larger share of total model information. To address both issues, I calculate the compensating asset that raises household utility as much as perfect SPA knowledge, effectively their willingness to pay to learn their SPA. For $\hat{\lambda} = 6 \times 10^{-8}$, compensating assets range from £6 at the 25th percentile to £14 at the 75th, with a mean of £11. For $\lambda = 5 \times 10^{-7}$, the mean is £31 (summary of compensating assets distributions for both λ values in the appendix Table XX).

The model finds such modest friction because the benefit from learning about your SPA is not of the order of the State pension benefit but of the difference this information makes in saving and labor supply decisions, and overall, these decisions are not that sensitive to this information. These results are similar in magnitude to findings from other settings, both in terms of the implied sensitivity of decisions to information about pensions and the estimated welfare effects. Mastrobuoni (2011), Dolls et al. (2018), Liebman and Luttmer (2015) find modest impacts of information about pension on choices. Luttmer and Samwick (2018) estimates that people would be prepared to part with 6% of their pension benefit to remove all pension uncertainty. Although not directly comparable to the estimates here since they estimate willingness to pay to remove *all* policy uncertainty, and I estimate willingness to pay to be perfectly informed of *current policy*, at the mean 6% of the pension benefit in my setting is £180, which is not of too dissimilar a magnitude to the estimate presented above.

The employment response to pension age reforms. Rising old-age dependency ratios make increasing older individuals' employment a global policy priority, with pension ages seen as a key tool (e.g. Kolsrud et al., 2024). This paper shows that misbeliefs from costly attention amplify employment responses at the SPA. This may naturally raise the question of whether misinformation makes the SPA a more effective tool to increase old-age employment. In fact, the model generally implies it does the opposite.

Column (2) of Table III shows the additional mean employment between ages 55-65 when the SPA is reformed from 60 to the value between 61 and 65 indicated for that row, in the model with $\hat{\lambda} = 6 \times 10^{-8}$ where prior beliefs and other state variables are initialized to the values of the SPA 60 cohort. Thus, this captures the response to an unanticipated SPA increase at age 55 from 60 to a later age, since households have savings, accrued pension entitlements, and the beliefs of a cohort that had a SPA of 60. Column (1) shows results from the model with policy uncertainty but no attention costs. Both versions show modest employment gains, but the increases are generally larger under costly attention. For example, for post-reform SPA 65, mean employment rises 0.27 years with attention costs vs. 0.31 without. So, employment rises by up to 15% more under costly attention, which may seem at odds with the finding that it causes a larger employment drop at SPA.

TABLE III

IMPACTS OF REFORMING SPA WITH INFORMED AND UNINFORMED HOUSEHOLDS

SPA increased from 60 to:	(1) - Informed Added Employment	(2) - Uninformed Added Employment	(3) MC	(4) WTP	(5) MR
61	0.07	0.06	£3.50	£4.22	£28.45
62	0.14	0.14	£4.00	£2.37	£11.78
63	0.18	0.16	£4.50	£18.34	£19.91
64	0.22	0.20	£5.00	£31.64	£4.31
65	0.31	0.27	£5.50	£44.41	£68.52

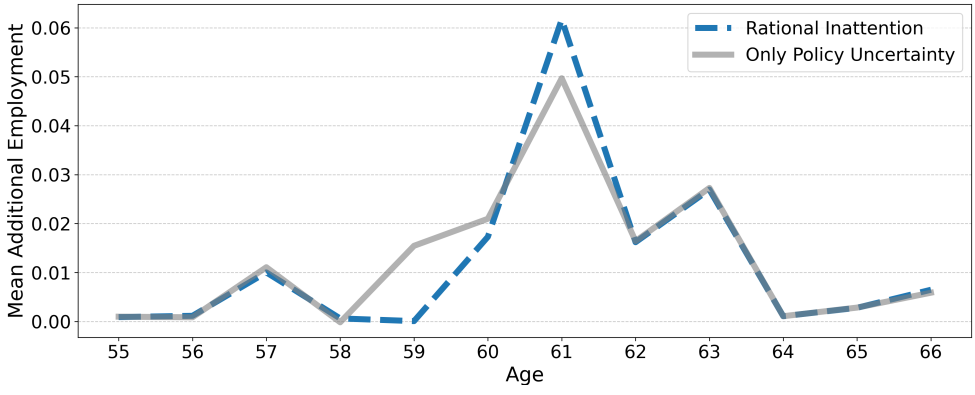
Note: Employment increases over 56-65 from raising SPA from 60 to the age in Column (1) with costly attention and in Column (2) without it. Columns 3-5 show the financial impacts of an accompanying information letter campaign that moves people from uninformed to informed. Column (3) shows the marginal cost, Column (4) the model implied mean willingness to pay, and Column (5) the model implied the marginal revenue.

This tension resolves when noting that rationally inattentive households respond less immediately to SPA increases. Fully informed households internalize the change early, increasing work in their 50s. Inattentive households react later when they realize their SPA has increased and they must compensate for lost earnings. This compensatory effort reduces but generally does not eliminate the difference over 55-59 due to imperfect intertemporal substitution and lower employment at older ages. This effect also inflates employment just before SPA, amplifying the drop at SPA. Thus, costly attention yields smaller overall employment gains but a larger response at SPA, with this bunching driven by intertemporal shifts. This may help us understand how small information frictions can generate relatively large employment effects, because much of the concentrated employment response reflects an intertemporal shifting of employment, and the welfare implications are not as high as they seem when we focus on the local effect. The echoes other results in the literature, such as Chetty (2012) and Choukhmane (2025), which find that apparently large deviations from optimizing behavior are explained by modest friction once we account for dynamics. Figure 4 illustrates the intertemporal shifting of employment for an SPA rise to 62, for which the two effect roughly cancel out.

The impact of information on response to pension age reforms. Columns (1) and (2) of Table III show added employment from an unanticipated SPA increase at age 55 in models with and without costly information. The only difference is in Column (1), households know the SPA, and in Column (2), they do not. Thus, the gap reflects the maximum potential impact of an annual information letter campaign. Columns (3)-(5) assess such a campaign.

Column (3) reports the marginal cost of the information letter campaign. After covering fixed costs, the only marginal cost is postage at £0.50/year (2013 prices, like the model). Column (4) shows the willingness to pay (WTP) for the information campaign under each post-reform SPA. Two forces drive WTP: higher SPAs reduce lifetime wealth (lowering WTP), but also, as it moves further from the pre-reform SPA of 60, the value of information rises. Initially, the first effect dominates, reducing WTP. From SPA 63 onward, the second dominates, and WTP increases. Comparing Columns (3) and (4) shows WTP for information exceeds the campaign's marginal cost for all post-reform SPAs except 62. For these reforms, the information campaign improves net welfare without accounting for added government revenue, but since the campaign also raises employment (see the difference between Columns (1) and (2)), the campaign is revenue-positive as quantified in Column (5), which presents marginal revenue from the campaign. Though modest (because 1950s-born women had low earnings), marginal revenue exceeds marginal cost for all SPA reforms except 64. Combining household and government gains, Columns (3)-(5) show the information campaign consistently raises total welfare, with

FIGURE 4.—Additional Employment resulting from Increasing the SPA from 60 to 62



Note: For the two versions, employment increases resulting from a reform of the female SPA from 60 to 62.

benefits exceeding costs by 3.5 to 20.5 times. Though absolute gains are modest, the experiment underscores a key point: informing individuals not only improves their welfare but also improves their policy responsiveness. Additionally, as pointed out by Dolls et al. (2018), the gains from information letter campaigns are not so modest when benchmarked against other policies that aim to increase old-age employment or retirement saving.

9. CONCLUSION

Mistaken beliefs are common, but their economic impacts are still not well-understood. Using UK data, this paper shows that incorporating costly attention, which endogenously generates misbeliefs, into a retirement model explains both observed misbeliefs and the sensitivity of employment to pension eligibility ages. Costly attention accounts for 51% of the employment response gap between the model and the data when calibrated to observed beliefs, and all of the gap when not so constrained. Given that both pension misbeliefs and excessive employment responses are across-country regularities, these insights may be cross-nationally relevant.

Endogenous information acquisition is key to explaining retirement behavior, but leads to the prior belief becoming a state variable. This high-dimensional state variable significantly increases computational demands. I propose a method for solving dynamic rational inattention models without suppressing beliefs as a state variable. From the belief data, I estimate the mean willingness to pay to learn the SPA as £11. Though small, this far exceeds the marginal cost of information letters. Policy experiments show that, after most SPA reforms, households' willingness to pay for such letters exceeds their cost, but also that sending letters increases employment by up to 15%. Hence, the campaign raises additional tax revenue, which, for most SPA reforms, also exceeds the cost. Considering total benefits to government and households, the campaign always improves welfare, with benefits outweighing costs by 3.5 to 20.5 times.

REFERENCES

ABADIE, ALBERTO, SUSAN ATHEY, GUIDO W IMBENS, AND JEFFREY M WOOLDRIDGE (2023): "When should you adjust standard errors for clustering?" *The Quarterly Journal of Economics*, 138 (1), 1–35. [9, 41, 42, 43, 44, 45, 46, 48, 49, 50]

AFROUZI, HASSAN AND CHOONGRYUL YANG (2021): "Dynamic rational inattention and the Phillips curve," . [4]

AIGNER, DENNIS J (1973): "Regression with a binary independent variable subject to errors of observation," *Journal of Econometrics*, 1 (1), 49–59. [14, 48]

ALMÅS, INGVILD, ORAZIO ATTANASIO, AND PAMELA JERVIS (2024): "Presidential address: Economics and measurement: New measures to model decision making," *Econometrica*, 92 (4), 947–978. [3]

AMIN-SMITH, NEIL AND ROWENA CRAWFORD (2018): "State pension age increases and the circumstances of older women," in *The Dynamics of Ageing: Evidence from the english longitudinal study of ageing*, 9–39. [5]

ARMENTER, ROC, MICHÈLE MÜLLER-ITTEN, AND ZACHARY R STANGEBYE (2024): "Geometric methods for finite rational inattention," *Quantitative Economics*, 15 (1), 115–144. [3, 4, 25, 27]

ARTMANN, ELISABETH, NICOLA FUCHS-SCHÜNDELN, AND GIULIA GIUPPONI (2023): "Forward-looking labor supply responses to changes in pension wealth: Evidence from germany," . [32]

ATALAY, KADIR AND GARRY F BARRETT (2015): "The impact of age pension eligibility age on retirement and program dependence: Evidence from an Australian experiment," *Review of Economics and Statistics*, 97 (1), 71–87. [8]

BAIROLIYA, NEHA AND KATHLEEN MCKIERNAN (2023): "Revisiting Retirement and Social Security Claiming Decisions," . [2]

BANKS, J., G.D. BATTY, J.J.F. BREEDVELT, K. COUGHLIN, R. CRAWFORD, M. MARMOT, J. NAZROO, Z. OLDFIELD, N. STEEL, A. STEPTOE, MARTIN. WOOD, AND P. ZANINOTTO (2021): "English Longitudinal Study of Ageing: Waves 0-9," 1998-2019 [data collection]. [7]

BARTOŠ, VOJTĚCH, MICHAL BAUER, JULIE CHYTILOVÁ, AND FILIP MATĚJKA (2016): "Attention Discrimination: Theory and Field Experiments With Monitoring Information Acquisition," *American Economic Review*, 106 (6), 1437–1475. [4]

BEHAGHEL, LUC AND DAVID BLAU (2012): "Framing Social Security Reform : Behavioral Responses to Changes in the Full Retirement Age," *American Economic Journal: Economic Policy*, 4 (4), 41–67. [4]

BERNHEIM, DOUGLAS (1988): *Social Security Benefits: An Empirical Study of Expectations and Realizations*. [5]

BHALOTRA, SONIA, ADELINE DELAVANDE, PAULINO FONT-GILABERT, AND JOANNA MASELKO (2025): "Maternal investments in children: The role of expected effort and returns," *The Economic Journal*, 135 (667), 712–747. [26]

BOEHM, EDUARD (2023): "Intermediation, Choice Frictions, and Adverse Selection: Evidence from the Chilean Pension Market," . [4]

BORUSYAK, KIRILL, XAVIER JARAVEL, AND JANN SPIESS (2024): "Revisiting event-study designs: robust and efficient estimation," *Review of Economic Studies*, rdae007. [9, 10, 42]

BOZIO, ANTOINE, ROWENA CRAWFORD, AND GEMMA TETLOW (2010): "The history of state pensions in the UK: 1948 to 2010," Tech. rep. [5, 17, 62]

BRONCHETTI, ERIN T, JUDD B KESSLER, ELLEN B MAGENHEIM, DMITRY TAUBINSKY, AND ERIC ZWICK (2023): "Is attention produced optimally? Theory and evidence from experiments with bandwidth enhancements," *Econometrica*, 91 (2), 669–707. [22]

BROWN, ZACH Y AND JIHYE JEON (2024): "Endogenous information and simplifying insurance choice," *Econometrica*, 92 (3), 881–911. [4]

BURTLESS, GARY (1986): "Social Security, Unanticipated Benefit Increases, and the Timing of Retirement," *Review of Economic Studies*, 53 (5), 781–805. [4]

CAPLIN, ANDREW, MARK DEAN, AND JOHN LEAHY (2019): "Rational Inattention, Optimal Consideration Sets, and Stochastic Choice," *Review of Economic Studies*, 86 (3), 1061–1094. [2, 3, 4, 22, 25, 26, 27, 52]

——— (2022a): "Rationally inattentive behavior: Characterizing and generalizing Shannon entropy," *Journal of Political Economy*, 130 (6), 1676–1715. [22, 25, 33]

CAPLIN, ANDREW, EUNGIK LEE, SØREN LETH-PETERSEN, AND JOHAN SÆVERUD (2022b): "Communicating social security reform," Tech. rep., National Bureau of Economic Research. [5, 15]

CARROLL, CHRISTOPHER AND ANDREW SAMWICK (1996): "The Nature of Precautionary Wealth," *NBER Working Paper*. [10]

CASANOVA, MARIA (2010): "Happy Together : A Structural Model of Couples ' Joint Retirement Choices," *Working Paper*, (5), 1–53. [17]

CESARINI, DAVID, ERIK LINDQVIST, MATTHEW J NOTOWIDIGDO, AND ROBERT ÖSTLING (2017): "The effect of wealth on individual and household labor supply: evidence from Swedish lotteries," *American Economic Review*, 107 (12), 3917–46. [32, 43]

CHETTY, RAJ (2012): "Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply," *Econometrica*, 80 (3), 969–1018. [4, 35]

CHETTY, RAJ, JOHN N FRIEDMAN, SØREN LETH-PETERSEN, TORBEN HEIEN NIELSEN, AND TORE OLSEN (2014): "Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from Denmark," *The Quarterly Journal of Economics*, 129 (3), 1141–1219. [33, 60]

CHOUKHMANE, TAHA (2025): "Default Options and Retirement Saving Dynamics," *American Economic Review*, 115 (11), 3749–3787. [4, 35]

Ciani, Emanuele, Adeline Delavande, Ben Etheridge, and Marco Francesconi (2023): "Policy uncertainty and information flows: Evidence from pension reform expectations," *The Economic Journal*, 133 (649), 98–129. [1]

Coyne, David, Itzik Fadlon, Shanthi P Ramnath, and Patricia K Tong (2024): "Household labor supply and the value of social security survivors benefits," *American Economic Review*, 114 (5), 1248–1280. [32]

Cribb, Jonathan, Carl Emmerson, and Gemma Tetlow (2016): "Signals matter? Large retirement responses to limited financial incentives," *Labour Economics*, 42, 203–212. [5, 6, 8, 9]

de Bresser, Jochem (2023): "Evaluating the Accuracy of Counterfactuals: Heterogeneous Survival Expectations in a Life Cycle Model," *Review of Economic Studies*, rdad088. [5]

De Bruin, Wändi Bruine, Alycia Chin, Jeff Dominitz, and Wilbert Van der Klaauw (2023): "Household surveys and probabilistic questions," in *Handbook of Economic Expectations*, Elsevier, 3–31. [12]

De Chaisemartin, Clément and Xavier d'Haultfoeuille (2020): "Two-way fixed effects estimators with heterogeneous treatment effects," *American economic review*, 110 (9), 2964–2996. [9]

De Nardi, Mariacristina (2004): "Wealth inequality and intergenerational links," *The Review of Economic Studies*, 71 (3), 743–768. [16]

De Nardi, Mariacristina, Eric French, and John Bailey Jones (2009): "Life expectancy and old age savings," *American Economic Review*, 99 (2), 110–15. [56]

De Nardi, Mariacristina, Eric French, and John Bailey Jones (2010): "Why Do the Elderly Save? The Role of Medical Expenses," *Journal of Political Economy*, 118 (1), 39–75. [29, 43]

Dean, Mark and Nathaniel Neligh (2023): "Experimental tests of rational inattention," *Journal of Political Economy*, 131 (12), 3415–3461. [22]

Delavande, Adeline (2008): "Pill, patch, or shot? Subjective expectations and birth control choice," *International Economic Review*, 49 (3), 999–1042. [5]

Deshpande, Manasi, Itzik Fadlon, and Colin Gray (2024): "How Sticky Is Retirement Behavior in the United States?" *Review of Economics and Statistics*, 106 (2), 370–383. [34]

Dolls, Mathias, Philipp Doerrenberg, Andreas Peichl, and Holger Stichnoth (2018): "Do retirement savings increase in response to information about retirement and expected pensions?" *Journal of Public Economics*, 158, 168–179. [34, 36]

Etgeton, Stefan, Björn Fischer, and Han Ye (2023): "The effect of increasing retirement age on households' savings and consumption expenditure," *Journal of Public Economics*, 221, 104845. [8]

Fang, Hanming, Michael Keane, Ahmed Khwaja, Martin Salm, and Dan Silverman (2007): "Testing the mechanisms of structural models: The case of the mickey mantle effect," *American Economic Review*, 97 (2), 53–59. [32]

French, Eric (2005): "The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behavior," *Review of Economic Studies*, 72 (2), 395–427. [2, 4, 29, 56]

French, Eric and John Bailey Jones (2011): "The Effects of Health Insurance and Self-Insurance on Retirement Behavior," *Econometrica*, 79 (3), 693–732. [4]

Fuster, Andreas, Ricardo Perez-Truglia, Mirko Wiederholt, and Basit Zafar (2022): "Expectations with endogenous information acquisition: An experimental investigation," *Review of Economics and Statistics*, 104 (5), 1059–1078. [22]

Gabaix, Xavier (2019): "Behavioral inattention," in *Handbook of Behavioral Economics: Applications and Foundations 1*, Elsevier, vol. 2, 261–343. [1]

García-Miralles, Esteban and Jonathan M. Leganza (2024): "Joint retirement of couples: Evidence from discontinuities in Denmark," *Journal of Public Economics*, 230, 105036. [10, 44]

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus (2021): "Five facts about beliefs and portfolios," *American Economic Review*, 111 (5), 1481–1522. [26]

Giupponi, Giulia (2019): "When income effects are large: Labor supply responses and the value of welfare transfers," . [32]

Golosov, Mikhail, Michael Graber, Magne Mogstad, and David Novgorodsky (2024): "How Americans respond to idiosyncratic and exogenous changes in household wealth and unearned income," *The Quarterly Journal of Economics*, 139 (2), 1321–1395. [32]

Gruber, Jonathan, Ohto Kanninen, and Terhi Ravaska (2022): "Relabeling, retirement and regret," *Journal of Public Economics*, 211, 104677. [5, 60]

Gruber, Jonathan and David A Wise (2004): "Social Security Programs and Retirement around the World: Micro-Estimation," Tech. rep., National Bureau of Economic Research. [1, 5]

Gustman, Alan L and Thomas L Steinmeier (1986): "A Structural Retirement Model," *Econometrica*, 54 (3), 555–584. [4]

——— (2005): "Imperfect knowledge of social security and pensions," *Industrial Relations: A Journal of Economy and Society*, 44 (2), 373–397. [1]

IMBENS, GUIDO W, DONALD B RUBIN, AND BRUCE I SACERDOTE (2001): "Estimating the effect of unearned income on labor earnings, savings, and consumption: Evidence from a survey of lottery players," *American economic review*, 91 (4), 778–794. [32]

JESSEN, ROBIN, DAVUD ROSTAM-AFSCHAR, AND SEBASTIAN SCHMITZ (2018): "How important is precautionary labour supply?" *Oxford Economic Papers*, 70 (3), 868–891. [32]

JOHNSEN, JULIAN VEDELER, KJELL VAAGE, AND ALEXANDER WILLÉN (2022): "Interactions in public policies: Spousal responses and program spillovers of welfare reforms," *The Economic Journal*, 132 (642), 834–864. [10]

KŐSZEGI, BOTOND AND FILIP MATĚJKA (2020): "Choice Simplification: A Theory of Mental Budgeting and Naive Diversification," *The Quarterly Journal of Economics*, 135 (2), 1153–1207. [22]

KALYVAS, CHRISTOS AND THEODOROS TZOURAMANIS (2017): "A survey of skyline query processing," *arXiv preprint arXiv:1704.01788*. [54]

KAPOR, ADAM J, CHRISTOPHER A NEILSON, AND SETH D ZIMMERMAN (2020): "Heterogeneous beliefs and school choice mechanisms," *American Economic Review*, 110 (5), 1274–1315. [26]

KOŞAR, GIZEM AND CORMAC O'DEA (2022): " Expectations Data in Structural Microeconomic Models," *National Bureau of Economic Research*. [5]

KOLSRUD, JONAS, CAMILLE LANDAIS, DANIEL RECK, AND JOHANNES SPINNEWIJN (2024): "Retirement consumption and pension design," *American Economic Review*, 114 (1), 89–133. [34]

LALIVE, RAFAEL, ARVIND MAGESAN, AND STEFAN STAUBLI (2023): "How social security reform affects retirement and pension claiming," *American Economic Journal: Economic Policy*, 15 (3), 115–150. [5, 14, 60]

LALIVE, RAFAEL AND PIERPAOLO PARROTTA (2017): "How does pension eligibility affect labor supply in couples?" *Labour Economics*, 46, 177–188. [10]

LIEBMAN, JEFFREY B AND ERZO FP LUTTMER (2015): "Would people behave differently if they better understood social security? Evidence from a field experiment," *American Economic Journal: Economic Policy*, 7 (1), 275–299. [34]

LOW, HAMISH, COSTAS MEGHIR, AND LUIGI PISTAFERRI (2010): "Wage risk and employment risk over the life cycle," *American Economic Review*, 100 (4), 1432–1467. [29]

LUMSDAINE, ROBIN L, JAMES H STOCK, AND DAVID A WISE (1996): "Why Are Retirement Rates So High at Age 65?" in *Advances in the Economics of Aging*, 61–82. [4]

LUSARDI, ANNAMARIA AND OLIVIA S MITCHELL (2011): "Financial literacy and retirement planning in the United States," *Journal of pension economics & finance*, 10 (4), 509–525. [1]

LUTTMER, ERZO F.P. AND ANDREW SAMWICK (2018): "The Welfare Cost of Perceived Policy Uncertainty: Evidence from Social Security," *American Economic Review*, 108 (2), 275–307. [34]

MACAULAY, ALISTAIR (2021): "Cyclical Attention to Saving *," . [3, 4, 25]

MACKOWIAK, BARTOSZ, FILIP MATEJKA, AND MIRKO WIEDERHOLT (2018): "Rational inattention: A disciplined behavioral model," *CEPR Discussion Papers*, 13243. [22]

MAĆKOWIAK, BARTOSZ AND MIRKO WIEDERHOLT (2009): "Optimal sticky prices under rational inattention," *American Economic Review*, 99 (3), 769–803. [4]

——— (2015): "Business cycle dynamics under rational inattention," *Review of Economic Studies*, 82 (4), 1502–1532. [4]

MANSKI, CHARLES (2004): "Measuring expectations," *Econometrica*, 72 (5), 1329–1376. [5]

MASTROBUONI, GIOVANNI (2009): "Labor supply effects of the recent social security benefit cuts: Empirical estimates using cohort discontinuities," *Journal of public Economics*, 93 (11-12), 1224–1233. [4]

——— (2011): "The role of information for retirement behavior: Evidence based on the stepwise introduction of the Social Security Statement," *Journal of Public Economics*, 95 (7-8), 913–925. [34]

MATĚJKA, FILIP AND ALISDAIR MCKAY (2015): "Rational inattention to discrete choices: A new foundation for the multinomial logit model," *American Economic Review*, 105 (1), 272–298. [2, 4, 24, 25]

MIAO, JIANJUN, JIERAN WU, AND ERIC R YOUNG (2022): "Multivariate rational inattention," *Econometrica*, 90 (2), 907–945. [4]

MIAO, JIANJUN AND HAO XING (2024): "Dynamic discrete choice under rational inattention," *Economic Theory*, 77 (3), 597–652. [3, 4, 22, 25]

NEWEY, WHITNEY K (1985): "Generalized method of moments specification testing," *Journal of econometrics*, 29 (3), 229–256. [30]

O'DEA, CORMAC (2018): "Insurance , Efficiency and Design of Public Pensions ," *WP*, (1), 1–63. [4, 29]

OECD (2011): "Net Pension Replacement Rates: Public and Private Schemes," in *Pensions at a Glance 2011: Retirement-income Systems in OECD and G20 Countries*, OECD Publishing. [6]

PISTAFERRI, LUIGI (2001): "Superior information, income shocks, and the permanent income hypothesis," *Review of Economics and Statistics*, 83 (3), 465–476. [5]

——— (2003): "Anticipated and unanticipated wage changes, wage risk, and intertemporal labor supply," *Journal of Labor Economics*, 21 (3), 729–754. [32]

PORCHER, CHARLY (2020): "Migration with Costly Information," . [3, 4, 25]
POWELL, MICHAEL JD (2009): "The BOBYQA algorithm for bound constrained optimization without derivatives," *Cambridge NA Report NA2009/06, University of Cambridge, Cambridge*, 26. [55]
RABATÉ, SIMON (2019): "Can I stay or should I go? Mandatory retirement and the labor-force participation of older workers," *Journal of Public Economics*, 180, 104078. [8, 11]
RABATÉ, SIMON, EGBERT JONGEN, AND TILBE ATAV (2024): "Increasing the retirement age: policy effects and underlying mechanisms," *American Economic Journal: Economic Policy*, 16 (1), 259–291. [10]
RABATÉ, SIMON AND JULIE TRÉGUIER (2024): "Labour supply and survivor insurance in the Netherlands," *Labour Economics*, 88, 102527. [32]
ROHWEDDER, SUSANN AND KRISTIN KLEINJANS (2006): "Dynamics of Individual Information about Social Security," *RAND WP*. [5]
RUST, JOHN AND CHRISTOPHER PHELAN (1997): "How Social Security and Medicare Affect Retirement Behavior In a World of Incomplete Markets," *Econometrica*, 65 (4), 781. [2, 4]
SCHENNACH, SUSANNE M (2016): "Recent advances in the measurement error literature," *Annual review of economics*, 8 (1), 341–377. [14]
SCHITTKOWSKI, K (2014): "NLPQLP-Nonlinear Programming with Non-Monotone and Distributed Line Search," . [54]
SEIBOLD, ARTHUR (2021): "Reference points for retirement behavior: Evidence from german pension discontinuities," *American Economic Review*, 11 (4), 1126–1165. [5, 14, 33, 59]
SIMS, CHRISTOPHER (2003): "Implications of rational inattention," *Journal of Monetary Economics*, 50 (3), 665–690. [2, 4]
STAUBLI, STEFAN AND JOSEF ZWEIMÜLLER (2013): "Does raising the early retirement age increase employment of older workers?" *Journal of public economics*, 108, 17–32. [8]
STEINER, JAKUB, COLIN STEWART, AND FILIP MATĚJKA (2017): "Rational Inattention Dynamics: Inertia and Delay in Decision-Making," *Econometrica*, 85 (2), 521–553. [3, 4, 22, 23, 25]
TUREN, JAVIER (2023): "State-dependent attention and pricing decisions," *American Economic Journal: Macroeconomics*, 15 (2), 161–189. [3, 4, 25]
VAN DER KLAAUW, WILBERT AND KENNETH I. WOLPIN (2008): "Social security and the retirement and savings behavior of low-income households," *Journal of Econometrics*, 145 (1-2), 21–42. [2, 4]
VESTAD, OLA LOTHERINGTON (2013): "Labour supply effects of early retirement provision," *Labour Economics*, 25, 98–109. [8]
VUL, EDWARD AND HAROLD PASHLER (2008): "Measuring the crowd within: Probabilistic representations within individuals," *Psychological Science*, 19 (7), 645–647. [28]

APPENDIX A: ADDITIONAL EMPIRICAL DETAILS

A.1. *Additional Institutional Details*

A.2. *Equity Acts*

The Equality Act (2006) banned mandatory retirement below age 65. Women observed reaching SPA in ELSA waves 1-7 did so after compulsory retirement at their SPAs (60-63) became illegal. The Equality Act (2010) banned all compulsory retirement ages with specific exceptions known as EJRA (Employer Justified Retirement Ages).

As EJRAs must be over 65 and all SPAs reached in the data are below this, EJRAs are not directly relevant. However, background and anecdotes may help illustrate the strictness of UK age discrimination laws on forced retirement. Seldon v Clarkson, Wright and Jakes (2012) clarified when EJRAs are justified, setting out three criteria. One, the justification must serve a public interest (e.g., intergenerational fairness), not just firm goals. Two, this objective must be consistent with the social policy aims of the state. Three, it must be a proportionate means to that end.

In Seldon v Clarkson, the plaintiff, a law firm partner, was subject to a justified EJRA. Documented EJRAs are rare; beyond law firm partners, the most debated cases involve Oxford and Cambridge. Most other UK universities have scrapped compulsory retirement. Notably, Oxford recently lost a tribunal where its EJRA was ruled unjustified. In Ewart v University of Oxford

TABLE IV

EFFECT OF SPA ON HAZARD RATE: HETEROGENEITY BY VERY LIQUID ASSETS (VLA)

	(1)	(2)	(3)	(4)
Over SPA	**0.128**	**0.120**	**0.128**	**0.140**
s.e	(0.0239)	(0.0320)	(0.0381)	(0.0237)
Over SPA× VLA above median			**-0.007**	
s.e			(0.0496)	
OVer SPA× VLA in £100K				**-0.012**
s.e				(0.0033)
Obs.	7,907	3,691	7,907	7,784

Note: Column (1) presents results from the specification in Equation 1 in the main text. Column (2) repeats the regression for those with above-median household Very Liquid Assets (VLA) in their last interview before SPA. Column (3) tests if treatment effects differ by fully interacting the specification with having above-median VLA. Column (4) adds an interaction between wealth and being above SPA. As well as the age, perdiod, and cohort dummies, all regressions control for marital status, years of education, highest qualification, self-reported health dummies, presence of a partner, partner's age, partner's age squared, partner's SPA eligibility, VLA, and a constant. Following Abadie et al. (2023), standard errors are clustered at the level of the treatment (i.e., birth cohort). Coefficients on controls and their interactions are not reported.

(2019), the court found Oxford's aim (intergenerational fairness) valid, but the EJRA disproportionate —its limited effectiveness didn't outweigh its clear harms. This underscores how seriously UK law treats forced retirement as age discrimination, with few, truly exceptional, exemptions.

A.3. *Robustness: Excess Employment Sensitivity*

A.3.1. *Restricted Asset Categorisation*

The aim of examining treatment effect heterogeneity by asset holdings is to assess the role of liquidity constraints. The main analysis uses NHNBW, but since parts of NHNBW may be illiquid, Table IV repeats the analysis using a narrower category—very liquid assets, i.e., those reasonably cashable within weeks. Results are qualitatively similar to those with NHNBW and do not suggest liquidity constraints alone explain the treatment effect. The effect remains positive for those above median assets; subgroup differences are still insignificant, and the continuous interaction shows heterogeneity is too weak for liquidity constraints to fully account for the effect.

A.3.2. *Bad Control Concerns*

Bad controls are a key concern in DID, with some arguing only time-invariant controls should be used, as controls imply parallel trends conditional on them. To address this, I take a broad approach and run the model without controls, showing that the main conclusions remain unchanged. Table V presents these results. As shown, they differ little from those with controls.

A.3.3. *Imputation Approach to DID*

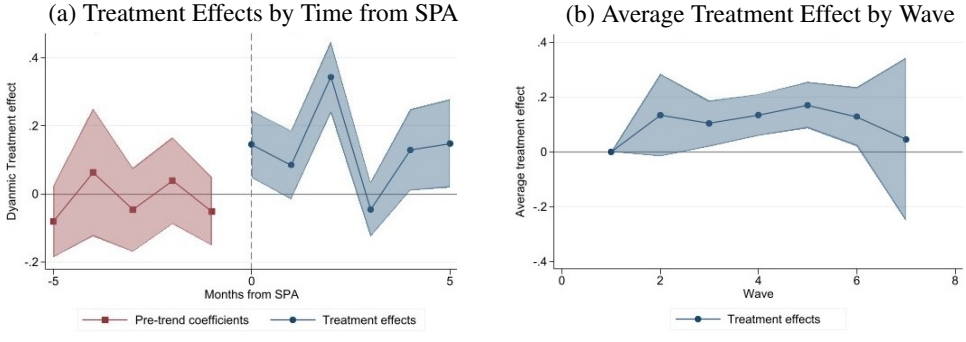
Two-way fixed effects DID models assume treatment effect heterogeneity across time and units. When treatment timing drives variation—as in this paper—violating these assumptions can yield nonsensical combinations of individual-level effects. Recent literature highlights this issue and, importantly, offers solutions that relax these assumptions.

TABLE V

EFFECT OF SPA ON HAZARD RATE: HETEROGENEITY BY NHNBW NO CONTROLS

	(1)	(2)	(3)	(4)
Over SPA	**0.123**	**0.093**	**0.161**	**0.136**
s.e	(0.02468)	(0.03155)	(0.03716)	(0.02599)
Over SPA× NHNBW above median			**-0.068**	
s.e			(0.04868)	
Over SPA× NHNBW in £100K				**-0.008**
s.e				(0.0023)
Obs.	8,119	3,963	8,119	7,898

Note: Column (1) presents results from the specification in Equation 1 in the main text. Column (2) repeats the regression for those with above-median Non-Housing Non-Business Wealth (NHNBW) in their last interview before SPA. Column (3) tests if treatment effects differ by fully interacting the specification with having above-median NHNBW. Column (4) adds an interaction between wealth and being above SPA. Following Abadie et al. (2023), standard errors are clustered at the level of the treatment (i.e., birth cohort). Coefficients on controls and their interactions are not reported.



Note: Panel (a) Average of the individual level treatment effects estimated using the imputation approach at a given time from SPA. Panel (b) shows the within-wave average of the individual-level treatment effects estimated using the imputation approach.

I apply the imputation method from Borusyak et al. (2024). Figure 5a shows dynamic treatment effects before and after SPA. No signs of violated parallel trends or anticipation effects appear as all pre-SPA effects are insignificant. A joint test confirms this ($p = .799$). Post-SPA, 4 of 6 effects are significant, and we reject the null of joint zero effects ($p = .000$). While the graph suggests limited variation among post-SPA effects, we cannot reject their equality ($p = .198$).

Figure 5b examines whether treatment effects vary by wave. They appear fairly uniform, though we can reject equality ($p = .137$). Neither violation of homogeneity seems severe, and overall, the graphs support the baseline assumption of a homogeneous treatment effect starting at SPA, though tests show this is an approximation.

These results show, allowing for arbitrary heterogeneity, something special is still happening at the SPA, which is difficult to explain in standard complete information models.

A.3.4. *Health, Wealth, Private Pensions, Joint Retirement, and Dismissals*

This section addresses alternative explanations for employment sensitivity to the SPA under a standard complete information framework. Specifically, it considers whether wealth, health, private pensions, joint retirement, or dismissals explain the labor supply response.

TABLE VI

HETEROGENEITY BY HEALTH

Over SPA	**0.113**
s.e	(0.0334)
Over SPA× V.good Health	**-0.003**
s.e	(0.0275)
Over SPA× Good Health	**0.035**
s.e	(0.0294)
Over SPA× Fair Health	**0.057**
s.e	(0.0458)
Over SPA× Poor Health	**0.025**
s.e	(0.0672)
Obs.	6,598

Note: Results of fully interacting specification in Equation 1 in the main text with a variable containing self-declared health status. As well as the age, perdiod, and cohort dummies, all regressions control for marital status, years of education, highest qualification, self-reported health dummies, presence of a partner, partner's age, partner's age squared, partner's SPA eligibility, NHNBW, and a constant. Following Abadie et al. (2023), standard errors are clustered at the level of the treatment (i.e., birth cohort). Coefficients on controls and their interactions are not reported.

Wealth effects influence labor supply, and women with later SPAs are lifetime poorer, so the puzzle isn't their higher labor supply but why it drops at the SPA, despite changes being announced 15 years or more in advance. In standard life-cycle models with complete information, a wealth-driven response should be spread over life, not concentrated at the SPA. In Equation 1 in the main text, lifetime wealth differences across birth cohorts (including those induced by SPA shifts) are absorbed by cohort effects. Thus, only within-cohort SPA-induced wealth differences are captured by the regressions. Additionally, the regression only captures the employment response at the SPA, so to explain the observed treatment effect via within-cohort wealth differences, the wealth effect would need to be enormous. Assuming a purely wealth-driven labor supply response implies a marginal propensity to earn (MPE) of about -0.3. This is on the high end of modern estimates (e.g. Cesarini et al., 2017), but becomes implausibly high given this captures just the final 2-3 months of a response that is spread out over 15-20 years. A wealth effect explanation also poorly explains the treatment's limited sensitivity to asset levels.

Health is a key driver of retirement decisions (e.g. De Nardi et al., 2010), but there's no reason to expect it to interact with the SPA or to explain employment's sensitivity to it. During the study period, the SPA was 60-63, while average health declines occur later. All the same, given health's importance, Table VI examines heterogeneity in labor supply response by health status. That table shows that those in the excluded category of excellent health have an 11% increase in their hazard of exiting employment at SPA and that this response does not differ significantly for those in any of the other self-reported health categories. A F-test of the joint significance of the interaction terms cannot reject the null of no significance ($p = 0.809$)

Private pension eligibility affects retirement decisions. However, occupational pension schemes likely didn't adjust pension ages with the female SPA, as private pensions are rarely differentiated by gender[33], and this reform only affected women. Still, checking for correlation between SPA and private pension normal pension ages (NPAs) is desirable. Checking this directly in ELSA is complicated by the fact that only self-reported NPAs are available. For the SPA, where alongside self-reports, we know an individual's true SPA, these self-reported

[33]This is likely illegal due to anti-discrimination law. The ECJ Test-Achats (2012) ruling barred gender-based pricing in insurance, and Barber v Guardian Royal Exchange (1990) found that employer pension constitute pay and so are subject to wage discrimination legislation.

TABLE VII

EFFECT OF SPA ON HAZARD RATE:
BELOW £2,000 IN DEFINED BENEFIT PENSION WEALTH

Over SPA	**0.180**
s.e	(0.0458)
Over SPA× NHNBW above median	**-0.088**
s.e	(0.0612)
Obs.	5,668

Note: Results of fully interacting specification in Equation 1 in the main text with an indicator for having above median non-housing non-business wealth in the last interview before SPA run on the population with above £2,000 in defined benefit wealth. As well as the age, perdiod, and cohort dummies, all regressions control for marital status, years of education, highest qualification, self-reported health dummies, presence of a partner, partner's age, partner's age squared, partner's SPA eligibility, NHNBW, and a constant. Following Abadie et al. (2023), standard errors are clustered at the level of the treatment (i.e., birth cohort). Coefficients on controls and their interactions are not reported.

ages are unreliable, as is documented in main text Section 4.2. However, only defined benefit pension systems have NPAs, as defined contribution schemes can be accessed from age 55. Hence, dropping those with > £2,000 in DB pensions removes any unlikely SPA-NPA correlation from explaining the results. Table VII shows that, despite reduced power, the treatment effect remains significant.

Turning to joint retirement Table VIII, repeat the analysis from the main text but only for single women and those with non-working husbands. The patterns are qualitatively similar. Crucially for the argument of this paper, the treatment effect in the subgroup is not significantly different from the treatment effect in the whole population.

Conversely if we restrict to those with a partner, we can look at heterogeinty of the treatment effect by whether that partner is working or not. In Table IX we see that this heterogeinty is not significant as any reasonable level, indicating that joint retirement is not a first-order consideration in understanding the employment response to one's own SPA. At some level, this is unsurprising; the literature on estimating joint retirement use the exogenous variation arising from aging past one's own pension eligibility on one's own employment to estimate the spillover effects onto partner's employment, and it finds these spillover effect to be an order of magnitude smaller. For example, García-Miralles and Leganza (2024) reviews this literature and uses this strategy to estimate spillover effects finding: "for every 100 individuals who retire when they reach pension eligibility, around 8 of their spouses adjust their behavior to retire at the same time". Since there is no discrete change in partner's employment at own pension eligibility age, and spillover effects are on order of magnitude smaller than the direct effect of own pension age on employment, it is unsurprising that these spillover effects would not be first-order to explain the employment response to one's own pension eligibility.

As mentioned in the main text age, age-based mandatory retirement is illegal, and as discussed at the start of this section, this is interpreted strictly by the courts. It is still possible that firms illegally force people to retire. To address this possibility, Table X drops all women who self-report having been forced out of their last job. Given the small numbers who self-report having been dismissed, the results do not change significantly.

A.3.5. *Placebo Test*

Table XI contains the result of placebo test in which I drop observations over SPA and replace the treatment in Equation 1 with indicators for being one or two years below SPA. As can be seen neither of the treatment indicators that do not align with the SPA are significant, indicating that the results in the main text do capture something specific happening at SPA.

TABLE VIII

EFFECT OF SPA ON HAZARD RATE: SINGLES AND NON-WORKING HUSBANDS

	(1)	(2)	(3)	(4)
Over SPA	**0.096**	**0.073**	**0.099**	**0.113**
s.e	(0.03788)	(0.04855)	(0.05523)	(0.03832)
Over SPA×NHNBW above median			**-0.026**	
s.e			(0.07366)	
Over SPA× NHNBW in £100K				**-0.016**
s.e				(0.0041)
Obs.	3,007	1,722	3,007	2,952

Note: Repeats first four columns of Table I from the main text on the subsample of singles and women with non-working husbands. As well as the age, perdiod, and cohort dummies, all regressions control for marital status, years of education, highest qualification, self-reported health dummies, presence of a partner, partner's age, partner's age squared, partner's SPA eligibility, NHNBW, and a constant. Following Abadie et al. (2023), standard errors are clustered at the level of the treatment (i.e., birth cohort). Coefficients on controls and their interactions are not reported.

TABLE IX

TREATMENT EFFECT HETEROGENEITY BY PRESENCE OF A NON-WORKING HUSBAND

	(1)	(2)
Over SPA	**0.115**	**0.101**
s.e	(0.03492)	(0.04150)
Over SPA× partner not-working	**0.021**	**0.019**
s.e	(0.05299)	(0.06380)
Obs.	6,028	3,266

Note: Results of fully interacting Equation 1 with an indicator of partner being out of work as captured by not having income from employment. The top row shows the size of treatment effect and the second the coffficeint on the interactor of the the indicator of having a non-working partner and the treatment indicator. Column one show results for all women with a partner. Column two shows results for women with a partner with wealth above the median wealth cut-off used main text to be consider in the wealthy subgourp. As well as the age, perdiod, and cohort dummies, all regressions control for marital status, years of education, highest qualification, self-reported health dummies, presence of a partner, partner's age, partner's age squared, partner's SPA eligibility, NHNBW, and a constant. Following Abadie et al. (2023), standard errors are clustered at the level of the treatment (i.e., birth cohort). Coefficients on controls and their interactions are not reported.

A.4. *Descriptive Analysis of SPA Beliefs*

Mistaken beliefs could take on many forms. People might simply not update from the pre-reform SPA of 60, or cling to other salient numbers, like the male SPA of 65. To explore this, Figure 6a shows reported SPAs for two cohorts, with true SPAs of 61 and 65. While there is some clustering around salient ages, reports mainly center on the true SPA, resembling a noisy signal. This is consistent with a model of costly information acquisition.

Figure 6b shows self-reported SPA errors at age 58 using monthly bins, rather than the yearly ones in the main text. Little model-relevance is gained from this; 31The main new insight is the spike in errors of 12 months.

Figure 6c shows SPA misbeliefs at age 58 by education. These rise with an individual's age left school up to the 19-years-or-over category, suggesting that more educated people (up to

TABLE X

EFFECT OF SPA ON HAZARD RATE: EXCLUDING SELF-REPORTED FIRED

	(1)	(2)	(3)	(4)
Above SPA	**0.129**	**0.104**	**0.160**	**0.145**
s.e	(0.02423)	(0.03086)	(0.03750)	(0.02451)
Above SPA× NHNBW above median			**-0.057**	
s.e			(0.04849)	
Above SPA× NHNBW in £100K				**-0.012**
s.e				(0.0027)
Obs.	7,799	3,738	7,799	7,676

Note: Repeat first four columns of Table I from the main text excluding women that self-report being fired. As well as the age, perdiod, and cohort dummies, all regressions control for marital status, years of education, highest qualification, self-reported health dummies, presence of a partner, partner's age, partner's age squared, partner's SPA eligibility, NHNBW, and a constant. Following Abadie et al. (2023), standard errors are clustered at the level of the treatment (i.e., birth cohort). Coefficients on controls and their interactions are not reported.

TABLE XI

PLACEBO TESTS

	Above Age One Year Below SPA	Above Age Two Years Below SPA
Placebo Test Coefficient	**0.031**	**0.005**
s.e	(0.0256)	(0.0230)
Obs.	4,279	4,279

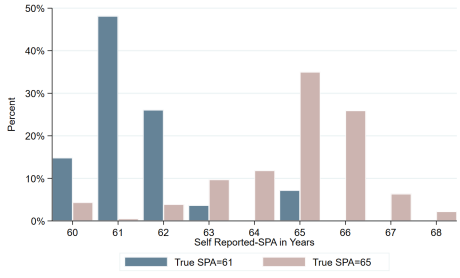
Note: Placebo test regressing the specification in Equation 1 after having dropped observations over SPA and replacing treatment indicator replaced with an indicator as per column heading. Column (4) adds an interaction between wealth and being above SPA. As well as the age, perdiod, and cohort dummies, all regressions control for marital status, years of education, highest qualification, self-reported health dummies, presence of a partner, partner's age, partner's age squared, partner's SPA eligibility, NHNBW, and a constant. Following Abadie et al. (2023), standard errors are clustered at the level of the treatment (i.e., birth cohort). Coefficients on controls and their interactions are not reported.

that point) are more mistaken. On the one hand, this is surprising, as we expect more educated people to have a higher information-processing capacity. On the other, the State Pension matters more for less educated individuals, giving them stronger incentives to learn. Thus, the pattern supports the modeling choices to focus on incentive heterogeneity rather than on ex-ante attention cost heterogeneity.

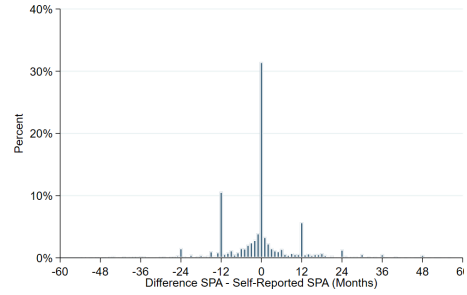
A.5. *Robustness of Treatment Effect Heterogeneity by Misbeliefs*

Section 4.2 investigated the relationship between SPA misbeliefs when respondents are in their 50s and their employment response to the SPA when they reached it in their 60s. In particular, it investigated the extent of support in the data for two model predictions. Firstly, that larger absolute error in self-reported SPAs should correlate with smaller employment responses to the SPA upon arriving at SPA, as individuals can select out of SPA knowledge, and those who do so will be those whose choices are not affected by the SPA. Secondly, that those whose misbeliefs over-predict their SPA should have a larger employment response upon reaching it, as they receive a positive wealth shock at that point. Columns (1)-(3) of Table XII provide robustness for the evidence presented in Section 4.2 of the first prediction, Columns (4)-(6) for the second prediction.

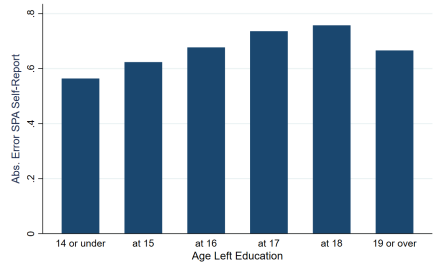
Column (1) of Table XII investigate heterogeneity in employment response to the SPA according to the size of discrepancy between self-reported SPA and the true SPA at, or around,



(a) SPA Beliefs by SPA-cohort



(b) Monthly SPA Misbeliefs Age 58



(c) Age 58 SPA Misbeliefs by Education

Note: Panel (a): self-Perceived SPA for two SPA-cohorts. One with a rounded SPA of 61 and one with a rounded SPA of 65. Panel (b): plot of error in self-reported SPA. The graph shows the frequency with which respondents gave mistaken answers about their SPA with errors at the true monthly level of SPA variation. Panel (c): SPA misbeliefs at age 58 by education.

age 58 without correcting for the fact individuals are asked about the age they will reach SPA and not its current value (i.e., without applying the drift correction). Without this correction, the point estimate of this heterogeneity is negative, consistent with the model predictions but only significant at the 10% level ($p = 0.091$).

Columns (2) and (3) use the square rather than the absolute value of this discrepancy to measure the size of misbeliefs. This is done because if self-reported SPAs are subject to classical measurement error, then squaring the difference between them and the true SPA will result in a variable subject to measurement error that is uncorrelated with the true squared discrepancy. Hence, unlike with the absolute value operator used to measure the size of misbeliefs in the main text, for the square, we recover the result that measurement error will only lead to attenuation bias. Columns (2) and (3) repeat the regression from the main text, but using the square of the error in SPA self-reports rather than the absolute value to investigate the heterogeneity in employment response to the SPA. Column (2) does this without drift-correction, and Column (3) does it with the drift-correction that is also applied to the results in the main text, as the square of the discrepancy was not used there. The results are negative, in line with the theory, and both are significant at the 10% level.

Finally, Columns (4) through (6) investigate the second prediction, that responses should be larger among people who over-predict their SPA. In the main text, this heterogeneity was investigated by applying the drift correction mentioned above to the data and by restricting the sample to address aggregation bias arising from a mixture of two populations with different responses. Here, the regression is run under other combinations of these assumptions. The point estimates are all negative in line with the theory, but the results are not significant, except Column (6), where the interaction is significant at the 10% level. As mentioned in the main text, for these threshold-type regressions used to investigate the impact of over- vs. under-

TABLE XII

ROBUSTNESS OF TREATMENT EFFECT HETEROGENEITY BY MISBELIEFS

	(1)	(2)	(3)	(4)	(5)	(6)
Over SPA	**0.184**	**0.178**	**0.178**	**0.175**	**0.175**	**0.178**
s.e	(0.0405)	(0.0345)	(0.0347)	(0.0360)	(0.0752)	(0.0458)
Over SPA × **|SPA - SPA Self-report|**	**-0.053**	—	—	—	—	—
s.e	(0.0310)					
Over SPA × **(SPA - SPA Self-report)**^{2}	—	**-0.010**	**-0.013**	—	—	—
s.e		(0.0058)	(0.0074)			
Over SPA × **SPA above Self-report**	—	—	—	**-0.127**	**-0.023**	**-0.185**
s.e				(0.1035)	(0.0926)	(0.1058)
Drift Corrected:	No	No	Yes	No	Yes	No
Restricted Sample:	No	No	No	No	No	Yes
Obs.	5,304	5,304	5,304	5,304	5,304	4,505

Note: The outcome variable is the hazard of exiting employment. Row 1 reports the treatment indicator for being over the State Pension Age (SPA). Rows 2–4 report interactions of this treatment indicator with: (2) the absolute difference between the true and self-reported SPA; (3) the square of the difference between the true and self-reported SPA; and (4) an indicator for the true SPA exceeding the self-reported SPA (i.e., under-estimation of the SPA). Self-reported SPAs are measured at age 58 or at the closest available age, and wealth is measured in the last interview before the SPA. Each column reports estimates from a specification of Equation 1 fully interacted with : (1) the absolute difference between the true and self-reported SPA, (2)-(3) the square of the difference between the true and self-reported SPA, and (4)-(6) an indicator of under-estimating the SPA. Whether the estimates reported in a column correct for the drift resulting from the difference between the SPA belief question and the object of interest or not is reported in the fifth row (i.e., it asks when you will arrive at SPA rather than your current age). The sixth row reports whether the sample has been restricted to exclude individuals making SPA self-report errors less than 5 quarters. As well as the age, period, and cohort dummies, all regressions control for marital status, years of education, highest qualification, self-reported health dummies, presence of a partner, partner's age, partner's age squared, partner's SPA eligibility, household non-housing non-business wealth, and a constant. Following Abadie et al. (2023), standard errors are clustered at the level of the treatment (i.e., birth cohort). Coefficients on controls and their interactions are not reported.

estimation, we know that measurement leads to attenuation bias, potentially explaining the lack of significance (Aigner, 1973). Additionally, we see in Column (6) that the results are sensitive to whether or not we apply the drift correction, since the only difference between this regression and the one producing a significant coefficient in the main text is the drift correction, and the significance level drops from 5% to 10%. This is perhaps unsurprising, since the drift correction shifts the point at which one is classified as over- or underestimating, and a threshold regression would likely be very sensitive to the correct placement of the threshold.

Table XIII repeats the analysis of Table XII but dropping controls from the regression to test if bad controls could be driving the results. The results are very similar, and the conclusions are essentially completely unchanged.

Table XIV attempts to test both predictions simultaneously in a single regression specification. In Columns (1) and (3), it does this by fully interacting all controls and dummies with both the absolute size of the error in self-reported SPA and the indicator of under-predicting your SPA. Column (1) does this without correcting for expected drift in SPA between interview and eligibility. Both coefficients are negative as predicted by the theory, but are not significant. Fully interacting all controls is extremely flexible, but it asks a lot of the data, and there is no ex-ante reason to suspect that the controls affect the different groups differently. So in Column (2), I do not allow the control to vary by misbeliefs but do control for the level of each interaction variable. In this specification, the coefficient on the size of misbeliefs becomes significant, but that on the direction remains negative and insignificant. Columns (3) and (4) repeat this analysis, but correcting for the expected drift in beliefs between the interview and SPA. The results are similar.

TABLE XIII

ROBUSTNESS OF TREATMENT EFFECT HETEROGENEITY BY MISBELIEFS: NO CONTROLS

	(1)	(2)	(3)	(4)	(5)	(6)
Over SPA	**0.184**	**0.178**	**0.178**	**0.175**	**0.175**	**0.178**
s.e	(0.0405)	(0.0345)	(0.0347)	(0.0360)	(0.0752)	(0.0458)
Over SPA$\times$ **|SPA - SPA Self-report|**	**-0.053**	—	—	—	—	—
s.e	(0.0310)					
Over SPA$\times$**(SPA - SPA Self-report)**^{2}	—	**-0.010**	**-0.012**	—	—	—
s.e		(0.0058)	(0.0074)			
Over SPA$\times$ **SPA above Self-report**	—	—	—	**-0.127**	**-0.022**	**-0.185**
s.e				(0.1035)	(0.0926)	(0.1058)
Drift Corrected:	No	No	Yes	No	Yes	No
Restricted Sample:	No	No	No	No	No	Yes
Obs.	5,304	5,304	5,304	5,304	5,304	4,505

Note: The outcome variable is the hazard of exiting employment. Row 1 reports the treatment indicator for being over the State Pension Age (SPA). Rows 2–4 report interactions of this treatment indicator with: (2) the absolute difference between the true and self-reported SPA; (3) the square of the difference between the true and self-reported SPA; and (4) an indicator for the true SPA exceeding the self-reported SPA (i.e., under-estimation of the SPA). Self-reported SPAs are measured at age 58 or at the closest available age, and wealth is measured in the last interview before the SPA. Each column reports estimates from a specification of Equation 1 fully interacted with : (1) the absolute difference between the true and self-reported SPA, (2)-(3) the square of the difference between the true and self-reported SPA, and (4)-(6) an indicator of under-estimating the SPA. Whether the estimates reported in a column correct for the drift resulting from the difference between the SPA belief question and the object of interest is reported in the fifth row (i.e., it asks when you will arrive at SPA rather than your current age). The sixth row reports whether the sample has been restricted to exclude individuals making SPA self-report errors of less than 5 quarters. Following Abadie et al. (2023), standard errors are clustered at the level of the treatment (i.e., birth cohort). Coefficients on controls and their interactions are not reported.

Finally, selection into SPA knowledge is proposed as a key driver of the first prediction: those who are more mistaken about their SPA have a smaller employment response to it. To more directly check this prediction, I reverse the format of the regression used in Tables XIII and XII. That is the outcome is the absolute size of the error in SPA self-reports at 58 or the closest age, and the regressor of interest is the interaction between being above SPA and the hazard of exiting employment controlling for each of the hazard and being above SPA separately as well as the age, period, and cohort dummies and the other controls used through this paper all interacted with the being above SPA. The results can be seen in Table XV. Column (1) shows results applying the drift correction, and Column (2) shows results without. Both versions indicate that exiting the labor force after SPA is associated with a reduction of SPA misbeliefs by about a fifth of a year.

A.6. *Savings and SPA Beliefs*

In the main text, it was shown that SPA belief data recorded when someone is in their 50s are predictive of employment responses to the SPA upon arriving at it in their 60s, in ways consistent with the model of costly information acquisition in this paper. The fact that SPA beliefs predict future labor supply behavior in ways consistent with the model provides some evidence for the model and also evidence that the ELSA's SPA belief data is not pure noise, as it is informative of behavior. In this section, I document that changes in those SPA beliefs are also predictive of contemporaneous saving behavior in ways one would expect. People who update their beliefs to a new belief that the SPA is later save more. Again, providing evidence that the belief data are not just pure noise, as they predict behavior in ways consistent with treating them as face-value assessments of beliefs.

TABLE XIV

INCLUDING BOTH TYPES OF HETEROGENEITY BY BELIEFS IN A SINGLE SPECIFICATION

	(1)	(2)	(3)	(4)
Over SPA	**0.185**	**0.180**	**0.268**	**0.216**
s.e	(0.0465)	(0.0335)	(0.0897)	(0.0562)
Over SPA × |SPA - SPA Self-report|	**-0.044**	**-0.029**	**-0.048**	**-0.045**
s.e	(0.0331)	(0.0125)	(0.0428)	(0.0215)
Over SPA × SPA above Self-report	**-0.156**	**-0.042**	**-0.094**	**-0.044**
s.e	(0.0993)	(0.0740)	(0.0993)	(0.0460)
Drift Corrected:	No	No	Yes	Yes
Fully Interacted:	Yes	No	Yes	No
Obs.	5,304	5,304	5,304	5,304

Note: The outcome variable is the hazard of exiting employment. Row 1 reports the treatment indicator for being over the State Pension Age (SPA), Row 2 reports interactions of this treatment indicator with the absolute difference between the true and self-reported SPA, and Row 3 with an indicator for the true SPA exceeding the self-reported SPA (i.e., under-estimation of the SPA). Self-reported SPAs are measured at age 58 or the closest available age, and wealth is measured in the last interview before the SPA. Column reports estimates from a specification of Equation 1 either fully interacted with both the absolute difference between the true and an indicator of under-estimating the SPA (Columns (1) and (3)) or with these variables added to the regression alonside the interaction terms reported in the table (Columns (2) and (4)). Columns (3) and (4) correct for the drift resulting from the difference in SPA belief question and object of interest, or not, is reported in the fifth row (i.e., it asks when you will arrive at SPA rather than your current age). Row 4 and 5 report respectively whether the drift correction has been applied and if the sample specification is fully interacted or not Following Abadie et al. (2023), standard errors are clustered at the level of the treatment (i.e., birth cohort). Coefficients on controls and their interactions are not reported.

TABLE XV

REGRESSION OF MISBELIEFS ON EMPLOYMENT RESPONSE AT SPA

	(1)	(2)
Over SPA × Hazard of Exiting Employment	**-0.207**	**-0.202**
s.e	(0.0791)	(0.0869)
Drift Corrected:	Yes	No
Obs.	5,304	5,304

Note: The outcome variable is the absolute size of the error in self-reported SPA at age 58 or the closest age recorded. The regressors of interest reported in this table are an interaction between being above SPA and the hazard of exiting employment, capturing the employment response to SPA. The other unreported regressors are being over SPA, the hazard of exiting employment, and the other controls and fixed effects used in other regression (age-period-cohort and demographic controls), interacted with being above SPA. Column (1) corrects for the drift resulting from the difference in SPA belief question and object of interest, or not, is reported in the fifth row (i.e., it asks when you will arrive at SPA rather than your current age), while Column (2) does not. Following Abadie et al. (2023), standard errors are clustered at the level of the treatment (i.e., birth cohort). Coefficients on controls and their interactions are not reported.

Table XVI shows this correlation. It shows the results of regressing changes in log NHNBW on a full set of quarterly age, period, and cohort dummies, the full set of controls used in the other regressions in this paper, and some measure of change of beliefs. Only the coefficients on change in beliefs are shown, as the other are treated as controls. Column (1) uses the change in SPA beliefs between one wave and the next; hence, it shows the resulting partial correlation between changes in log wealth and changes in SPA beliefs between one wave and the next. The results may be interpreted as a semi-elasticity, and we see that increasing your SPA belief by one year is associated with savings that result in a 4% increasing in total NHNBW. So, consistent with the responses to these questions reflecting people's true beliefs about the SPA, those who believe they have less State Pension wealth save more. We might expect a difference in response between those whose beliefs about the SPA have increased, and so believe themselves

TABLE XVI

REGRESSION OF DIFFERENCE IN LOG WEALTH ON SPA BELIEFS

Outcome variable:	(1) First Difference of Log Wealth	(2) Second Difference of Log Wealth
Δ(Error SPA)	**0.039**	**0.218**
s.e	(0.0185)	(0.1042)
Δ(Error SPA) × Δ(Error SPA) < 0	—	**-0.305**
s.e		(0.1787)
Obs.	1,752	858

Note: Regressions with difference in log wealth as the dependent variable and with SPA Beliefs as the main regressor of interest. All specifications include age, period, and cohort dummies, as well as the same controls as in the other regression. In column (1), SPA beliefs are measured as the change in self-reported SPA. In column (2), SPA beliefs are measured by an indicator equal to one if the change in self-reported SPA is positive (i.e., the respondent believes their SPA has increased).

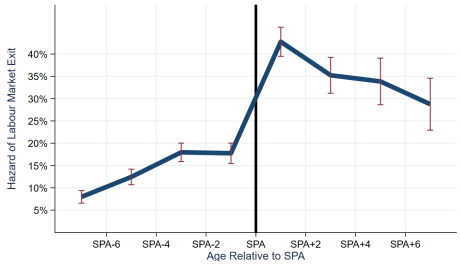
poorer, and those whose beliefs about the SPA have decreased, and so believe themselves richer. Since these are not causal estimates and these individuals are approaching retirement and often at their peak earning capacity, it seems unreasonable to expect them to dissave in response to a change in their SPA, but we might expect to see a change in how much they save. To investigate this, Column (2) of Table XVI constructs changes in saving as the second difference of log wealth and uses this as the outcome variable. It then adds to the regressors used in Column (1) an interaction between the change in SPA beliefs and the sign of that change. The first row shows that those who believe their SPA has gone up also increase their rate of saving, and row two shows that this effect is smaller for those who think it goes down, although this interaction term is only significant at the 10%. Hence, consistent with a face-value interpretation of people's responses, those who receive a negative pension wealth shock appear to increase their saving rate, while those who receive a positive shock appear to decrease it.

A.7. *Other Pension Belief and Knowledge Questions*

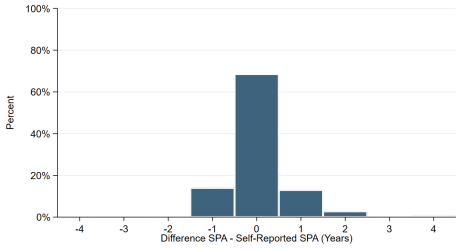
ELSA includes more data on beliefs about the State Pension and awareness of the reform than just SPA beliefs. I briefly discuss two question sets that seem highly relevant but are ultimately less useful than the SPA belief information.

From wave 3, ELSA asked if individuals were aware of the female SPA reform. Interestingly, a total lack of awareness of the reform does not drive SPA misbeliefs, with only 6.62% reported being unaware. While the unaware were more misinformed on average (mean error at age 58 of 1.4 vs. 0.9 years), the error distributions overlap. Moreover, dropping the unaware 6.62% does not materially change the patterns. Thus, I conclude that total unawareness is less informative than a nuanced view, allowing for partial misinformation and an intensive margin of information acquisition, as per the main text.

During a single wave (wave 3), ELSA collected subjective probability distributions on the level of pension benefits, but, as this was a single wave, using this data loses the panel dimension. Additionally, as those below SPA were asked these questions, the number of observations is very small.: 548 reported upper and lower bounds on expected State Pension income, and just 221 provided probabilities. Moreover, the complexity of the benefit formula makes it harder to identify mistakes than with SPA beliefs. While we cannot observe mistaken beliefs directly, the narrowing range of responses as people near SPA mirrors the decline in mean squared error in SPA reporting. The average expected income range drops from £14.48 at age 55 to £6.39 at 59. Given the small sample size, the difficulty of computing true entitlements, and the computa-



(a) Fraction exiting labor employment - Men



(b) Mistaken SPA Beliefs of Men at Age 58

Note: Panel (a): pooled average fraction exiting employment market at ages relative to the SPA. Data were plotted at two-yearly intervals because ELSA waves are biennial. Panel (b): plot of error in self-reported State Pension Age (SPA). The graph shows the frequency with which respondents gave mistaken answers about their SPA, with errors binned at the yearly level.

tional challenges of incorporating two sources of pension uncertainty into the model, I focus on SPA beliefs in this project.

A.8. *Men: Misbeliefs and Employment around SPA*

Due to the lack of policy variation, the employment response to SPA cannot be causally estimated for men. Thus, the main text focuses on how misbeliefs affect women's employment response. That does not mean that similar mechanisms are not at play for men.

Figure 7a shows a similar jump in men's hazard rate at SPA. While it is not possible to separate the SPA effects from aging, it is notable that the jump also occurs for men at SPA. Figure 7b shows mistaken beliefs for men at age 58. Despite no SPA reform and the male SPA remaining unchanged since 1948, nearly 40Though lower than the 60% for women, it supports the idea that misbeliefs are relevant in the absence of reform. If attention is costly, the mere possibility of reform could lead to mistaken beliefs. Thus, this evidence is consistent with the paper's proposed mechanism.

APPENDIX B: ADDITIONAL MATHEMATICAL DETAILS

B.1. *Finding Unique Actions Using Kuhn-Tucker Conditions*

Using the Kuhn-Tucker conditions of Equation 14 from the main text Caplin et al. (2019) provide an alternative formulation of the solution of the model. If the CCP satisfy Equation 13 from the main text and for all possible actions ($\forall d = (c, l) \in \mathcal{C}$)

$$\sum_{spa} \pi_t(spa) \frac{\exp\left(n^{(k)} \frac{((c/n^{(k)})^{\nu} l^{1-\nu})^{1-\gamma}}{\lambda(1-\gamma)} + \beta \overline{V}^{(k)}_{t+1}(d, X_t, spa, \underline{\pi_t})\right)}{\sum_{d' \in \mathcal{C}} q_t(d') \exp\left(n^{(k)} \frac{((c'/n^{(k)})^{\nu} l'^{1-\nu})^{1-\gamma}}{\lambda(1-\gamma)} + \beta \overline{V}^{(k)}_{t+1}(d', X_t, spa, \underline{\pi_t})\right)} \leq 1, \tag{16}$$

with equality if $q_t(d) > 0$, then the CCPs solve the model. This new condition from (Caplin et al., 2019) replaces the need for the unconditional choice probabilities to solve the log-sum-exp of Equation 14 from the main text.

If an action $d^\star = (c^\star, l^\star)$ satisfies Equation 15 from the main text repeated here:

$$\sum_{spa} \pi_t(spa) \frac{\exp\left(n^{(k)} \frac{((c/n^{(k)})^\nu l^{1-\nu})^{1-\gamma}}{\lambda(1-\gamma)} + \beta \overline{V}^{(k)}_{t+1}(d, X_t, spa, \underline{\pi}_t)\right)}{\exp\left(n^{(k)} \frac{((c^\star/n^{(k)})^\nu l^{\star 1-\nu})^{1-\gamma}}{\lambda(1-\gamma)} + \beta \overline{V}^{(k)}_{t+1}(d^\star, X_t, spa, \underline{\pi}_t)\right)} < 1, \tag{17}$$

for all $d = (c, l) \in \mathcal{C}$. That is, $d^\star$ produces such a high utility in all states that, in expectation, the exponentiated utility of any other payoff divided by its exponentiated utility is below 1.

If such a $d^\star$ exists then it automatically satisfies 16 to have $q_t(d^\star) = 1$, because substituting $q_t(d^\star) = 1$ into 16 yields 15 from the main text with an non-binding constraint.

APPENDIX C: ADDITIONAL COMPUTATIONAL DETAILS

C.1. *Solving the Models without Costly Attention*

The models are solved by backward induction starting at age 101 when the household dies with certainty. The household problem is modeled as a discrete choice. When rational inattention does not complicate this within-period discrete choice optimization, it is solved by grid search, selecting the value that maximizes the household's utility. States are discretized with 30 grid points for assets (a_t), 4 for average earnings ($AIME_t$), 5 for wages (w_t), two for the unemployment shock (ue_t), and in the model with policy uncertainty the state pension age (SPA_t) has 8 grid points as it ranges from 60 to 67.

A finer grid of 500 points is offered to the household when making their saving choice. This keeps the size of the state space manageable whilst not unduly constraining households and is equivalent to having a finer grid for consumption than for assets. When evaluating continuation values of off-grid values, I use linear interpolation of the value function.

C.2. *Solving the Models with Costly Attention*

Belief Distribution Costly attention introduces a high-dimensional state variable: the belief distribution ($\underline{\pi}_t$). To discretize it, I consider all ways of reallocating fixed-size probability masses across the eight possible SPAs (60-67). Since Bayesian updating cannot shift probability from zero, I want to avoid having beliefs assigning zero weight to any SPA in the belief grid. So, each SPA gets a minimum probability of 0.01, with the movable masses allocated on top.

Specifically, I use four movable probability masses. In the absence of the minimum probability requirement, each mass would be 0.25. With the minimum probability requirement, the size of the movable masses changes as the support of SPA_t changes. For example, in periods where all eight SPAs are possible (because $t < 60$ and the women have not aged past any possible SPA), these probability masses are of the size $\frac{1-0.08}{4} = 0.23$. These four masses distributed over eight SPAs yield $\binom{7+4}{4} = 330$ grid points (because each combination can be thought of as an ordering of the four masses and the seven breaks between the eight grid points). As individuals age and fewer SPAs remain, the grid shrinks—e.g., to $\binom{1+4}{4} = 5$ at $t = 65$ when only SPAs of 66 and 67 are possible. With no natural ordering over Δ^7, I cycle through combinations lexicographically. As robustness increases, I increase the number of movable probability masses to five, finding that it does not materially change results.

High Dimensional Interpolation When the prior with which a household starts the next period is off this grid, I use k-nearest neighbor inverse distance weighting to do multidimensional interpolation. I use the difference in means between the distributions as an approximation to the Wasserstein, or earth mover, metric as the distance concept in the inverse distance weighting. High-dimensional interpolation is computationally intensive and prone to approximation error. To mitigate this, I start with the two nearest grid points; if the fixed point loop for the unconditional choice probabilities (q_t) fails to converge within 25 iterations, I incrementally increase the number of neighbors up to a maximum of $2^8 = 256$.

Range of Attention Costs When rational inattention matters because $t < SPA_t$ the main equation to solve to find the household's optimal decision is:

$$\max_{\underline{q_t}} \sum_{spa} \pi_t(spa) \log\left(\sum_{d' \in \mathcal{C}} q_t(d) \exp\left(n^{(k)} \frac{((c/n^{(k)})^{\nu} l^{1-\nu})^{1-\gamma}}{\lambda(1-\gamma)} + \beta \overline{V}_{t+1}^{(k)}(d, X_t, SPA_t, \underline{\pi_t})\right)\right) \quad (18)$$

Following the random utility literature, I normalize the payoff in this equation by dividing it by the highest payoff across SPAs. However, the presence of λ complicates exponentiation. While data—not computation—should guide $lambda$'s value, its role as a denominator in the exponent causes exponentiated payoffs to diverge as λ shrinks, but these vanishingly small $lambda$'s values cannot be ignored. Since earlier SPAs are preferred, terms tied to SPA=60 are larger, and lower attention costs amplify these differences. Still, very small exponentiated payoffs associated with high SPAs when $\lambda issmall$ cannot be ignored: as $\lambda \to 0$, $\log()$ terms diverge to $-\infty$, and their gradients explode. Thus, even tiny exponentiated payoffs materially affect the objective function. Given this and the small attention costs implied by belief data, I carefully optimized code to retain very small utility values rather than dropping them—as might be acceptable with other objective functions. I use quadruple precision floating-point numbers to store the utility values (min value $\sim 10^{-4965}$), but since compilers are optimized for double precision, this greatly slows computation. So, I resort to quadruple precision only when necessary, checking first whether normalization causes underflow in double precision.

Solving the within-period problem Culling actions that are never chosen is central to the solution method. One of the two ways this is done is by dropping strictly dominated actions (for the other, see Section B.1). While identifying strictly dominated actions is an interesting problem studied in computer science (Kalyvas and Tzouramanis, 2017), the choice set here is modest (max 1,500 resulting from 3 labor and 500 asset options), so a simple Block Nested Loop algorithm is most efficient. When culling alone does not yield a solution, I solve Equation 18 using sequential quadratic programming (Schittkowski, 2014).

High-level Pseudo Code

Remove d from choice set $\mathcal{C}$ that are strictly dominated across all possible combinations of SPA_t and $\underline{\pi_{t+1}}$
if $|\mathcal{C}| = 1$ **then**
    Set $\underline{q_t}$ to degenerate distribution at unique $d \in \mathcal{C}$
else
    Set initial value of $\underline{\tilde{q}_t}$ and Error > Tolerance
    while Error > Tolerance **do**
        Solve for $\overline{V}_{t+1}$ (Equation 12 from the main text) given $\underline{\tilde{q}_t}$
        Remove d from $\mathcal{C}$ that are strictly dominated across all possible SPA_t given $\overline{V}_{t+1}$
        if $|\mathcal{C}| = 1$ **then**
            Set Error = 0 < Tolerance and $\underline{q_t}$ to degenerate distribution at $d \in \{d\} = \mathcal{C}$

else
  if there is an action d that satisfies 15 from the main text **then**
    Set Error = 0 < Tolerance and $\underline{q_t}$ to degenerate distribution at d
  else
    Solve 14 from the main text using sequential quadratic programming for $\underline{q_t}$
    Set Error to distance between $\underline{q_t}$ and $\underline{\tilde{q}_t}$
    Update $\underline{\tilde{q}_t} = \underline{q_t}$
  end if
end if
end while
end if
Substitute $\underline{q_t}$ into 13 from the main text to solve for $\underline{p_t}$.

C.3. *Simulating and Estimating*

My simulated sample consists of 50,000 randomly drawn individuals aged 55. Since the simulated sample exceeds the data size, state variables initialized directly from the empirical distribution (assets and average lifetime earnings) are sampled multiple times using random Monte Carlo draws from their joint distribution. I initialize wages with draws from its estimated distribution. I simulate one SPA cohort at a time, setting SPA_t to match the cohort's SPA. I assume the SPA response reflects draws from the individual prior belief distributions, with every one of the same type starting at age 55 with identical beliefs. Thus, I initialize beliefs using the type-specific distribution of SPA point estimates.

Given these initial conditions, I simulate the choice of the individual households using the decision rule found when solving the model and the exogenous process estimate in the first stage. I aggregate the simulated data in the same way as with observed data to construct the moment conditions, detailed in Appendix D. The method of simulated moments estimates model parameters by minimizing a GMM criterion, also described in Appendix D. To minimize the objective function, I first sample the parameter space via Sobol sequencing, then apply the BOBYQA routine (Powell, 2009) at promising starting points.

APPENDIX D: ADDITIONAL ECONOMETRIC DETAILS

D.1. *Imputing AIME*

Average lifetime earnings are observed only for women present in wave 3 who consented to link their National Insurance records. To initialize the model from the joint distribution of $AIME_{55}$ and a_{55} without introducing selection into a_{55}, I impute missing values. I first regress $AIME_{55}$ on a quintic in NHNBW and a rich set of controls, including variables on health, education, location, labor market behavior, housing tenure, cohort, age, wage, and cognitive ability.

Using predicted values alone would overstate the correlation between $AIME_{55}$ and a_{55}, so I add noise to the imputations to match observed heteroscedasticity. I regress non-imputed $AIME_{55}$ on a quintic in NHNBW (excluding controls, as they're absent in the model), then regress the squared residuals on the same polynomial. Since imputed $AIME_{55}$ is homoscedastic by construction, adding noise with variance from the second regression replicates the observed heteroscedasticity.

D.2. *Type-specific Mortality*

Life expectancy heterogeneity affects older individuals' behavior (e.g. De Nardi et al., 2009), but death is often poorly recorded in surveys. To address this, I estimate type-specific mortality without relying on ELSA death records, instead combining them with ONS survival tables following French (2005). I do this using Bayes' rule:

$$Pr(death_t|type=k) = \frac{Pr(type=k|death_t)}{Pr(type=k)} Pr(death_t).$$

Where $Pr(type=k|death_t)$ and $Pr(type=k)$ are taken from ELSA and $Pr(death_t)$ are taken from the ONS life-tables. If measurement error affects all types equally, estimates of $Pr(type=k|death_t)$ from ELSA are unbiased, unlike those of $Pr(death_t|type=k)$, and deal with the measurement error issue.

D.3. *Generating Profiles*

To avoid contamination by cohort effects or macroeconomic circumstances, targeted profiles were generated with a fixed effect age regression, which included: year of birth effects, SPA-cohort specific age effects, the aggregate unemployment rate rounded to half a percentage point, and an indicator of being below the SPA. Specifically, the following regression equation was estimated:

$$y_{it} = U_t + \sum_{c\in C} \gamma_c \mathbb{1}[cohort_i = c] + \sum_{s\in S} \mathbb{1}[SPA_i = s](\sum_{a\in A} \delta_{a,s} \mathbb{1}[age_{it} = a])$$

where $cohort_i$ is the year-of-birth cohort of an individual, SPA_i is her final SPA, $age_{i,t}$ her age in years, U_t aggregate unemployment to half a percent, and the outcome variable y_{it} is either assets or employment depending on which profile is being calculated. The profiles targeted were then predicted from these regressions using average values for the pre-reform cohorts.

APPENDIX E: Additional Results

E.1. *First Stage Estimates*

Model Types A woman is classified as highly educated if she exceeds the compulsory schooling for her generation. She is considered married if married or cohabiting, as household structure matters more than legal status for the questions considered in this paper. As the model abstracts from separation, any woman ever observed as married is treated as married in all periods. This accounts for the likely receipt of alimony or a widow's pension, making 'married' the most model-consistent classification for previously married women. The resulting type shares are: 34% married/low education, 11% single/low education, 44% married/high education, and 11% single/high education.

Initial conditions Initial assets a_{55} and average earnings $AIME_{55}$ are drawn from the type-specific empirical joint distribution (summary statistics in Table XVII). As expected for this generation, married women have weaker labor market attachment, resulting in lower $AIME_{55}$ but higher household assets. Higher education raises both variables.

TABLE XVII

SUMMARY STATISTICS OF INITIAL CONDITIONS (£)

Type	Variable	Mean	SD
Married, Low Education	Initial Assets	76,226	163,320
	Initial AIME	4,889	2,915
Single, Low Education	Initial Assets	13,231	30,471
	Initial AIME	6,015	4,334
Married, High Education	Initial Assets	148,440	218,143
	Initial AIME	9,358	6,264
Single, High Education	Initial Assets	97,495	186,362
	Initial AIME	10,663	6,676
...total	Initial Assets	102,680	189,801
	Initial AIME	7,618	5,199

Note: Means and standard deviations of the initial distribution of assets and average lifetime earnings.

TABLE XVIII

TYPE SPECIFIC UNEMPLOYMENT TRANSITION PROBABILITIES

Type	Transition	Probability(%)
Married, Low Education	From employment to unemployment	2.37
	From unemployment to employment	57.75
Single, Low Education	From employment to unemployment	3.20
	From unemployment to employment	27.03
Married, High Education	From employment to unemployment	1.72
	From unemployment to employment	71.08
Single, High Education	From employment to unemployment	3.25
	From unemployment to employment	37.78

Note: Unemployment and reemployment transition probabilities.

Labor market conditions Type-specific transition probabilities—estimated by classifying individuals as unemployed when claiming benefits—are shown in Table XVIII. Parameters of the stochastic wage component (persistence, innovation variance, measurement error, and initial draw) appear in Table XIX. The deterministic wage component generates the profiles in Figure 8a. Spousal income is shown in Figure 8b

Social Insurance As noted in the main text, State Pension income differs more by marital status than by education. Among State pension claimants, high-education women receive £92.52 on average, low-education £87.11, while single women receive £112.50 and married women £80.89. To capture this key distinction and maximize power, I restrict State Pension heterogeneity to marital status only. The resulting functions of average lifetime earnings are shown in Figure 9a."

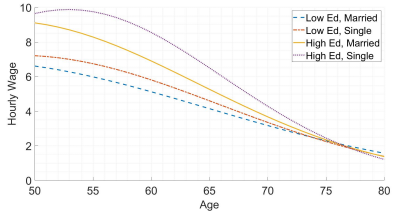
Conversely, private pension income varies more by education than by marital status. Among those with non-zero private pension income, high-education women receive £118.50 on average, low-education £61.42, while single women receive £100.78 and married women £94.24.

TABLE XIX

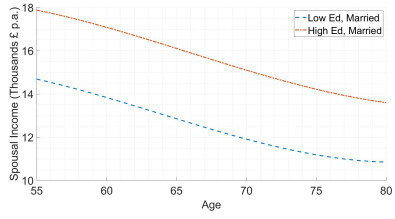
PARAMETERS OF THE STOCHASTIC COMPONENT OF THE WAGES

Type	ρ_w	σ_ϵ	σ_μ	$\sigma_{\epsilon,55}$
Married, Low Education	0.911	0.039	0.249	0.266
Single, Low Education	0.901	0.042	0.255	0.178
Married, High Education	0.945	0.035	0.351	0.322
Single, High Education	0.974	0.025	0.358	0.224

Note: Notes: Estimates of the persistence of wages and the variance of their transitory and persistent components as well as initial distribution.

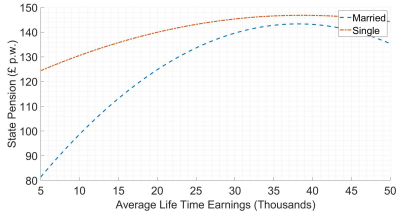


(a) Wage Profiles

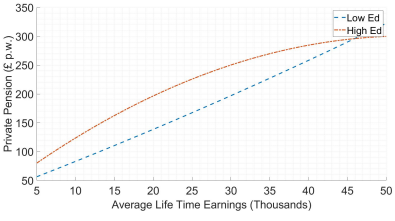


(b) Spousal Income

Note: Panel (a): the deterministic component of female hourly wages for the four model types plotted against female age. Panel(b): spousal income plotted against female age.



(a) State Pension Function



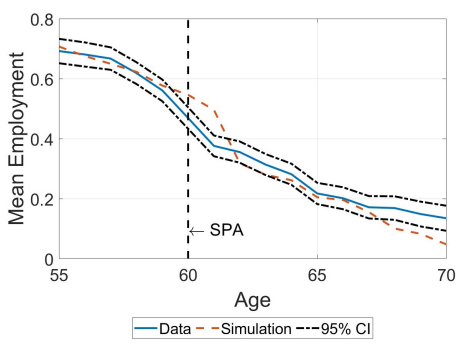
(b) Private Pension Function

Note: Panel (a): state Pension income as a function of average lifetime earnings (AIME) for married and single women. Panel(b): Private Pension income as a function of average lifetime earnings (AIME) for high and low-education women.

To capture this key difference and maximize power, I restrict private pension heterogeneity to education only. The resulting functions are in Figure 9b.

E.2. *Model Fit*

As mentioned in the main text, although the different model specifications have different predictions about the labor supply response to the dynamic SPA, the static profiles are not very sensitive to model specifications. All versions are able to match the static profiles. Figures 10a-11b show the employment and asset profiles for the baseline and the version with rational inattention with the parameter estimates of Table 3d from the main text.

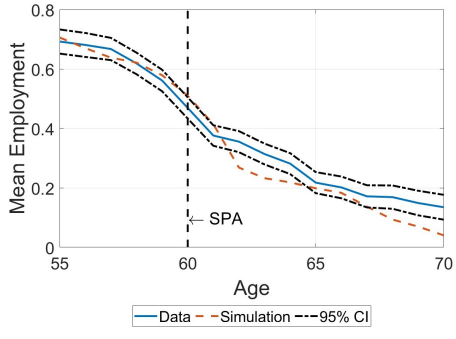


(a) Employment Profile Baseline

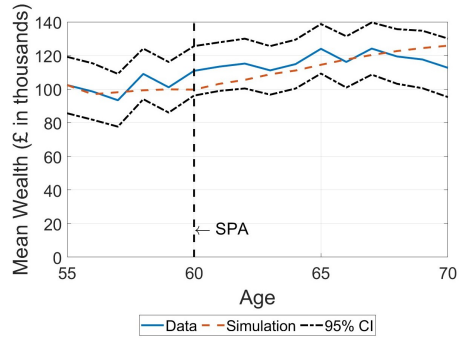


(b) Asset Profile Baseline

Note: Panel (a): model fit to targeted labor supply profile. The empirical profile is for the pre-reform SPA cohort with a SPA of 60. The model was simulated with an unchanging SPA of 60, mimicking the conditions faced by this cohort. Panel (a): model fit to targeted asset profile. The empirical profile is for the pre-reform SPA cohort with a SPA of 60. The model was simulated with an unchanging SPA of 60, mimicking the conditions faced by this cohort.



(a) Employment Profile Model with Rational Inattention



(b) Asset Profile Model with Rational Inattention

Note: Panel (a): model fit to targeted labor supply profile. The empirical profile is for the pre-reform SPA cohort with a SPA of 60. The model was simulated with an unchanging SPA of 60, mimicking the conditions faced by this cohort. Panel (b): model fit to targeted asset profile. The empirical profile is for the pre-reform SPA cohort with a SPA of 60. The model was simulated with an unchanging SPA of 60, mimicking the conditions faced by this cohort.

TABLE XX

SUMMARY STATISTICS OF ATTENTION COST CONVERTED TO COMPENSATING ASSETS (£)

λ	Mean	S.D	Median	25th-Percentile	75th-Percentile	Semi-elasticity (per 10k Assets)
6×10^{-8}	£11.00	£9.00	£9.00	£6.00	£14.00	-1.82%
5×10^{-7}	£31.00	£32.00	£22.00	£15.00	£33.00	-2.7%

Note: Distribution of compensating assets equivalent to the utility of learning your SPA today, shown for two attention costs.

E.3. *Reference Point Retirement Literature*

Seibold (2021) supports reference dependence through a process of elimination, ruling out alternatives. He rejects misbeliefs as an explanation, on the basis that less-educated individuals,

TABLE XXI

IMPACTS OF REFORMING SPA WITH INFORMED AND UNINFORMED HOUSEHOLDS (FRACTION OF PASSIVE HOUSEHOLDS)

SPA increased from 60 to:	(1) - Informed Added Employment	(2) - Uninformed Added Employment
61	0.16	0.13
62	0.30	0.26
63	0.39	0.33
64	0.48	0.41
65	0.62	0.52

Note: Results of raising SPA from 60 to the age in Column (1) with costly attention and in Column (2) without it.

who he argues are more prone to confusion, show a smaller employment response at eligibility. While they likely have higher processing costs, they may also be more incentivized to learn about this particular issue, dedicating more resources to it (something my model implies and belief data supports). Lalive et al. (2023) provide survey evidence in support of reference dependence, finding that eligibility is the main reason for stopping work and that many claim benefits simply because "it seems natural." Mapping survey responses to model construct is difficult, however. Eligibility could be interpreted as an implicit recommendation in the presence of costly attention, leading people to describe claiming at this age as natural. Gruber et al. (2022) presents compelling Finnish evidence: relabeling a pension age, despite minimal financial changes, caused significant employment shifts. On the one hand, this seems to strongly support reference dependence, yet many who retired due to relabeling later returned to work, which they interpret as suggesting regret. In contrast, inattention could explain both phenomena. Confusion about the relabeling prompts exit, while belief updates drive re-entry. As Gruber et al. (2022) note," since [information about optimal retirement] is always attached to the statutory age itself, it is difficult to disentangle this effect empirically". I would add the caveat without gathering belief data.

E.4. *Introducing a Fraction of Passive Agents*

As a simple way of capturing a behavioral bias like reference dependence preference, I introduce a fraction of passive agents that retire at SPA but do not anticipate this fact, in the vein of (as in Chetty et al., 2014). I use this fraction to match the employment responses to the SPA of the whole population and the richer subgroup. I find that in the model with only policy uncertainty, I need 14% to be passive to match the data (treatment effect 0.119 and 0.108 for above median assets) but due to the better initial fit of the rationally inattentive version, only 10% (treatment effect 0.118 and 0.122 for above median assets)in that version of the model. Table XXI shows employment responses to SPA increases in the two versions of the model with the different fraction of passive agents required to match the treatment effects. Since the difference between these two columns is not just being informed or not (because the fraction of passive agents changes, it does not make sense to analyze the impact of an information letter campaign as is done in the main text. In this table, we see that due to the mechanical effect of a fraction of passive agents, the additional employment from increasing the SPA is larger, but the relative difference between the two columns is similar to that found in the experiment without passive agents in the main text.

APPENDIX F: EXTENSION: DEFERRAL PUZZLE

Attributing all policy uncertainty to the stochastic State Pension Age (SPA) understates overall pension uncertainty. This section introduces uncertainty and learning about another key feature: the actuarial adjustment from deferring benefits. Combined with a claiming decision, this addition improves realism and helps explain the deferral puzzle (discussed below). Since the adjustment rate becomes irrelevant after claiming, rational inattention speaks directly to this puzzle. While deferral may appear actuarially favorable, this overlooks the benefit of claiming due to removing the need to monitor the adjustment rate. The model in Section 5.2 omits this mechanism for two reasons: it lacks a benefit-claiming decision and assumes SPA is the only uncertainty incurring attention costs, and once SPA is reached, this uncertainty resolves, irrespective of claiming. The rest of this section presents a simple extension that introduces this new incentive and its implications.

F.1. *Deferral Puzzle*

By deferral puzzle, I refer to the rarity of deferring state pension benefits despite highly generous terms between April 2005 and April 2016. Between those dates, benefits rose by 1% for every 5 weeks deferred—an annual adjustment of 10.4%. This is an extremely generous actuarial adjustment, and yet, 86.7% of women observed past SPA in ELSA during this period had claimed by their first post-SPA interview.

What exactly constitutes actuarially fair depends on life expectancy and the interest rate, but at all plausible levels, this adjustment was generous. For women reaching SPA during this period, life expectancy ranged from 23 to 25 years. Even using a conservative estimate of 23 years, a 10.4% annual adjustment was advantageous at interest rates up to 9%. The Bank of England base rate never exceeded 5.75% and was 0.5% from March 2009 onward. Thus, the 10.4% adjustment was actuarially favorable at any realistic rate.

Even among the few women who deferred, deferral durations were short. With a conservative life expectancy of 23 years and a 5.75% interest rate, the optimal deferral is 9 years. Yet, the median deferral was 2 years, and 99.54% claimed within 8 years.

These calculations use mean life expectancy, which masks heterogeneity. However, heterogeneity alone is not a plausible explanation. It would mean 86.7% of women had significantly below mean life expectancy, implying implausible skewness.

F.2. *Model and Estimation*

Benefit claiming is a binary decision, and having claimed is an absorbing state: once an individual claims the state pension, they cannot unclaim. Benefit claiming is only an option once past the SPA. To keep the problem tractable, an upper limit of 70 is placed on deferral.

Stochastic deferral adjustment is modeled as iid with two points of support. Having only two points of support limits the growth of the state space resulting from solving the model with different values of the adjustment rate to a factor of two. Having the uncertainty be iid means that beliefs do not enter as a state variable. Instead, the true probabilities form beliefs in each period: yesterday's learning is not relevant to today's state of the world. This also avoids a fundamental identification problem as there is no data on beliefs about adjustment rates. As benefit claiming is an absorbing state, an indicator of having claimed or not also expands the state space.

The two points of support are chosen as 10.4% and 5.8%, the actuarial adjustment from 2006 to 2016 and post-2017, respectively. The probability of being offered the higher actuarial adjustment of 10.4% is chosen to match the average actuarial adjustment since 1955, resulting in

TABLE XXII

PARAMETER ESTIMATES - EXTENSION

ν: Consumption Weight	0.5310
β: Discount Factor	0.9852
γ: Relative Risk Aversion	2.0094
θ: Warm Glow bequest Weight	20,213

Note: Estimated parameters from method of simulated moments for the model extension with a stochastic deferral rate and a benefit claiming decision.

TABLE XXIII

MODEL PREDICTIONS - EXTENSION WITH BENEFIT CLAIMING AND UNCERTAIN DEFERRAL

	Costly Attention	Data
Population	Treatment Effect for being below SPA on employment	
Whole Population	0.0416	0.080
Assets >Median(£29,000)	0.0903	0.061
Age	Variance of SPA Answers	
55	2.985	2.852
58	1.795	1.180
Coefficient	Treatment Effect Heterogeneity by SPA Error	
Treatment Effect	0.0532	0.157
Interaction	-0.0111	-0.023

Notes: Costly attention refers to the model with, additionally, a cost of information acquisition about the stochastic policy. The top panel shows labor supply response across the wealth distribution in a similar spirit to Table II from the main text. The second panel shows the reduction in self-reported SPA between 55 and 58. The bottom panel shows, in the interaction term, the heterogeneity in labour supply response to the SPA by self-reported SPA error at age 58.

a probability of 0.415. Deferral rules are taken from Bozio et al. (2010), and since earlier deferral rules were previously stated in absolute rather than percentage terms, the ONS time series of state pension spending going back to 1955 is used to work out implied average percentage deferral adjustments.

The model with policy uncertainty, to the stochastic SPA and adjustment rate, is then re-estimated to match the same pre-reform employment and assets profiles with a constant realization of 10.4% for the deferral adjustment, which was the deferral rate these cohorts faced. Parameter estimates are in Table XXII and, for these values, only 6.2% of individuals claim the state pension before the mandatory claiming age of 70, much lower than the 99% plus claiming seen in the data.

Next, I introduce costly attention with a cost of attention corresponding to approximately £10 of consumption to the median consuming household to be fully informed. This increased the number voluntarily claiming to 22.2%, approximately a fourfold increase on the model without informational frictions, but still short of the rate observed in the data. As can be seen in Table XXIII, this cost of attention produced a relatively good fit in terms of employment response to the SPA and patterns in the belief data. Note that the treatment effect displayed is the effects of being below the SPA on the probability of being in employment, rather than being above the SPA on the hazard of exiting employment used in the main text.