Network for Studies on Pensions, Aging and Retirement

NETSPAR ACADEMIC SERIES

Costly Attention and retirement puzzles

Jamie Hentall MacCuish

DP 01/2020-006

Costly Attention and Retirement Puzzles

Preliminary and Incomplete *

Jamie Hentall MacCuish†

Abstract

This paper presents evidence that accommodating costly attention, modelled with rational inattention, and inherent policy uncertainty might solve three well-established puzzles in the retirement literature. The first puzzle is that, given incentives, the extent of bunching of labour market exits at legislated state pension ages (SPA) seems incompatible with complete information (e.g. Cribb, Emmerson, and Tetlow, 2016). Adding to the evidence for this puzzle, this paper includes an empirical analysis focusing on whether liquidity constraints can account for this bunching and find they cannot. The nature of this puzzle is clarified by exploring a life-cycle model with informed rational agents that does match aggregate profiles. This model succeeds in matching these aggregates only by overestimating the impact of the SPA on poorer individuals whilst underestimating its impact on wealthier people. The second puzzle is that people are often mistaken about their own pension provisions (e.g. Gustman and Steinmeier, 2001). Concerning this second puzzle, I incorporate rational inattention to the SPA into the aforementioned life-cycle model, thus allowing for mistaken beliefs. To the best of my knowledge, this paper is the first not only to incorporate rational inattention into a life-cycle model but also to assess a rationally inattentive model against non-experimental individual choice data. This facilitates another important contribution to the rational attention literature: discipling the cost of attention with subjective expectations data. Preliminary results indicate costly attention better matches the response of participation to the SPA across the wealth distribution, hence offering a resolution to the first puzzle. The third puzzle is that despite more than actuarially fair options to defer receipt of pension benefits in some countries, take up remains extremely low (e.g. Shoven and Slavov, 2014). A simple extension of the main model generates, by extending the source of uncertainty and including a claiming decision, an explanation of this last puzzle: the actuarial calculations implying deferral being preferable ignore the utility cost of paying attention to your pension which can be avoided by claiming. This paper researches these puzzles in the context of the ongoing reform to the UK female state pension age.

Keywords: Rational inattention, Labour supply, Retirement, Pension provision, Learning

JEL Classification: D14, D83, D91, E21, J26, H55

*I thank Fabien Postel-Vinay and Eric French for detailed discussions and encouragement. I also thank George-Marios Angeletos, James Banks, Antonio Cabrales, Andrew Chesher, Monica Costas Dias, Jonathan Cribb, Martin Cripps, Mariacristina De Nardi, Rory McGee, Robert Moffitt, Cormac O'Dea, Ran Spiegler, Morten Ravn, and participants at a number of seminars for their helpful comments. At different times funding for this work has been received from ESRC Grant Inequality and the insurance value of transfers across the life cycle (ES/P001831/1) and ESRC studentship (ES/P000592/1), and this funding is gratefully acknowledged.

†University College London and Institute for Fiscal Studies, correspondence to jamie.maccuish.15@ucl.ac.uk

1 Introduction

This paper explores retirement decisions and whether accommodating costly attention and inherent policy uncertainty explains observed regularities better than the dominant complete information paradigm, which abstract away from these costs. In particular, it focuses on three puzzles for a complete information explanation of retirement choices. Firstly, it is difficult to reconcile the large number of labour market exits at legislated pension entitlement ages with RE. Secondly, that individuals are frequently mistaken about their own retirement provisions. Thirdly, when deferral incentives are more than actuarially fair pension, deferral rates remain very low. I find that acknowledging costly thought accommodates the second puzzle whilst generating mechanisms that help explain the first and third.

Ageing populations have forced many governments to increase the state pension age. These reforms revealed the first puzzle: labour market exits are more sensitive to legislated state pension ages than standard models with complete information (SM) can accommodate given the limited incentives to retire exactly at these ages (Behaghel, and Blau, 2012; Cribb, Emmerson, and Tetlow, 2016; Seibold, 2019; Lalive, Magesan and Staubli, 2017). This paper contributes to the evidence for this puzzle, within the context of the ongoing reform to the UK female state pension age (SPA), by studying the dependence of the response to the SPA on asset holdings. I find an indistinguishable participation response to the SPA across the wealth distribution largely which largely rules out liquidity constraints as a SM consistent explanation of this puzzle.

This claimed oversensitivity is in tension with studies that successfully match observed retirement decisions within SM (French, 2005; French and Jones, 2011; O'Dea, 2018), although these are not studies of a pension reform. One contribution of this paper is to offer a resolution of this tension. This is achieved by investigating whether a SM that successfully matches aggregate labour market profiles around SPA also matches the response to changes in the SPA across the wealth distribution. I find that although the model matches the static aggregate profile it fails to match the dependence of retirement on the SPA both in terms of magnitude of response and of the correlation of this response and wealth.

The second puzzling regularity at odds with the SM is that individuals are frequently mistaken about their pension provision (Gustman and Steinmeier, 2001; Manski, 2004; Rohwedder and Kleinjans, 2006; Crawford and Tetlow, 2010; Amin-Smith and Crawford, 2018). Traditionally SM life-cycle models treat institutional factors like the SPA as static parameters known, without cost, by everyone; an approach which precludes any explanation of why people are systematically mistaken. By acknowledging the stochasticity of government pension policy and incorporating costly thought, this paper attempts to explain this observed ignorance. Costly thought is modelled using rational inattention (RI), an approach that includes a utility cost of information acquisition.

Allowing for these incorrect beliefs explains the bunching of labour market exits with greater success than the benchmark model. The mechanism behind this result is as follows: rational inattention to the SPA introduces additional uncertainty implying greater precautionary saving which leads to greater labour market participation. As this uncertainty is resolved upon reaching SPA, this induces bunching of labour market exit at the SPA. This mechanism is not strictly dependent on rational inattention and exists with the introduction of a stochastic SPA alone. However, for reasonable levels of stochasticity that match how the government has historically reformed the SPA, I find that the amplification of uncertainty inherent in rational inattention is required to produce a discernible difference from the baseline.

The third puzzle for a explanation of retirement behaviour consistent with SM is that despite clear incentives to defer in both the UK and USA, deferral is very uncommon in both countries (Shoven and Slavov, 2012; Shoven and Slavov, 2014; Crawford and Tetlow, 2010). The calculation showing that deferral is more than actuarially fair do not, however, take account of the attention cost of tracking pension entitlements. Once we allow for costly attention claiming immediately can be optimal as claiming removes one demand on our limited attention: tracking pension entitlements. I present an extension to my main model with that includes a mechanism capable of generating such results.

The reform to the UK female SPA provides me with the opportunity to investigate these question. This reform has a staggered implementation which creates individual level variation in SPAs allowing the effect of the SPA on employment to be identified separately from effects of ageing. Additionally, the UK institutional context has two features advantageous to identifying motivations behind retirement decisions. Firstly, receipt of the UK state pension is not conditional on employment status and only provides an incentive to retire for liquidity constrained individuals. Secondly, forcing someone to retire purely due to age is illegal, ruling out firm mandated retirement as an explanation for the bunching of labour market exits. The dataset I will use is the English Longitudinal Study of Ageing (ELSA) which is a detailed panel survey of older individual.

The rest of the paper is structured as follows. In section 2, I review the literature and, in section 3, I outline the institutional context and the data used. In section 4, I present evidence of the three puzzles in my dataset focusing on the first puzzle that bunching of labour market exits at SPA cannot be explained with SM. In particular, as other SM explanation are largely ruled out by the UK institutional context, I investigate whether liquidity constraints can explain this bunching and find indications they cannot. To see if some SM compatible mechanism might explain the bunching, in section 5, I take a rich structural model that matches aggregate profile and investigates its mechanisms. I find that it matches the aggregate profiles in the pre-reform sample but fails to match the dependence on the SPA. It fails to match both the magnitude of the response to the SPA and its correlation with wealth. This model then serves as a baseline into which I incorporate rational inattention in section 6. This rational inattention to the SPA mechanically explains the second puzzle but in section 5 I present preliminary result from a calibrated version of the model that the introduction of RI better approximates both the magnitude of the response to the SPA and its distribution over assets thus offering a explanation of the first puzzle. In section 7 I present an extension that increases the source of uncertainty that demand agents attention and introduces a claiming decision and argue that this generates a mechanism that can help explain the third puzzle. Section 8 concludes.

2 Literature Review

This paper contributes to the literatures identifying the three puzzles mentioned earlier by proposing a solution to them. This solution brings together two disparate literatures, on rational inattention and on quantitative retirement models, hence making a contribution to each of these. In bringing these two literatures together this paper is the first work, to the best of the author's knowledge, to apply rational inattention to non-experimental individual choice data. This work can also be more broadly construed as part of the growing literatures on behavioural public economics or the economic implications of attention.

The first puzzle is that labour market participation appears more sensitive to SPA than SM predicts given the incentives. It has long been known that labour force participation responds strongly to the eligibility ages of social security programs. Gruber and Wise (2004) survey evidence from 11 developed countries and find labour force exits concentrated around legislated retirement ages. The response of labour market participation to reforms in the social security eligibility ages has been widely studied (Mastrobuoni, 2009; Staubli and Zweimuller, 2013; Manoli and Weber, 2012; Atalay and Barrett, 2015). Often this literature has focused on estimating elasticities and fiscal impacts; however, recently a group of studies has argued that SM are unable to match the magnitude of labour market exits at legislated pension ages; this is the principal puzzle this paper attempts to explain (Behaghel, and Blau, 2012; Cribb, Emmerson, and Tetlow, 2016; Seibold, 2019; Lalive, Magesan and Staubli, 2017). Lalive et al. (2017) study a Swiss reform that increased women's full retirement age (FRA), the age they can claim their full pension, and introduced an Early Retirement Age (ERA), the earliest eligibility age to a reduced pension. They find, whilst incentives encourage most women to claim early, most delay. Seibold (2019) studies retirement in Germany where there exists an ERA, a FRA, and a NRA (normal retirement ages, a purely nominal age), along with multiple pathways into retirement; this arrangement creates kinks and notches in life-time budget constraints. These non-differentiabilities can be classified into statutory retirement ages where some expectation of retirement exists (i.e. ERA, FRA and NRA) and pure financial incentives. Siebold (2019) documents that the bunching of labour market exits is higher at the non-differentiability associated with a statutory retirement age rather than those only associated with a financial incentive and interprets this as evidence for reference dependent preference. Both Seibold (2019) and Lalive et al. (2017) present structural model that attempt to explain these findings. Cribb et al. (2016) study the same increase to the UK female SPA as this paper. Since the UK state pension age is the earliest that the pension can be claimed it is an ERA for international comparison. They produce reduced form estimates of the impact of reaching ERA on labour force participation and, although their paper is purely empirical, argue that their findings are difficult to explain with RE. As I borrow from and build on the methodology of Cribb et al. (2016) their methods are discussed in more detail in section 4.

The second puzzle, that people hold mistaken beliefs about their own pension provision is well documented going back to, at least, Bernheim 1987. He finds individuals' social security forecast errors are systematically related to information individuals posses. Gustman and Steinmeier (2001) compare reported expected benefits to objective calculations based on social security records and employer provided pension description and find misinformation the norm. Manski (2004) finds most uncertainty about future social security entitlement can be explained by a lack of understanding of current social security arrangements.Rohwedder and Kleinjans (2006) study the dynamics of these mistaken beliefs and find that their expectations become increasingly accurate as individuals approach retirement. Finally, Crawford and Tetlow (2010) look at women subject to the UK female SPA reform and find they hold substantially incorrect beliefs about their own SPA; Amin-Smith and Crawford (2018) update this analysis finding broadly similar results. An advantage of the structural model presented here over those of Seibold (2019) and Lalive et al. (2017) is that it accommodates both the fact people hold incorrect beliefs and that labour market exits bunch at SPA. I also rely on a different mechanism, rational inattention to the SPA, rather than reference dependent preferences (Seibold, 2019) or controlling the proportion of agents being unable to choose when to retire (Lalive et al., 2017).

The literature on the third puzzle is, to the best of my knowledge, more recent. This is because what constitutes an actuarially favourable pension deferral is dependent on mortality probabilities and interest rates. Shoven and Slavov (2014) look at the US and find that although whether deferral was beneficial or not in the 1960s was dependent on family status, by the 21st century deferral was preferable for nearly all the groups considered. This change can be attributed to increased life-expectancy and historically low interest rates. Shoven and Slavov (2012) extend this analysis to look at optimal deferral choices for individuals allowing for heterogeneous survival probabilities and find that most households – even those with mortality rates that are twice the average – benefit from some delay. Despite these strong incentives to defer, these paper find in the US most individuals appear to claim Social Security soon after they reach the eligibility threshold age. This same puzzle has been observed in the UK context that I study. Crawford and Tetlow (2010) find that only $\approx$2% take up this option to defer receipt of the state pension, despite it being more than actuarially fair.

Addressing these puzzles, this paper brings together two disparate literatures. One being on rational inattention, developed by Sims (1998; 2003; 2006), which originally found application to macroeconomic problems, exemplified by Maćkowiak and Wiederholt (2009; 2015) explanations of differential business cycle responses and of sticky prices. Recently its domain of application has expanded for example: in a decision theory context, Caplin and Dean (2015) develop a revealed preference test that characterizes all patterns of choice "mistakes" consistent with a generalisation of RI; Ravid (2018) applies RI in a game theory context; Bartoš et al. (2016) apply RI to explain job market discrimination in a field experiment; and Matejka and McKay (2014) and Steiner, Stewart, and Matejka (2017) prove results linking RI to a logit model of choice. This paper contributes to this expansion of the domain of application of RI by applying it to a life-cycle model of saving and labour supply and in doing so it is the first work, to the best of my knowledge, to apply rational inattention to non-experimental individual choice data. To achieve this I make extensive use of the theoretical advance made by Steiner, Stewart, and Matejka (2017) who solve a general class of dynamic discrete choice problems.

The second literature this paper makes a substantial contribution to is that on quantitative retirement models. Early work in this literature, such as Gustman and Steinmeier (1986) and Burtless (1986), abstracted away from uncertainty and borrowing constraint but more recent work take these into account. Rust and Phelan (1997) introduce incomplete markets. French (2005) estimate a model that allows for dynamics, health transitions and fixed costs of work to explain the retirement phenomena and French and Jones (2011) allow for medicare cost and greater individual heterogeneity. Imrohoroglu and Kitao (2012) bring in general equilibrium and Erosa, Fuster, and Kambourov (2012) use quantitative retirement models to make cross-country policy comparisons.The paper most related to the current one from this literature is O'Dea (2018) who estimates a structural retirement model in the UK context to investigate the differing pension provision policies. As well as contributing to this literature by integrating new modelling techniques from RI, this paper is also contribute to this literature by comparing the predictions of these model to the ex-post data from a pension age reform.

More broadly this paper fits into the behavioural public economics literature and the growing literature on attention. This paper speak to the call in Gabaix (2019) for more structural estimation in the attention literature.

3 Data and Institutional Context

The context for this research project is the UK reform to the female state pension age (SPA). The Pensions Act 1995 legislated for the female SPA to rise gradually from 60 to 65 over the ten years from April 2010, rising by one month every two months, reaching age 65 by April 2020. The Pension Act 2011 accelerated the rate of change of the female SPA from April 2016 so that it equalises with men's at 65 by November 2018. It additionally legislated an increase to both the male and female SPA to 66 years phased in between December 2018 to October 2020. Figure 1 from Cribb et al. (2016) summarises how these changes affect women in different birth cohorts.

An advantage of studying the this reform is UK law prohibits age-based compulsorily retirement; employment equality regulations introduced in 2006 classified forcing employees to retire purely based on age as discrimination. So firms mandate retirement cannot explain the bunching of labour market exits at SPA. A second

Figure 1: SPA as a Function of Date of Birth for Women in the UK

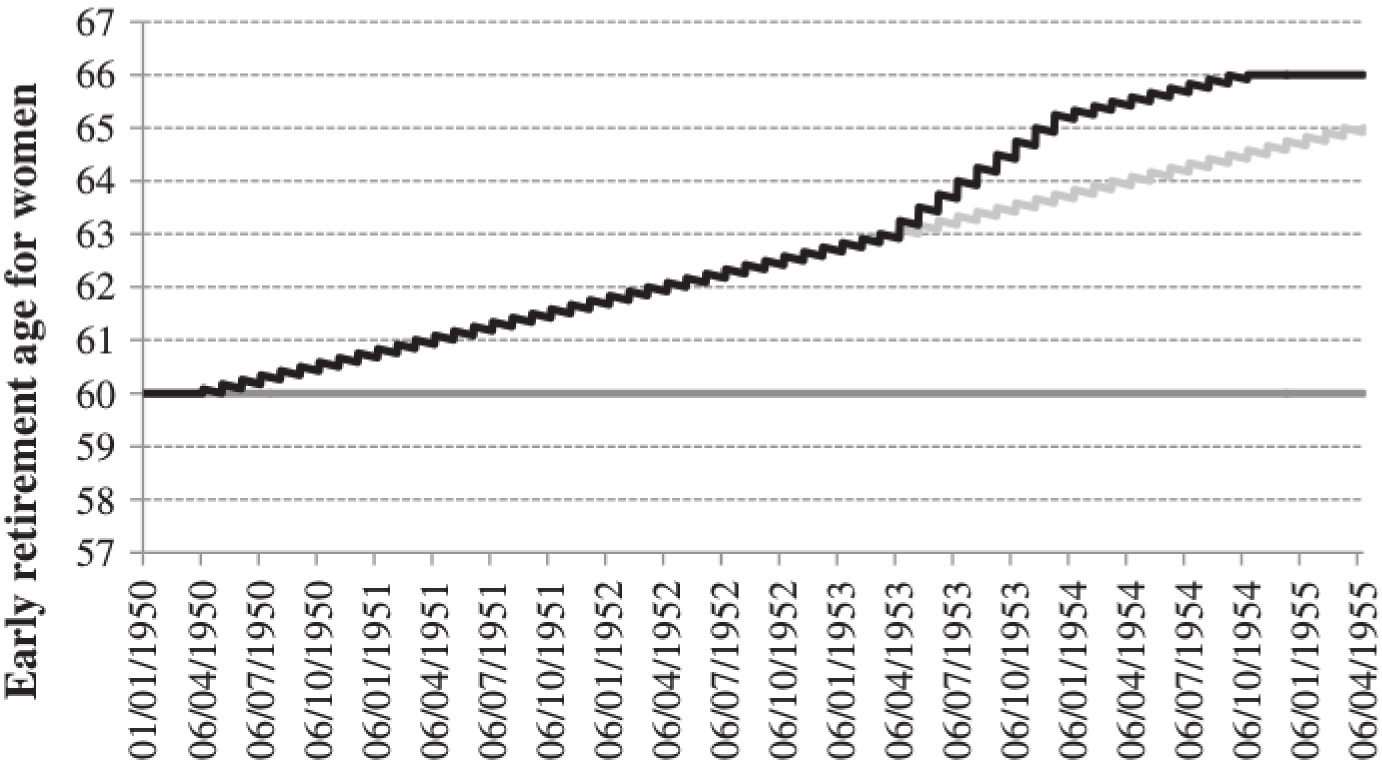

advantage of the UK context is the state pension is not conditional on employment status so that the benefit claiming and participation decision are completely divorced. It is common to carry on working and claim the state pension. A final useful feature of the UK pension system is that, contemporaneously with eligibility to a cash transfer called the state pension, liability for a part of income tax, termed National Insurance, is removed at the SPA. Hence, the tax incentives induced by the SPA are ambiguous and Cribb et al. (2016) find it does not provide major tax incentive to exit from the labour market at this age. Together these features seem to removed removes financial incentives to retire at the SPA for all but the liquidity constrained. A disadvantage with the UK reform is the UK has a single retirement age at which pension benefits are claimed. This makes it difficult to rule out liquidity constraints driving the bunching of retirement at SPA, as the ability to borrow against future pension benefits is severely limited. Cribb et al. (2016) argue against credit constraints being the primary driver on the basis that, whilst homeowners are less likely to be liquidity constrained than renters, the effects of SPA on their labour market participation are indistinguishable. Homeownership, however, is a coarse measure of wealth and equity in one's own home is an illiquid asset. To investigate the role liquidity constraints play, in section 4, I repeat the analysis of Cribb et al. (2016) on the English Longitudinal Study of Ageing (ELSA) dataset which contains detailed asset holding information this allows me to control more precisely for assets.

ELSA is the principal dataset used it this investigation. It is a panel dataset at a biennial frequency containing a representative sample of the English private household population aged 50 and over modelled on the Health and Retirement Study (HRS) in the USA. ELSA contains detailed data on: labour market circumstances, earnings, and the amount and composition of asset holdings. The first wave of ELSA after the start of the initial implementation of the female SPA reform was wave 5 which covered 2010/11 and I use all waves of ELSA from wave 1 (2002/03) thought to wave 7 (2014/15).

4 Reduced Form Evidence

In this section I present reduced form evidence motivating the structural model. In section 4.1 I present reduced form evidence on the first puzzle. Section 4.1 is detailed as I contribute to the literature establishing the existence of a puzzle by focusing the analysis on the ability of liquidity constraints to explain this puzzle in the UK context and find this explanation falls short. In section 4.2, I present some evidence for puzzles two and three. Section 4.2 is shorter as I simply document the nature of the puzzle and their presence in my data.

4.1 First Puzzle

The reduced form analysis presented in this section builds on Cribb et al. (2016) who study the impact of the UK female SPA reform on labour market participation. They regress probability of participation (y_{it}) on: an indicator of being below the SPA; a full set of age, and year of birth dummies; and a vector of controls leading to the following specification:

$$P(y_{it}=1) = \alpha \mathbb{1}[age_{it} \leq SPA_{it}] + \sum_{c=1}^{K} \gamma_c \mathbb{1}[YOB_i = c] + \sum_{a=1}^{A} \delta_a \mathbb{1}[age_{it} = a] + X_{it}\beta + \epsilon_{it} \tag{1}$$

They interpret the parameter α as a difference-in-difference estimator of the effect of the treatment of being below the SPA and find a statistically significant increase in the probability of being employed from the treatment. They estimate this equation as both a linear probability model and as a probit model. As mentioned, they argue that liquidity constraints cannot explain the treatment effect α because a similar size effect is observed for home-owners and renters. Here I repeat their analysis using ELSA which allows me to more carefully control for assets.

In this section, I present results of estimating equation 1 as a random effect linear probability model with errors clustered at the level of the individual. I use a random effect specification because the small sample size means controlling for both autocorrelation and heteroskedascitiy by clustering and arbitrary fixed effect leads to imprecise estimates. The random effects assumption was tested with a Durbin-Wu-Hausman test on the treatment effect and the null, of no difference between the random effect and fixed effects coefficients, was not rejected. For those still uncomfortable with the random effects assumption, I repeat all regression as fixed effects regression with clustered standard errors in appendix A which also contains other robustness checks on the empirical results presented here. Amongst theses robustness checks appendix A repeats the analysis presented here with a random effects probit specification model to address the problems of linear probability models. The complete list of controls used is: a full set of marriage status, years of eduction, education qualifications, and self reported health dummies; an indicator for having a partner, partner's age; partner's age squared; the aggregate unemployment rate during the quarter of interview; dummies for partner eligible for SPA, partner being 65, partner being 60-64, and for being one and two years above and below SPA; and assets of household according to one of the two specification discussed below.

The first column of the top panel of table 1 contains the results of regressing equation 1 as a random effect linear probability model on the entire sample from ELSA. A significant treatment effect is found implying a 10% increase in the probability of being in work from being below the SPA. Then to try and address the question of whether liquidity constraints can explain this treatment effect the next two columns show the results of restricting to the subpopulation of those with above median assets. There are two columns as the subpopulation is selected using two different categories of assets namely the two most liquid categories in Carrol and Samwick (1996). These are, firstly, very liquid assets (VLA) which cover any assets that could be liquidated almost immediately. This includes bank account balances, money market funds, certificates of deposit, government savings bonds, mutual funds, and publicly traded stocks. The secondly categorisation is non-business wealth (NHNBW) which consist of VLA and all other assets and liabilities not related to the primary residence or personally owned businesses; these have in common that the household could liquidate them without losing their home or primary income. Column 2 of the top panel of table 1 shows the results of running the regression on the subpopulations who have more than the median assets, taken over the whole sample, in the interview immediately before their SPA using the NHNBW asset category. Column 3 does the same for the the VLA asset categories. The treatment effect is significant at the 5% level for the whole population and for the population restricted to having above median NHNBW. For the population restricted by VLA the treatment effect is only significant at the 10% level. The size of the treatment effects also does not appear to vary much between the groups; to rigorously test this claim the second panel of table 1 introduces interaction term between being in these subpopulations and being below the SPA. As can be seen, the interaction terms is no significant at any reasonable level supporting the claim treatment effect does not vary between the whole population and the subpopulations having above median assets. These results differ from the predictions of a liquidity constraint explanation for the bunching of labour market exits at SPA.

As it points towards liquidity constraint not being the casual mechanism, I will take the approximate invariance of this treatment effect to restriction based on a median asset split as the stylised fact representing the first puzzle. However, as the median split of asset is arbitrary, in appendix A I consider other data driven

Table 1: Treatment Effect different subpopulations: Random Effects Model

	Whole Population	NHNBW > Median (£24,846)	VLA > Median (£23,800)
Below SPA	0.102	0.099	0.084
s.e.	(0.0312)	(0.0441)	(0.0443)
P>|z|	0.001	0.059	0.025
Below SPA	-	0.136	0.156
s.e.	-	(0.0461)	(0.0457)
Below SPA×(Assets > Median)	-	-0.036	-0.069
s.e.	-	(0.0640)	(0.0639)
P>|z|	-	0.572	0.284
Observations (Individuals)	5,710 (2,882)	2,573 (1,154)	2,563 (1,155)

categorisation of whether individuals are liquidity constrained. Finally, appendix A also contains an analysis in which assets measured continuously are interacted with the indicator of being below the SPA to address concerns about loss of information or bias arising from the dichotomisation of assets. This analysis supports a very limited role for wealth in explaining the treatment effect of being below the SPA on being in work and certainly does not support the idea that alone liquidity constraints can explain all of this effect.

4.2 Second and Third Puzzles

Figure 2 shows the difference between the SPA of 58-year-old women subject to the reform as calculated from their date of birth and their self reported SPA. Although , at age 58 by far the largest group of individuals are those who know their true SPA to within a year, over 40% hold beliefs that are incorrect by a year or more. I take this fact as confirmation that the second puzzle, mistaken beliefs about pension provision, is relevant to understanding the UK female SPA reform. Figure 3 shows the SPA beliefs of two cohorts of women with different SPA. This demonstrates that SPA beliefs appear to be a noisy signal of the true SPA, rather than the error reported in figure 2 being driven by women reporting one or two salient dates such as the pre-reform or male SPA.

With regards to the third puzzle, those reaching their SPA between April 2005 and April 2016 could receive a 1% increase in their subsequent weekly state pension for every five weeks they deferred. Despite this very generous deferral option, Crawford and Tetlow (2010) find that only around 2% of individuals aged between the SPA and 75 receiving a state pension income had deferred receipt. This is difficult to reconcile with SM as incentives clearly indicate that most would be better off deferring.

Figure 2: Mistaken SPA Beliefs of Women Subject to the Reform at Age 58

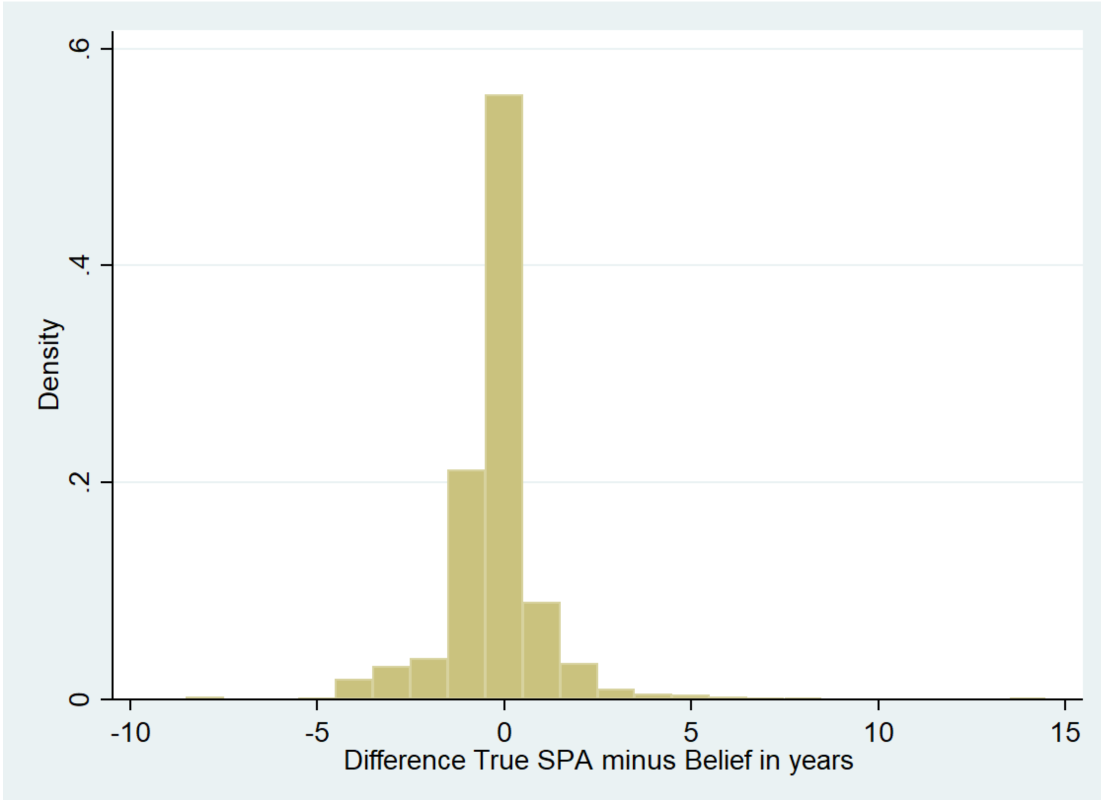

5 Baseline Rational Expectation Model

In this section I use a model to investigate the ability of complete information SM to match the treatment effect of being below the SPA on labour market participation across the wealth distribution observed in section 4.1. This model builds on O'Dea (2018) who develops an life-cycle model that closely matches asset and participation profiles of older individuals in the UK.

The model contains 4 types of individual, $i \in \{1, 2, 3, 4\}$, differentiated by high or low education and being single or in a couple. Agents are rational expected utility maximisers who choose how much to consume c_t, how much to invest in a risk-less asset a_t with return r, and whether to work, dependent on not being involuntarily unemployed. The agent, conditional on not receiving a negative unemployment shock $u_t = 1$, receives a stochastic income offer y_t each period. Unemployment status is considered verifiable so that only if $u_t = 1$ can the agent claim benefit b. If the agent is not unemployed, $u_t = 0$, she receives a stochastic income offer y_t and accepting the offer gives them an income of y_t and leisure time of $l_t = 1 - h_w$. Her partner is modelled deterministically and earns a type and age dependent income $spouseInc_i(t)$ which captures employment income as well as income from other sources such as private and state pensions. The agent receives the state pension, p once she reaches the SPA which is a parameter that is varied to mimic the UK reform. There is a types specific private pension P_i which is accessed from some type specific private pension age PPA_i onward. The value of both the private

Figure 3: SPA Beliefs by Cohort

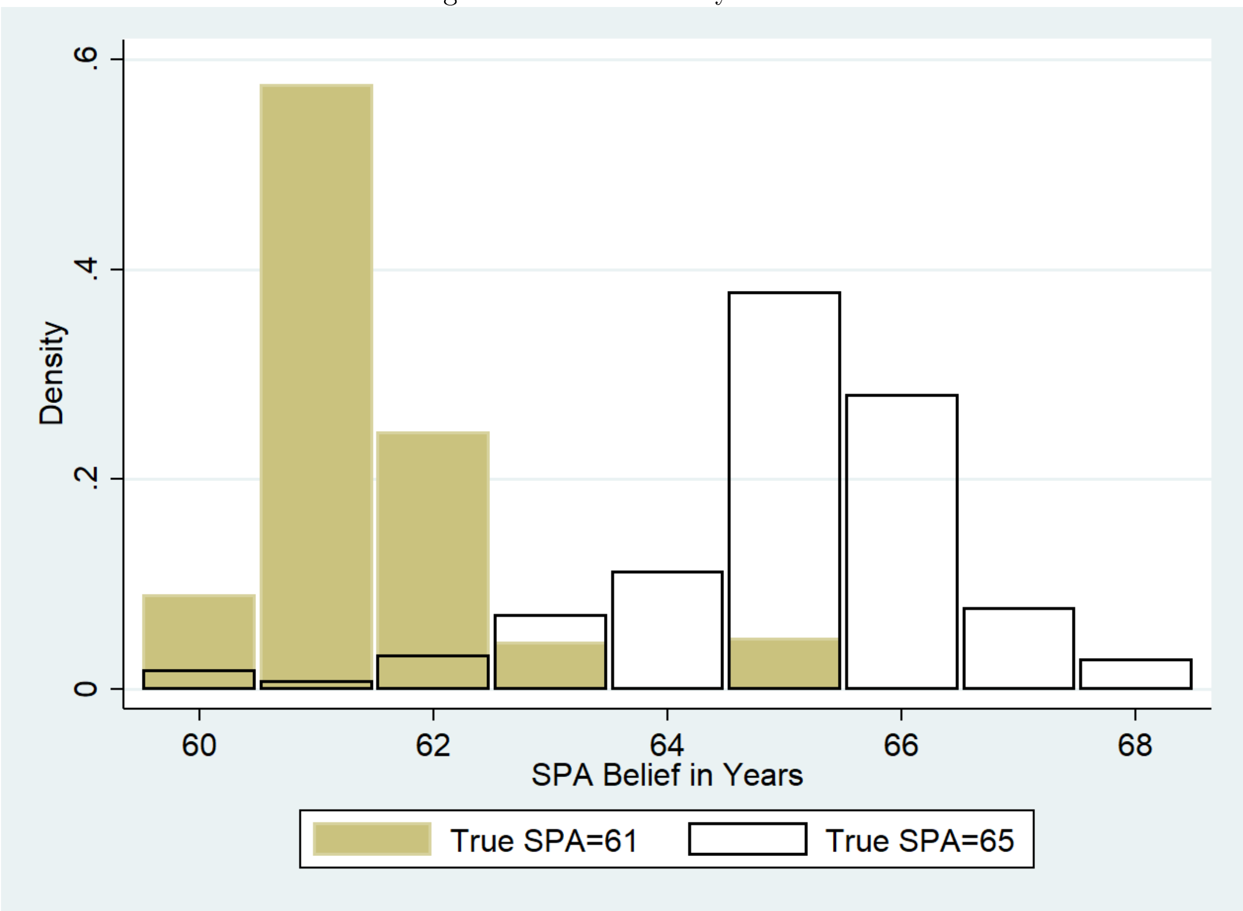

and state pension are functions of average life time earning $AIME_t$. From age 60 the agent is exposed to a survival probability s_t. Finally, agents value bequest through a warm glow bequest function (De Nardi, 2004; French, 2005).

Since the impact of the state pension on retirement decisions is the focus of this paper, it is worth interrupting outlining the rest of the model to explain why I model the state pension as I do. Modelling the state pension as a transfer received upon reaching the SPA abstracts away from the benefit claiming decision: individuals can delay claiming the state pension in exchange for higher benefits in the future. This issue is addressed in an extension in section 7 but it left out here as the objective in this section is to investigate the ability of SM to accommodate the first puzzle: excess sensitivity of participation. However, the third puzzle, that despite more the actuarially fair deferral option only ≈2% take up the option, means that including a deferral option would introduce another puzzle SM have trouble explaining. Hence, to give the baseline model the best chance of explaining the first puzzle deferral is ignored here, but is discussed in an extension in section 7.

So the model can be summarised as the agent solving the following recursive problem to find policy functions for consumption $c_t(X_t)$, leisure $l_t(X_t)$, and assets $a_{t+1}(X_t)$ in terms of states $X_t = (a_t, y_t, AIME_t, u_t, i)$:

$$V_t(X_t) = \max_{c_t, l_t, a_{t+1}} \{u(c_t, l_t, a_t) + \beta s_{t+1} E[V_{t+1}(X_{t+1})|X_t]\} \tag{2}$$

s.t

$$\begin{aligned} c_t + a_{t+1} = {} & (1+r)a_t + y_t \mathbb{1}[l_t = 1 - h_w] + b\mathbb{1}[u_t = unemploy] \\ & + \mathbb{1}[t \geq SPA] S_i(AIME_t) + spouseInc_i(t) + \mathbb{1}[t \geq PPA_i] P_i(AIME_t) \end{aligned} \tag{3}$$

$$a_t \geq 0,\ l_t|(u_t = 0) \in \{1 - h_w, 1\}, \text{ and } l_t|(u_t = 1) \in \{1\}$$

$$\text{where } u(c_t, l_t, a_t) = \frac{(c_t^{\nu} l_t^{1-\nu})^{1-\gamma}}{1-\gamma}|alive + \theta \frac{(a_t + K)^{\nu 1-\gamma}}{1-\gamma}|deceased \tag{4}$$

The log income offer, y_t, is the sum of a deterministic component, quadratic in age and specific to the agent's type, and a stochastic component:

$$log(y_t) = \delta_{i0} + \delta_{i1} t + \delta_{i2} t^2 + \epsilon_t \tag{5}$$

where ϵ_t follows an AR1 with normal error term and an initial distribution $\epsilon_1 \sim N(0, \sigma^2)$.

The private and state pension have the functional form:

$$f_i(AIME_t) = \begin{cases} a_{i,1} AIME_t - a_{i,2} AIME_t^2 & AIME_t < A\bar{IME} \equiv \frac{a_{i,1}}{2a_{i,2}} \\ a_{i,1} A\bar{IME} - a_{i,2} A\bar{IME}^2 & \text{otherwise} \end{cases} \tag{6}$$

Average earning evolving according to:

$$AIME_t = \begin{cases} \frac{AIME_{t-1}(t-1) + \mathbb{1}[l_t = 1 - h_w] y_t}{t} & t < 60 \\ AIME_{t-1} & t \geq 60 \end{cases} \tag{7}$$

The income offer can be conceptualised as being equal to some underlying productivity which the agent maintains during unemployment spells. The unemployment status of the agent u_t evolves according to a conditional Markov process, where the probability of unemployment is dependent on current productivity y_t and the type of the agent. Spousal income is a fourth order polynomial in t,adjusted to account for male mortality.

The model starts with agents aged 52. The reasons to start agents so far into the life-cylce are, firstly, the ELSA dataset only starts interviewing people over 50 and, secondly, the period I am interested in is around retirement and so modelling early life-cylce behaviour would be computationally wasteful. The reason to start at 52 rather than 50 is that this is the youngest age with interviews from a large number of people some with SPA equal to 60 and some with an SPA strictly greater than 60. The agents start life with a draw from the empirical distribution of assets at age 52. To make sure the endogeneity of the SPA to the quantity of assets chosen by age 52 does not bias the model I used a Kolmogorov–Smirnov test to test the null that assets conditional on distinct SPAs are drawn from the same distribution and found that the data do not reject this null. If the age 105 is reached the agent dies with certainty. From age 80 the agent no longer has the choice of working; this is to model some of the limitation imposed by declining health.

The model is estimated using two-stage method of simulated moments. In a first stage the parameters of the earning process, the private pension system, the unemployment transition matrix, and the state pension function are estimated outside the model. The curvature of the warm-glow bequest is taken from De Nardi, French and Jones (2010) and the interest rate from O'Dea (2018). All prices are deflated to a 2013 values using the RPI. In the second step, moments are matched to estimates the preference parameters γ, ν, β, and θ. The moments used are the 48 pre-reform moments of mean labour market participation and asset holdings between 52 and 75.

Figures 4 and 5 show the match of the participation and asset profile which are acceptably close given the coarseness with which the model is currently estimated (for example 30 grid points for assets). Table 2 contains the parameter estimates.

Figure 4: Participation Profile

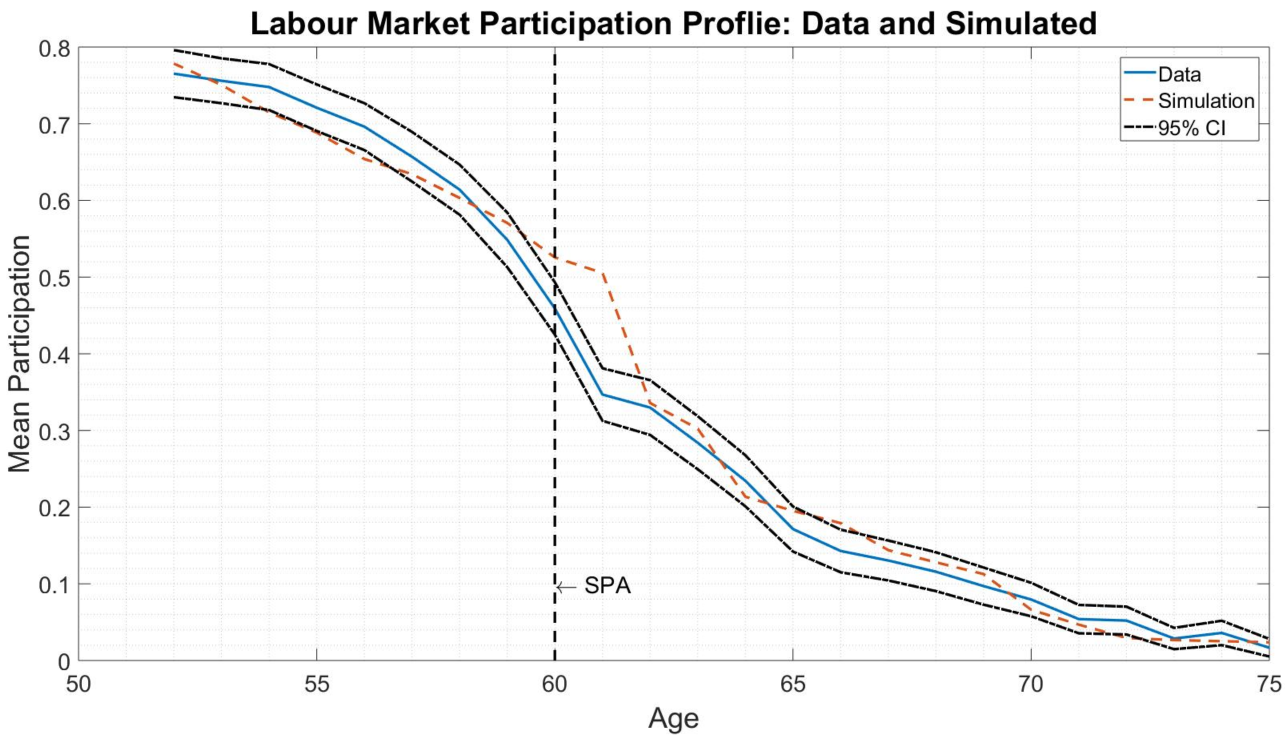

Figure 5: Asset Profile

Table 2: Parameter Estimates

ν: Consumption Weight	0.4319
	(0.0034)
β: Discount Factor	0.9745
	(0.0031)
γ: Relative Risk Aversion	6.1770
	(0.1663)
θ: Warm Glow bequest Weight	28,991
	(5,308)

Table 3: Regression Analysis on Baseline Model Simulations Data

	Baseline Model	Data
	Treatment Effect for being below SPA on participation	
Whole Population	0.0181	0.102
Assets > Median	-0.0016	0.099

These profiles were estimated with $SPA = 60$ and against the moments of the pre-reform data. Once the parameters were estimated, I re-ran the model to generate simulated data with $SPA = 60$, $SPA = 61$, and $SPA = 62$ and re-ran the regression analysis from section 4 on this data. As can be seen in table 3 a significant treatment effect of being below the SPA is observed for the whole sample but unlike the real world data this treatment collapse to negative and insignificant when we restrict to people with sufficient assets. Moreover the model fails to match the magnitude of the treatment effect over the whole population seen in the data. So the baseline model struggles to match both the aggregate response to SPA and the correlation of this response with wealth. In the next section I attempt to ameliorate this issue by introducing rational inattention to the SPA.

6 Rational Inattentive Model

In this section I consider the implications that costly thought has for this retirement decision. First in subsection 6.1, I describe rational inattention and how I incorporate it into the model presented in section 5 and in subsection 6.2, I solve and analyse the model.

6.1 Incorporating Incorrect Beliefs with Rational Inattention

The second retirement puzzle this paper seeks to address is that people are consistently mistaken about their own retirement provision (Gustman and Steinmeier, 1999 & 2001; Rohwedder and Kleinjans, 2006). In particular, women subject to the UK SPA reform are misinformed about their SPA (Crawford and Tetlow, 2010; Amin-Smith and Crawford, 2018); for example, 41% of 58-year-old women with a SPA between 60-64 don't know their own SPA to within a year (Amin-Smith and Crawford, 2018). This is a difficult fact to accommodate in a standard complete information life-cycle model for two reasons: policy parameters are not stochastic and information is acquired without cost. To address these difficulties I introduce inherent policy uncertainty, by making SPA stochastic, and costly attention, by making acquisition of information about the SPA costly using a rational inattention approach.

Making a policy parameter like the SPA a stochastic variable is unusual. Hence, treating the SPA as stochastic requires some justification. Firstly the state pension does changes; a women born in 1954 would have had an SPA of 60 when she entered the workforce in the 1970s but this would have been changed to 65 in 1994 and subsequently to 66 in 2011. Assuming the SPA is a parameter known from entry into the labour force is counter factual.

With regards to interpreting the attention cost, RI takes an agnostic stance which can be both a strength and a weakness of the approach. Hence, multiple interpretations of these cost are possible. However at a minimum the attention cost should be understood as including the full cognitive cost of remembering and assimilating this information into future plans and should not be interpreted as just the hassle cost of finding out your SPA. As illustration, the author has paid the hassle cost of looking up his own SPA online but has not paid the cognitive cost of remembering this information and so would show up in survey data as someone with a mistaken belief and could not currently use his accurate SPA in decision making. This indicates that the full cognitive cost of remembering and assimilating information is both the minimum model and data consistent conceptualisation of these costs.

Although the SPA changes throughout time, reforms have not been frequent enough to allow the estimation of a very flexible stochastic SPA process. Hence I impose a parametric form on the SPA process the capture essential features:

$$SPA_{t+1} = \begin{cases} SPA_t + e_t & SPA_t < 70 \\ SPA_t & SPA_t = 70 \end{cases} \tag{8}$$

$$\text{where } e_t \in \{0,1\} \text{ and } e_t \sim Bern(p)$$

When the SPA is below 70, the process is a random walk with a skewed innovation as $e_t \in \{0,1\}$. So the SPA either increases by one year with probability p in a given year or it stays the same. This process accommodates the idea that in recent history governments have reformed SPA upwards but generally not downward. This process is a parsimonious and analytically tractable model of pension reform. Although it does have some counter factual predictions, in particular, that the SPA can increase by at most one year per year whilst many individuals saw their SPA rise by multiple years in 1995, I believe it captures the essence of pension reform.

I incorporate this stochastic SPA into the model described in section 5, so the model starts with agents aged 52 but they are understood to start working life at age 20 with a SPA of 60. The agents are imperfectly informed of their SPA, which is explained in more detail below, and I make the assumption they are unable to acquire additional information before the start of the model at age 52. So they start with the posterior belief that arises from entering the workforce at age 20 believing with certainty that SPA=60 and then applying Bayesian updating in each period given the process above. Once an agent reaches their SPA, their pension cannot be taken away from them. The probability p is estimated outside the model to match the actual SPAs of women born between 1950 and 1954 assuming these were generated by the process in equation 8. The probability of the SPA increasing taken from the data by this process is 6%.

A stochastic SPA alone would not explain people being mistaken about their SPA. For this I need to incorporate costly thought and this is modelled using the rational inattention approach pioneered by Sims (1998, 2003). This acknowledges the costly nature of thought whilst abstracting away from psychological details. In rational inattention agents suffer a disutility cost for receiving more information as measured by the entropy of the signal they receive about the state variables.

For ease of exposition, I will use a simplified model to explain how rational attention works in my life-cycle setting, although later I introduce rational inattention into the full model described in section 5. In this simplified model the only states are the income offer y_t, assets a_t, and the stochastic SPA SPA_t,

$$\max_{c_t,l_t,a_{t+1}} \sum_{t=0}^{T} \beta^t s_t E[u(c_t,l_t)]$$

$$\text{s.t } a_{t+1} = (1+r)a_t + y\mathbb{1}_t[l_t = 1 - w_h] - c_t$$

In a fully rational model the agent solves for policy functions for consumption $c_t(a_t, y_t, SPA_t)$, leisure $l_t(a_t, y_t, SPA_t)$, and assets $a_{t+1}(a_t, y_t, SPA_t)$ to solve the problem above. In the rationally inattentive model the agent is not able to directly observe the SPA but can only perceive a noisy signal of it $Z_t \sim f_t(z_t|SPA_t, a_t, y_t)$. She can choose the distribution of the signal and make it as precise as she likes but receives a disutility for receiving a more precise signal proportional to the mutual information between signal and SPA $I(Z_t, SPA_t)$. Her policy function can no longer depend on the SPA but only on her beliefs as to what her SPA is. As all agents start with the same prior belief at age 52 their belief in period t is determined by the histories of draws of Z_t they have receive up to that point, z^t. So now the agent chose $Z_t \sim f_t(z_t|SPA_t, a_t, y_t)\}$ and policy functions $c_t(a_t, y_t, Pr(SPA_t|z^t))$, $l_t(a_t, y_t, Pr(SPA_t|z^t))$, and $a_{t+1}(a_t, y_t, Pr(SPA_t|z^t))$ to solve:

$$\max_{c_t,l_t,a_{t+1},f_t} \sum_{t=0}^{T} \beta^t s_t E[u(c_t,l_t) - \lambda I(SPA_t; Z_t)] \tag{9}$$

$$\text{s.t } a_{t+1} = (1+r)a_t + y_t\mathbb{1}[l_t = 1 - w_h] - c_t$$

The penalty for receiving a more precise signal is proportional to the mutual information $I(Z_t, SPA_t)$ which is a concept from information theory. It measures the expected reduction in uncertainty from receiving a signal, where uncertainty is measured by entropy $H(.)$:

$$I(SPA_t; Z_t) = H(SPA_t) - E_Z[H(SPA_t|Z_t)]$$

Entropy is the central concept of information theory and is defined as $H(Z) \equiv -E_Z[\log(f_Z(Z)]$ where $Z \sim f_z$. If the base of the logarithm is taken to be 2, then the entropy is the minimum number of bytes required to communicate the information contained in a random variable; if the logarithm is to a different base then entropy represents the same quantity but measured in a different unit. As such it is an easily understandable measure of uncertainty and the most precise measure of the amount of information received by an agent.

The introduction of rational inattention greatly complicates this model for two reasons. Firstly, it introduces a very high dimensional choice variable in f_t. Since SPA_t has finite support we can restrict, without loss of generality, the support of f_t to be discreet and finite. In this case, f_t is a finite dimensional object but is still of very high dimension, having a dimension of $Dim(SPA_t) - 1$. Secondly, it introduces a large and unobservable state to the agents decision problem in the form of their posterior belief $P(SPA_t|z^t)$.

Most approaches either make a lot of simplifying assumptions, like quadratic utility (Maćkowiak, Matějka and Wiederholt, 2016), or use numerical methods that assume interior solutions (Maćkowiak and Wiederholt, 2015). The method I use builds from theoretical results from Steiner, Stewart, and Matějka (2017) (henceforth SSM) who analytically solve a general class of dynamic discrete choice models without additional simplifying assumptions. However, as SSM's result is for a discreet choice model it requires that I discretise the choice variable. As labour choice is already discreet, this only implies a need to discretise assets. Some other minor extension and adaptations of SSM's result were required and I explain these in Appendix B.

Before outlining the results from SSM which I rely on to solve this model, it is convenient to introduce some notation. Firstly, for brevity I will denote by d_t the agent's decision $d_t = (a_{t+1}, l_t)$ and then I will re-express the agent's utility function as a function of d_t and the states, X_t, by substituting out consumption via the budget constraint to give $u(d_t, X_t)$. Secondly, I let W_t denote the states that the agent freely and costlessly observes, that is all the states except the SPA, $W_t = (a_t, y_t, AIME_t, u_t)$.

The principal result of SSM that I use to solve this model is that a general class of dynamic discreet choice RI problems can be reduced to collections of static RI problems. In the full generality of SSM framework each history of actions represent a different decision node and a different static RI problem. However, there are two ways in which the full history of actions matters in SSM: one through the posterior belief as actions stand in for the history of signals received (see appendix B for details), and two because they allow utility to directly depend upon the full history of actions. Obviously this second channel is not present in the model in this paper, and so the full history of actions only matters above the states of the model due to the posterior belief. Hence if we include the posterior belief as a state variables this removes the dependence on the full history of actions. So solving the model can be reduced to a sequence of static RI problems conditional on an expanded state space including the posterior belief. Then the static RI problems can be solved using the logit result of Matějka and McKay (2015).

More formally at a point in the state space X_t given posterior belief in period $\pi_t = Pr(SPA_t)$ the solution to the model is given by:

$$p_t(d_t|X_t, \pi_t) = \frac{q_t(d_t|W_{t,}, \pi_t)\exp(v_t(d_t, X_t; \pi_t))}{\sum_{d_t' \in D} q_t(d_t'|W_t, \pi_t)\exp(v_t(d_t', X_t; \pi_t))} \tag{10}$$

where

$$v_t(d_t, X_t; \pi_t) = u_t(d_t, X_t) + \beta E[V_{t+1}(X_{t+1}, \pi_{t+1})|d_t, X_t, \pi_t]$$

$$V_t(X_t, \pi_t) = \log(\sum_{d_t \in D} q_t(d_t|W_t)\exp(v_t(d_t, X_t)))$$

and $q_t \in \Delta(d_t)$ solves

$$\arg\max_{q_t \in \Delta(d_t)} E_{\pi_t}[V_t(X_t, \pi_t)] \tag{11}$$

and the posterior next period is calculated using Bayesian updating given the agents actions and the exogenous transition rules. This solution has suppressed the choice of signal function as actions reveal signals received when signals are chosen optimally and so we can directly solve for the conditional choice probabilities $p_t(d_t|X_t, \pi_t)$. However this still leaves the high dimensional state variables π_t which if we wish to approximate each element in the probability vector by n grid point there are $\binom{n+dim(SPA)-1}{n}$ possible values for the posterior π_t. This is too large to be practically manageable in most situations of interest and so I make the simplifying approximation of predicting the posterior base on the other states W_t. Then during simulations, given beliefs in the initial period, I can predict the true beliefs implied by the resulting decision rules and test if this is a reasonable approximation. I find that this approach introduces little approximation error with the average individual error as measured by the total variation metric (i.e. the largest possible mistake possible given the two different beliefs) is of the order 1×10^{-6}.

With this simplification of the posterior belief object this solution of the within period problem can be characterised as:

$$p_t(d_t|X_t) = \frac{q_t(d_t|W_{t,})\exp(v_t(d_t, X_t))}{\sum_{d_t' \in D} q_t(d_t'|W_t)\exp(v_t(d_t', X_t))} \tag{12}$$

where

$$v_t(d_t, X_t) = u_t(d_t, X_t) + \beta E[V_{t+1}(X_{t+1})|d_t, X_t]$$

$$V_t(X_t) = \log(\sum_{d_t \in D} q_t(d_t|W_t)\exp(v_t(d_t, X_t)))$$

$$q_t = \arg\max_{q_t \in \Delta(d_t)} E_{\pi_t(W_t)}[V_t(X_t)] \tag{13}$$

I solve the within period problem by first solving equation 13 using sequential quadratic programming as suggest by Armenter, Müller-Itten and Stangebye (2019) and then substitute the solution for q_t into the analytic formula for the CCP p_t equation 12. I solve the life-cycle problem that nest the within period static RI problem by backwards induction; until age 68 is reached the problem in each period is a fully rational problem and can be solved by maximising utility. At age 68 RI begins to play a part as at ages 70 and above the agent is in receipt of the state pension with certainty and at age 69 if she has not received her pension she knows the SPA must currently be 70. Once age 68 is reached I solve the within period static RI problem as outlined above with the posterior belief about the SPA modelled as a function of the freely observed states W_t

6.2 Rational Inattention Results

I apply the results and methods described in the last section to solve the RI version of the model described in section 5. That is, households who choose the distribution of the signal they will receive $Z_t \sim f_t$ and their actions d_t given costlessly observed states and their history of signal draws to solve:

$$\max_{d_t, f_t} \sum_{start}^{finish} \beta^t s_t E[u(d_t, X_t) - \lambda I(SPA_t, Z_t)] \tag{14}$$

subject to the constraints and exogenous process outlined in equations 3-8.

I have yet to estimate this model but instead present the dynamics of the model for a value of λ, choose to approximately match the subjective belief data, with all other parameter values taken from the estimates of the baseline model. To discipline the cost of attention using the subjective belief data, I treat people's response to question about their SPA as the mode of their posterior distribution in the period which I solve for in the model using the results of SSM. Treating subjective expectation response as representing the mode of a belief distribution is a standard interpretation of such data (Juster, 1966; Manski, 1990; Blass et al. 2010, van der Klaauw, 2012). The value of the cost of attention settled on is $\lambda = 0.0001$. This cost of attention parameter implies to be fully informed of the SPA would incur a cost per period 1.76 times larger than the life-time utility gain to the median household of a 1% increase of consumption in all contingencies. This may seem large but this value was chosen to approximately match the share of people who are mistaken about their own SPA by

Table 4: Treatment Effects Baseline vs RI

	Baseline	Policy Uncert.	Policy Uncert. + Attention Cost	Data
Treatment Effect for being below SPA on participation				
Whole Population	0.0181	0.0065	0.0643	0.102
Assets > Median	-0.0016	0.0008	0.0428	0.099

more than a year: in fact by this metric this value of λ is slightly too small implying 71% of people correctly know their SPA when in the data the truth is slightly less the 60%.

As the RI uncertainty is resolved upon reaching SPA, the SPA functions as a positive exogenous uncertainty shock in the RI model. This might help explain the bunching of exits at SPA because this reduces the precautionary labour supply motive upon reaching SPA. Repeating the regression used in section 6.2 we see that RI does explain part of the excessive bunching puzzle. In table 4, there is a much smaller reduction in the treatment effect in the RI model when we restrict to those with above median assets in the period before SPA with the treatment effect no longer becoming negative and insignificant. Additionally, there is an increase in the size of the treatment effect in the whole population moving it closer to the value seen in the data. Although RI moves the treatment effects closer to their empirical counterparts, for this calibrated version of the model the results are still very different to those seen in the data. One of the main goals of future iterations of this paper is to investigating how much of this discrepancy RI is capable of explaining. As can be seen the introduction of a stochastic SPA alone does not do help resolve this bunching puzzle.

An intuition for these results can be gleaned from considering the trade-off faced by the agents. Figure 8 displays in the top panel a schematic representation of the utility function of an agent at an arbitrary point in the state space if they choose to work and if they do not. Any point on the x-axis is a choice of next periods assets with the value increasing from right to left. This can be thought of roughly as consumption expressed in the standard left-to-right direction although the exact consumption bundle is different for the two functions, working and not working, as the same level of next period assets implies a higher level of consumption if you work. The reason for having the x-axis in these terms is to accord with the bottom pane which shows expected marginal utility next period as a function of the asset choice. This bottom pane is used to find the asset choice level that equalise today's marginal utility with the expected marginal utility next period. As the participation decision is discrete, the agent then determines her working status by selecting the conditionally optimal asset which levels produces the highest utility from these two alternatives. This can be found in the top panel by comparing which intersection point is higher. I have done this for high uncertainty, representing RI in the discussion above, and low uncertainty, representing RE. As the utility function here displays prudence, marginal utilities are convex and so, by Jensen's inequality, increasing uncertainty shifts the marginal value of assets next period upwards. In this diagram we see how this increase in uncertainty can flip the participation of an agent

Figure 6: Schematic Representation of the UMP facing a RI agent

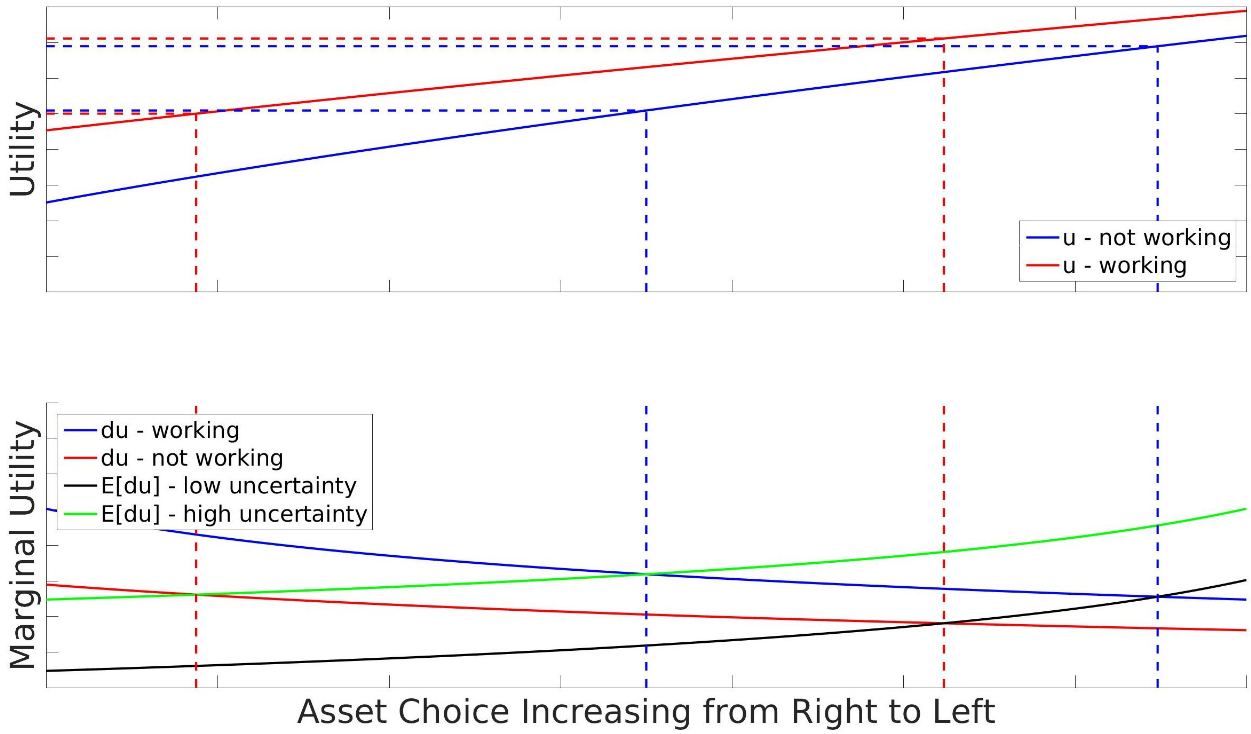

with marginal labour market attachment.

RI is more than the introduction of more uncertainty. It also introduces another channel for the agent to optimise over: the precision of the signal. This channel can be understood in this diagram as the ability to shift the expected marginal value of assets tomorrow down by reducing the uncertainty but only at the cost of also shifting the utility function down. This is the central trade-off introduced by RI but my exploration of this channel is still very preliminary. Investigating this channel, and the informativeness of the signals chosen by agents at different points in the state space, has implications for why people are misinformed about their SPA. One interesting result is that, except at a handful of points, agents choose to receive very little information and these points tend to be found in the upper half of the asset distribution. This prediction agrees with a finding by Rohwedder and Kleinjans (2006) that richer people are more likely informed of their social security provision.

7 Extension

So far this work has ignored the third retirement puzzle identified: that people do not take up more than actuarially fair option to defer pension receipt. However, rational inattention can speak very directly to this puzzle. The calculation implying actuarial favourable deferral ignore the attention cost of an uncertain future pension entitlement, which benefit claiming convert into certain and salient present income.

The version of the model presented in section 6, does not incorporate such a mechanism for two reason.

Firstly, the model does not include a benefit claiming decision. Secondly, the only source of uncertainty subject to an attention cost is the SPA and once this age is reached the attention cost disappears whether the agent claims or not. Including more sources of uncertainty subject to an attention cost would make the model more realistic. If one of these additional source were uncertainty concerning the level of the state pension and a benefit claiming decision were added, then the model would include an additional incentive not to defer resulting from cognitive costs. As long as current income is still treated as salient claiming, removes the attention cost related to the level of state pension by converting uncertainty future pension benefits into certain current income. Hence this provides an incentive not to defer which is ignored in the claims that deferral is more than actuarially fair. The level of uncertainty in pension benefits could be estimated directly from the data outside the model as was done in section 6 for the SPA.

8 Conclusion

This paper offers a solution to three puzzles in the retirement literature by incorporating policy uncertainty and rational inattention into a quantitative retirement saving model. In doing so it makes contributions beyond offering a solution to these three puzzles. Firstly, it adds to the evidence for the first puzzle that retirement choices are more sensitive to legislated pension ages than rational expectations can account for. It does this in the context of the UK female state pension age reform by more carefully controlling for assets and by demonstrating that a state-of-the-art RE model fails to match the observed both the magnitude of treatment effect and correlation with the asset distribution. Secondly, the paper contributes to the rational inattention literature by offering a new solution approach to dynamic RI model and being the first, to the best of my knowledge, not only to incorporate rational inattention into a life-cycle model but also to assess a rationally inattentive model against non-experimental individual choice data. Doing this allows the life-cycle model to accommodate the second puzzle: that people hold mistaken beliefs about their own provisions for retirement. Results from the RI model are preliminary but they offer some insight into the dynamics RI introduces and indicate that RI has the potential to illuminate bunching of labour market exits at SPA as well as people's ignorance of their pension provision. The key mechanism behind this result is that by allowing for uncertainty in the SPA, resolved upon reaching SPA, introduces additional precautionary saving; thus, inducing greater labour market participation pre-SPA. Crucially, introducing reasonable levels of uncertainty about the SPA without RI only negligible increases the sensitivity of labour force participation to the SPA. It is the differential and endogenous amplification of this uncertainty by rational inattention that allows the model to better match the sensitivity of labour market choices to the SPA across the wealth distribution. The approach taken to discipline the cost of attention introduced in the RI model is to use subjective expectations data and, to the best of my knowledge, this paper is the first to use the techniques of subjective expectations literature to discipline attention costs in a rational inattentive model. Finally, the paper propose an extension of the main model with a mechanism to explain the third puzzle for a rational expectations explanation of retirement: that people do not take up more that actuarially advantageous deferral options. The insight offer by this extension is that assertions that deferral is actuarially advantageous omit attention cost which can be avoid by claiming, hence omit an incentive not to defer.

References

[1] Amin-Smith, N. and Crawford, R. 2018. Unpublished Manuscript, 2018

[2] Atalay, K. and Barrett, G.F., 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), pp.71-87.

[3] Bartos, V., Bauer, M., Chytilova, J., & Matejka, F. (2016). Attention discrimination: Theory and field experiments with monitoring information acquisition.The American Economic Review , 106(6), 1437-1475.

[4] Behaghel, B. L., & Blau, D. M. (2012). Framing Social Security Reform : Behavioral Responses to Changes in the Full Retirement Age. American Economic Journal: Economic Policy, 4(4), 41–67.

[5] Bernheim, B.D., 1987. Social Security benefits: An empirical study of expectations and realizations.

[6] Blass, Asher A., Saul Lach, and Charles F. Manski. 2010. Using elicited choice probabilities to estimate random utility models: Preferences for electricity reliability. International Economic Review 51, no. 2:421–40.

[7] Bozio, A., Crawford, R. and Tetlow, G., 2010. The history of state pensions in the UK: 1948 to 2010.

[8] Burtless, G., 1986. Social security, unanticipated benefit increases, and the timing of retirement. Review of Economic Studies 53 (5), 781–805.

[9] Caplin, A. and Dean, M., 2015. Revealed preference, rational inattention, and costly information acquisition. American Economic Review, 105(7), pp.2183-2203.

[10] Carroll, C. D., & Samwick, A. A. (1996). The Nature of Precautionary Wealth. NBER Working Paper.

[11] Crawford, R., & Tetlow, G., 2010. Employment, retirement and pensions. Financial Circumstances, Health and Well-Being of the Older Population in England: ELSA 2008 (Wave 4), 11–75.

[12] Cribb, J., Emmerson, C. and Tetlow, G., 2013. Incentives, shocks or signals: labour supply effects of increasing the female state pension age in the UK (No. W13/03). IFS working papers.

[13] Cribb, J., Emmerson, C. and Tetlow, G., 2016. Signals matter? Large retirement responses to limited financial incentives. Labour Economics, 42, pp.203-212.

[14] De Nardi, M., 2004. Wealth inequality and intergenerational links. The Review of Economic Studies, 71(3), pp.743-768.

[15] French, E., 2005. The effects of health, wealth, and wages on labour supply and retirement behaviour. The Review of Economic Studies, 72(2), pp.395-427.

[16] French, E. and Jones, J.B., 2011. The effects of health insurance and self-insurance on retirement behavior. Econometrica, 79(3), pp.693-732.

[17] Gabaix X. Behavioral Inattention. In: Handbook of Behavioral Economics, Vol 2, edited by D Bernheim, S DellaVigna and D Laibson, Forthcoming. Elsevier ; 2019.

[18] Gustman, A., Steinmeier, T., 1986. A structural retirement model. Econometrica 54 (3), 555– 584.

[19] Gustman, A.L. and Steinmeier, T.L., 2001. Imperfect knowledge, retirement and saving (No. w8406). National Bureau of Economic Research.

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

[21] Jaynes, Edwin T. 1957. "Information Theory and Statistical Mechanics." Physical Review 106 (4): 620–30.

[22] Juster, Thomas. 1966. Consumer buying intentions and purchase prob- ability: An experiment in survey design. Journal of the American Sta- tistical Association 61 (September): 658–96.

[23] Lalive, R., Magesan, A. and Staubli, S.,2017.Raising the Full Retirement Age: Defaults vs Incentives.WP

[24] Maćkowiak, B., & Wiederholt, M. 2009. Optimal sticky prices under rational inattention. American Economic Review, 99(3), 769–803.

[25] Maćkowiak, B. and Wiederholt, M., 2015. Business cycle dynamics under rational inattention. The Review of Economic Studies, 82(4), pp.1502-1532.

[26] Mackowiak, B., Matejka, F., & Wiederholt, M. (2016). The Rational Inattention Filter. CEPR Discussion Papers.

[27] Manoli, D. and Weber, A., 2012. The effects of increasing the early retirement age on social security claims and job exits. Unpublished manuscript, version of May, 11, p.2012.

[28] Manski, Charles F. 1990. The use of intentions data to predict behavior: A best case analysis. Journal of the American Statistical Association 85, no. 412:934– 40.

[29] Matêjka, F., & McKay, A. 2015. Rational inattention to discrete choices: A new foundation for the multinomial logit model. American Economic Review, 105(1), 272–298. https://doi.org/10.1257/aer.20130047

[30] Mastrobuoni, G. (2009), 'Labor Supply Effects of the Recent Social Security Benefit Cuts: EmpiricalEstimates Using Cohort Discontinuities',Journal of Public Economics93(11–12), 1224–1233.

[31] O'Dea, C., 2018. Insurance, Efficiency and Design of Public Pensions. WP

[32] Ravid, D., 2018. Bargaining with rational inattention. WP

[33] Rohwedder, S. and Kleinjans, K.J., 2006. Dynamics of Individual Information about Social Security. Unpublished manuscript, RAND.

[34] Rust, J. (1987). Optimal Replacement of GMC Bus Engines : An Empirical Model of Harold Zurcher. Econometrica, 55(5), 999–1033.

[35] Rust J, Phelan C. How social security and medicare affect retirement behavior in a world of incomplete markets. ECONOMETRICA-EVANSTON ILL-. 1997 Jul 1;65:781-832.

[36] Seibold, A. 2017. Reference Dependence in Retirement Behavior: Evidence from German Pension Discontinuities. WP

[37] Shannon, Claude E. 1959. "Coding Theorems for a Discrete Source With a Fidelity Criterion" In Institute of Radio Engineers, International Convention Record, Vol. 7, 325–50. New York: Institute of Radio Engineers.

[38] Shoven, J.B. and Slavov, S.N., 2014. Does it pay to delay social security?. Journal of Pension Economics & Finance, 13(2), pp.121-144.

[39] Shoven, J.B. and Slavov, S.N., 2012. When does it pay to delay social security? The impact of mortality, interest rates, and program rules (No. w18210). National Bureau of Economic Research.

[40] Sims, C. 1998. Stickiness. Carnegie-Rochester Conference Series on Public Policy, 49, 317–356. https://doi.org/10.1016/S0167-2231(99)00013-5

[41] Sims, C. A. 2003. Implications of rational inattention. Journal of Monetary Economics, 50(3), 665–690. https://doi.org/10.1016/S0304-3932(03)00029-1

[42] Sims, C.A., 2006. Rational inattention: Beyond the linear-quadratic case. American Economic Review, 96(2), pp.158-163.

[43] Staubli, S. and Zweimüller, J., 2013. Does raising the early retirement age increase employment of older workers?. Journal of public economics, 108, pp.17-32.

[44] Steiner, J., Stewart, C. and Matějka, F., 2017. Rational Inattention Dynamics: Inertia and Delay in Decision-Making. Econometrica, 85(2), pp.521-553.

Table 5: Treatment Effect different subpopulations: Random Effects Model

Population	Observations	People	Treatment Effect	P>|z|	95%CI
Whole Population	5,710	2,882	0.101 (0.0312)	0.001	[0.0402,0.1623]
NHNBW > Median	2,573	1,154	0.099 (0.0441)	0.025	[0.0125,0.1855]
VLA > Median	2,563	1,155	0.084 (0.0443)	0.059	[-0.0033,0.1702]
LC Classification 1 NHNBW	4,750	2,384	0.077 (0.0347)	0.027	[0.0085,0.1446]
LC Classification 1 VLA	4,723	2,372	0.077 (0.0348)	0.027	[0.0086,0.1454]
LC Classification 2 NHNBW	2,871	1,260	0.106 (0.0422)	0.012	[0.0234,0.1890]
LC Classification 2 VLA	2,842	1,251	0.101 (0.0423)	0.017	[0.0176,0.1836]
LC Classification 3 NHNBW	2,539	1,139	0.100 (0.0443)	0.025	[0.0125,0.1865]
LC Classification 3 VLA	2,557	1,153	0.084 (0.0444)	0.057	[-0.0025,0.1714]

[45] Van der Klaauw, W., 2012. On the use of expectations data in estimating structural dynamic choice models. Journal of Labor Economics, 30(3), pp.521-554.

[46] Wiederholt, M. (2010). Rational Inattention. The New Palgrave Dictionary of Economics, (August), 1–8. https://doi.org/10.1016/j.red.2010.01.002

A Empirical Results - Robustness

Obviously, the households whose retirement decisions are least likely to be affected by liquidity constraints are those with substantial liquidity assets. I consider two categories of assets which are the two most liquid categories in Carrol and Samwick (1996). Firstly, I look at very liquid assets (VLA) which cover any assets that could be liquidated almost immediately. This includes bank account balances, money market funds, certificates of deposit, government savings bonds, mutual funds, and publicly traded stocks. Secondly, I consider non-housing, non-business wealth (NHNBW) which consist of VLA + all other assets and liabilities not related to the primary residence or personally owned businesses; these have in common that the household could liquidate them without losing their home or primary income.

However, we cannot tell a priori how far up the wealth distribution liquidity constraints affect choices. For this reason, I construct three additional classifications of whether an individual is liquidity constrained each addressing this issue in distinction ways. The first classification considers an individual to be liquidity constrained if they are from a household without assets greater or equivalent to the total of their wage from their interview until their SPA. This classification, however, takes no account of exogenous risk or precautionary saving. The second classification takes account of the precautionary saving motive. It classifies an individual as liquidity constrained if their household has insufficient assets to cover their wage plus the level of asset decumulation at the bottom 25th percentile of the distribution of changes in asset between the periods covering their SPA. The idea behind this classification is that this amount of asset decumulation represent a measure of the exogenous

Table 6: Differences in Treatment Effects: Random Effect Model

	Baseline Treatment Effect	Interaction	P>|z| (Interaction)	95%CI (Interaction)
NHNBW > Median	0.103 (0.0332)	-0.004 (0.0194)	0.852	[-0.0417,0.0344]
VLA > Median	0.109 (0.0332)	-0.013 (0.0194)	0.500	[-0.0512,0.0250]
LC Classification 3 NHNBW	0.102 (0.0331)	-0.001 (0.0194)	0.963	[-0.0389,0.0371]
LC Classification 3 VLA	0.109 (0.0332)	-0.013 (0.0194)	0.502	[-0.0511,0.0250]
Continuous Interaction NHNBW	0.110 (0.0314)	-8.05e-08 (3.18e-08)	0.011	[-1.43e-07,-1.83e-08]
Continuous Interaction VLA	0.109 (0.0314)	-8.18e-08 (3.22e-08)	0.011	[-1.45e-07,-1.86e-08]

risk a household faces at this point in their life-cycle. However, as both of these classifications select on wages when labour force participation is the dependent variable, they produce biased estimates. As people without a job before SPA are less likely to be excluded, the estimates are biased downwards. As I still find a significant treatment effect despite this bias, it is not as large an issue as it would first seem. I still, however, consider a third classification which does not select on wages: having sufficient assets to cover decumulation at the bottom 15th percentile of the distribution of changes in asset between the periods covering their SPA. Rows 4-9 of table 5 show the results of regressing equation 1 excluding liquidity constrained individuals; each row corresponds to a different combination of classifications and asset categories. As can been seen all treatment effects are positive and significant at the 5% level bar one which is significant at 10% level. The magnitudes are also little changed between subpopulation giving little indication that liquidity constraints even form part of the explanation of the observed effect of the SPA on labour force participation.

In rows 3 and 4 of table 6, I test the difference between the whole population and those classed as liquidity constrained according to classification 3 above. There is no statistically significant difference for either grouping. I do not test the difference in treatment effect for the first two classification of liquidity constrained individuals because as mentioned selecting based on these two classification introduces selection bias and so the difference is not interpretable. The final two rows test the significance of a continuous interaction term for the two asset categorisations. For both of them the interaction term is negative and significant indicating that having more assets decrease the impact of being below SPA on the probability of being in work. However, the magnitude of the effect is tiny implying only a 50% reduction of the treatment effect to $\approx 5\%$ for someone at the 99th percentile of the wealth distribution according to NHNBW assets. With this tiny change in the treatment effect due to increased assets it is hard to argue that the treatment effect is completely explained by liquidity constraints.

Table 7 and 8 replicate tables 1 and 6 for the fixed effects specification. As can be seen in table 7, the treatment effects are now significant at the 5% level in four of the populations, at the 10% level in another two and insignificant in two populations. As the magnitudes of the point estimator are little changed, this lack of significance seems to be mostly driven by a lack of power. This is supported by table 8 where the difference between the treatment effects in the two subpopulations remains insignificant. Moreover, the impact of assets

Table 7: Treatment Effect different subpopulations: Fixed Effects Model

Population	Observations	People	Treatment Effect	P>|t|	95%CI
NHNBW > Median	2,573	1,154	0.105 (0.0514)	0.041	[0.0043,0.2060]
VLA > Median	2,563	1,155	0.082 (0.0521)	0.116	[-0.0202,0.1840]
LC Classification 1 NHNBW	4,750	2,384	0.076 (0.0412)	0.066	[-0.0049,0.1566]
LC Classification 1 VLA	4,723	2,372	0.075 (0.0415)	0.067	[-0.0055,0.1573]
LC Classification 2 NHNBW	2,871	1,260	0.108 (0.0501)	0.031	[0.0100,0.2064]
LC Classification 2 VLA	2,842	1,251	0.100 (0.0504)	0.047	[0.0012,0.1988]
LC Classification 3 NHNBW	2,539	1,139	0.107 (0.0516)	0.038	[0.0057,0.2083]
LC Classification 3 VLA	2,557	1,153	0.084 (0.0522)	0.108	[-0.0185,0.1862]

Table 8: Differences in Treatment Effects: Fixed Effect Model

	Baseline Treatment Effect	Interaction	P>|t| (Interaction)	95%CI (Interaction)
NHNBW > Median	0.095 (0.0395)	0.016 (0.0261)	0.534	[-0.0349,0.0673]
VLA > Median	0.103 (0.0396)	0.002 (0.0260)	0.939	[-0.0491,0.0530]
LC Classification 3 NHNBW	0.091 (0.0394)	0.022 (0.0261)	0.399	[-0.0292,0.0731]
LC Classification 3 VLA	0.103 (0.0395)	0.002 (0.0260)	0.938	[-0.0490,0.0531]
Continuous Interaction NHNBW	0.112 (0.0367)	-7.02e-08 (3.14e-08)	0.026	[-1.32e-07,-8.49e-09]
Continuous Interaction VLA	0.112 (0.0367)	-7.14e-08 (3.19e-08)	0.025	[-1.34e-07,-8.77e-09]

on the treatment effect is still tiny.

Table 9 and 10 replicate tables 1 and 6 for the probit specification. For interpretability, table 9 cites the average marginal effect at the median of the subpopulations wealth distribution and quotes the p values of this marginal effect; the coefficients themselves are always significant at the same level as their corresponding marginal effect. As 10 is intended to test the difference of the treatments the p values of the coefficients are quoted. The results are largely comparable to the random effect model both in terms of significance level and in terms of the magnitude of the effect.

Table 9: Treatment Effect different subpopulations: Probit Model

Population	Observations	People	Average Marginal Effect	P>|z|	95%CI
Whole Pollination	5,706	2,881	0.091 (0.0294)	0.002	[0.0342,0.1496]
NHNBW > Median	2,570	1,152	0.083 (0.0396)	0.037	[0.0051,0.1602]
VLA > Median	2,560	1,153	0.070 (0.0395)	0.075	[-0.0071,0.1478]
LC Classification 1 NHNBW	4,745	2,382	0.069 (0.0320)	0.031	[0.0061,0.1315]
LC Classification 1 VLA	4,718	2,370	0.069 (0.0321)	0.031	[0.0064,0.1322]
LC Classification 2 NHNBW	2,868	1,258	0.092 (0.0384)	0.017	[0.0167,0.1672]
LC Classification 2 VLA	2,839	1,249	0.089 (0.0385)	0.021	[0.0133,0.1642]
LC Classification 3 NHNBW	2,536	1,137	0.085 (0.0397)	0.032	[0.0072,0.1627]
LC Classification 3 VLA	2,554	1,151	0.070 (0.0396)	0.073	[-0.0067,0.1484]

Table 10: Differences in Treatment Effects: Probit Model

	Baseline Treatment Effect	Interaction	P>|z| (Interaction)	95%CI (Interaction)
NHNBW > Median	0.859 (0.2703)	-0.159 (0.1598)	0.319	[-0.4724,0.1540]
VLA > Median	0.923 (0.2691)	-0.260 (1599)	0.105	[-0.5730,0.0539]
LC Classification 3 NHNBW	0.847 (0.2695)	-0.142 (0.1591)	0.372	[-0.4538,0.1698]
LC Classification 3 VLA	0.919 (0.2688)	-0.255 (0.1595)	0.111	[-0.5674,0.0582]
Continuous Interaction NHNBW	0.855 (0.2520)	-9.18e-07 (3.48e-07)	0.008	[-1.60e-06,-2.35e-07]
Continuous Interaction VLA	0..854 (0.2519)	-9.28e-07 (3.54e-07)	0.009	[-1.62e-06,-2.34e-07]

B Extending Steiner, Stewart, and Matejka (2017) and Mapping to Model

In this section I adopt much of the notation of SSM and notation is not related to the rest of the paper.

My variant of their setup is the following. There is a payoff relevant exogenously evolving state $\theta_t \in \Theta_t$ according to measure $\pi \in \Delta(\prod_t \theta_t)$ and agents must make a payoff relevant decision from a choice set D_t. Before making a decision, the agent first observes a costless signal $y_t \in Y_t$, $y_t \sim g_t(y_t|\theta^t, y^{t-1})$ and then can choose any costly signal about θ_t on signal space X_t. Agents get gross flow utilities $u(d^t, \theta^t)$ that can depend on the whole history of state and actions but suffer a utility cost for more precise information $\propto I(\theta^t, x_t|z^{t-1})$ where $z^t = (x^t, y^{t+1})$. It is assumed that $y_{t+1} \perp (x^t, d^t)|(\theta^t, y^t)$. The sets Θ_t, D_t,Y_t, and X_t are finite and that $|D_t| \leq |X_t|$. My setup differs from SSM's in that I adapt the timing assumption so that the costless signal is received before the action is taken each period rather than after it. This change in timing only affects the proof of lemma 1 from SSM's paper and I show below that this results still holds using a slightly different strategy to prove it.

The agent chooses information strategy $f_t(x_t|\theta^t, z^{t-1})$ and action strategies $d_t = \sigma_t(z^{t-1}, x_t)$, collectively referred to as their strategy $s_t = (f_t, \sigma_t)$ to solve

$$\max_{f,\sigma} E[\sum_{t=0}^{T} \beta^t (u(\sigma_t(z^{t-1}, x_t), \theta^t) - I(\theta^t, x_t|z^{t-1})] \tag{15}$$

where the expectation is taken with respect to the distribution over sequences (θ_t, z_t) induced by the prior π together with the strategy $s_t = (f_t, \sigma_t)$ and the distributions g_t of costless signals. The function $u(.,.)$ is assumed continuous. For notational convenience, let $\omega^t = (\theta^t, z^{t-1})$ be the current state and the agents current decision node, or information about the state, then:

Proposition 1. *(Lemma 1 in SSM) Any strategy s_t solving the dynamic RI problem generates a choice rule $p_t(d_t|\omega^t)$ solving*

$$\max_{p} E[\sum_{t=0}^{T} \beta^t (u(d^t, \theta^t) - I(\theta^t, d_t|z^{t-1})] \tag{16}$$

where we redefine $z^{t-1} = (d^{t-1}, y^t)$ the expectation is with respect to the distribution over sequences (θ_t, z_t)

induced by p, the prior π, and the distributions g. Conversely, any choice rule p solving 16 induces a strategy solving the dynamic RI problem.

Proof. We precede in steps.

Step 1:First note that for random variable $\zeta_t \in \{x_t, b_t\}$

$$E[\sum_{t=1}^{\infty} \beta^t I(\theta^t, \zeta_t | z^{t-1})] = E[\sum_{t=?}^{\infty} \beta^t (H(\theta^t|\zeta^{t-1}, y^t) - H(\theta^t|\zeta^t, y^t))] \tag{17}$$

But then by the entropic chain rule and that $\theta_t \perp \zeta^{t-1}|\theta^{t-1}$

$$\begin{aligned} H(\theta^t|\zeta^{t-1}, y^t) &= H(\theta^{t-1}|\zeta^{t-1}, y^t) + H(\theta_t|\theta^{t-1}, \zeta^{t-1}, y^t) \\ &= H(\theta^{t-1}|\zeta^{t-1}, y^t) + H(\theta_t|\theta^{t-1}, y^t) \end{aligned}$$

Also $y_{t+1} \perp (x^t, b^t)|(\theta^t, y^t) \Rightarrow H(y_{t+1}|\theta^t, x^t, y^t) = H(y_{t+1}|\theta^t, y^t) = H(y_{t+1}|\theta^t, b^t, y^t)$, so by symmetry of mutual information

$$\begin{aligned} H(\theta^t|\zeta^t, y^t) - H(\theta^t|\zeta^t, y^{t+1}) &= I(\theta^t; y_{t+1}|\zeta^t, y^t) = I(y_{t+1}; \theta^t|\zeta^t, y^t) \\ = H(y_{t+1}|\zeta^t, y^t) - H(y_{t+1}|\theta^t, \zeta^t, y^t) &= H(y_{t+1}|\zeta^t, y^t) - H(y_{t+1}|\theta^t, y^t) \end{aligned}$$

So 17 becomes

$$\begin{aligned} &E[\sum_{t=1}^{\infty} \beta^t (H(\theta^{t-1}|\zeta^{t-1}, y^t) - H(\theta^t|\zeta^t, y^{t+1}) - H(y_t|\zeta^t, y^t) + H(y_t|\theta^t, y^{t-1}) + H(\theta_t|\theta^{t-1}, y^t))] \\ &= E[\sum_{t=1}^{\infty} (\beta^{t+1} - \beta^t) H(\theta^t|\zeta^t, y^{t+1}) - \beta^t H(y_t|\zeta^t, y^t) + \beta^t (H(y_t|\theta^t, y^{t-1}) + H(\theta_t|\theta^{t-1}, y^t))] \end{aligned}$$

Step 2: Given strategy s and the choice rule generated by it p by construction they generate the same gross utilities. Hence by step 1, 16-15 is:

$$E[\sum_{t=1}^{\infty} (\beta^t - \beta^{t+1})(H(\theta^t|b^t, y^{t+1}) - H(\theta^t|x^t, y^{t+1})) + \beta^t (H(y_{t+1}|b^t, y^t) - H(y_{t+1}|x^t, y^t))]$$

But then $|B| \le |X| < \infty \Rightarrow b^t$ is measurable wrt x^t and hence $E[H(\theta^t|b^t, y^{t+1})] \ge E[H(\theta^t|x^t, y^{t+1})]$ and $E[H(y_{t+1}|b^t, y^t)] \ge E[H(y_{t+1}|x^t, y^t)]$ and therefore 16$\ge$15 .

Step 3: As $B \subset X$ if p is a probability choice rule then $f_t(x_t|w^t) = p_t(b_t|\omega^t)$ and $x_t = \sigma_t(z^{t-1}, x_t)$ is a viable solution to 15. For this strategy generated by this mapping, the probability choice rule makes equation16=equation15

Step 4: If s solves 15 the corresponding PCR p must solve 16, as by step 2 the value from p in 16$\geq s$ in 15, so if p doesn't solve 16 $\exists$ PCR producing greater net lifetime utiltiy than s in 15. But by step 3 this produces a viable solution to 15 with greater net life-time utility contradicting s being a solution to 15.

Step 5: If p solve 16 then by step 3 it produces a viable solution to 15 but then 16$\geq$15 so this strategy must be the optimal solution to 15 $\square$

The remainder of the proof follow as stated in SSM for the case where the choice variables are discrete as is the case in this paper.

It is worth saying a few words about how the model in this paper maps to the class of models in this appendix based on SSM as the correspondence is not obvious. The clearest difference between the SSM setup and the model presented in section 6 is that SSM only allow for exogenous states whilst I have an endogenous state in the form of assets a_t. However, since utility can depend upon the entire history of choices and states there is a simple mapping from the endogenous states without history dependent preferences to the world of exogenous states with history dependent preferences. The state in the sense of SSM now only contains the exogenous states $\Theta_t = Supp(SPA) \times Supp(Y_t) \times Supp(AIME_t)$, $(SPA_t, y_t, u_t) = \theta_t \in \Theta_t$ but since $a_t \in d_{t-1}$ and $AIME_t = g(d^{t-1}, \theta^{t-1})$ for the function g that follows from the definition of $AIME_t$ given in section 5. Hence, we can re-express the the utility given in terms of section 6 states X_t and the current decision $u(d_t, X_t)$ in terms of the history of exogenous state θ^t and the history of decisions $u(d^t, \theta^t)$. And since the SSM agent condition their action on everything useful from $z^{t-1} = (d^{t-1}, y^t)$, they can condition on all states.