Policy Uncertainty about State Pension Reform[*]

PRELIMINARY

Emanuele Ciani, Adeline Delavande, Ben Etheridge, Marco Francesconi[†]

March 7, 2016

Abstract

This paper estimates the effect of uncertainty about pension reform in Europe both before and after the financial crisis. Using data on subjective beliefs from SHARE, we provide evidence of considerable uncertainty about the future path of pension provision. This uncertainty exists along multiple dimensions — such as the generosity of benefits and changes to pensionable age — and beliefs about future reforms are highly heterogeneous. To assess the effects of this uncertainty we estimate a rich life-cycle model of saving, labour supply and pension wealth accumulation. A new and important contribution is that the model allows for the evolution of subjective beliefs. We find that, across the whole population, the different types of uncertainty have moderate effects, even during (or after) the financial crisis, and even allowing for high levels of risk aversion. Some subgroups of the population, however, face higher welfare costs of uncertainty, particularly those approaching retirement in unemployment. Conversely, for those in work, pension uncertainty is overshadowed by (current) income uncertainty and is reduced by flexibility about retirement choice. We show that, within different types of reform, uncertainty about benefit generosity is more costly than uncertainty about pensionable age. Finally, we assess the consequences of belief heterogeneity.

JEL Classification: D84; D91; J26.

Keywords: Consumption; Life-cycle models; Subjective expectations; Public pension reform.

[*]Thanks to Tom Crossley, Cormac O'Dea, Jayant Ganguli, seminar participants at Essex, ESPE 2015. The views expressed in this paper are those of the authors and do not necessarily reflect those of the institutions to which they are affiliated. Errors are our own etc.

[†]Ciani: Banca d'Italia; Delavande: ISER, University of Essex; Etheridge: Essex, E-mail: bsethe@essex.ac.uk, www: https://sites.google.com/site/benetheridge/; Francesconi: University of Essex.

1 Introduction

Across much of the world populations continue to age rapidly. In most developed countries, governments have inherited pension systems that were designed when populations were on average younger, and dependency ratios were lower. To maintain fiscal stability, these pension systems have been, and continue to be, reformed. Such pension reform remains one of the biggest policy challenges of the current era. However, given that the extent of reform required is unknown and given the vagaries inherent in the political process, it seems that reform is surrounded with great uncertainties.[1] This paper contributes to the literature assessing the nature of these uncertainties, and their effects on welfare.

This uncertainty can take many forms. It depends not just on uncertainty about the government's long-run budgetary position. Uncertainty also arises because reform might happen along many different paths (Kitao, 2014; Haan and Prowse, 2014). Uncertainty also affects individuals differently depending on their situation in life. For this reason we build a detailed structural model that includes uncertainties across multiple dimensions and allows for life-cycle saving and labour supply, in the spirit of, and building on Haan and Prowse (2014) and French and Jones (2011). Our paper is the first to address these uncertainties using a model with such detail.

We measure uncertainty using data on reported subjective beliefs. Our data come from the Survey of Health, Ageing and Retirement in Europe (SHARE). The dataset contains extensive information on labour supply and saving for a large sample of individuals aged over 50 across Europe. Importantly within the present context, the dataset also elicits subjective beliefs of the chance of reform to state pensions along two margins: the chance of a drop in pension generosity and the chance of an extension in retirement age. Our dataset also elicits expectations about policy outcomes. By combining these two sets of responses, we obtain measures of the perceived size of reforms.

Our data cover 20 countries and come in 4 waves, covering 2004 to 2013. Not only does our dataset therefore cover a variety of countries with different initial policies and needs for reform, it also straddles the pre- and post- crisis period, when state fiscal positions deteriorated. We first document that individual expectations correlate extremely well with existing pension entitlements and have responded to past pension reforms. We also document strikingly that perceived chances of reform increased after the crisis. We further provide evidence that, in spite of these responses, considerable uncertainty exists about the future path of pension provision. This uncertainty exists along both measured dimensions - the generosity of benefits and changes to state retirement age - and beliefs about future reforms are highly heterogeneous.

To assess the effects of this multi-dimensional uncertainty we formulate and estimate a rich model of the latter life-cycle, from age 50 onwards. Much of our model is in the spirit of French and Jones (2011) and Haan and Prowse (2014). The model is further formulated to match the institutional details of pension systems across Europe. An important contribution, however, is that we model the evolution of subjective beliefs within a subjective expected utility framework. Specifically, our model features heterogeneous

[1] See, for example, Giavazzi and McMahon (2012) in the case of German pension reform, and the literature discussed below.

beliefs about future pension policy; these beliefs are updated in each period based on a stylized model of learning. More conventionally, our model also features private saving, pension wealth accumulation, and several channels of labour supply adjustment including retirement. In addition, we model other risks faced by the individual, specifically those to wages, employment and mortality. Our model, therefore, focusses on capturing factors that might influence, attenuate or amplify uncertainties about the state pension system.

We use this model to address two main sets of questions about individual outcomes. First, and similarly to the literature on earnings risks (such as Low, Meghir, and Pistaferri, 2010), we assess the effect of risk on welfare and on prior behaviour. Second, we assess the effect of the observed belief heterogeneity. In our context, this belief heterogeneity implies that at least some individuals' beliefs must be mistaken. This is because the true probability of reform is naturally modelled as common; reform affects everybody equally within countries. We not only examine the welfare costs of these mistakes, we also assess to what extent heterogeneity in beliefs drives heterogeneity in outcomes.

Before answering these questions we fit the model to the data in two steps. We first estimate the model of belief formation. To proceed, we assume that the true chance of reform is captured by the average belief. In our econometric specification we therefore allow for beliefs to converge on average, even though they are determined idiosyncratically. The estimation itself is challenging because the data are measured with significant error.[2] We therefore adapt the approaches in Bernheim (1987) and Benitez-Silva and Dwyer (2005). We find persistence in individual beliefs, with a one-year autoregressive coefficient of 0.7. Therefore, for example, individuals who are pessimists about pension reform remain pessimists over several waves, but do eventually become more neutral. The rate of convergence is informative about the speed at which individuals learn about the chance of reform.

We then fit the retirement model to data on labour supply and wealth using simulated method of moments as in, for example, French and Jones (2011). We fit preferences for consumption and non-linear costs of work. Similarly to French and Jones, we pay attention to an initial conditions problem that arises because our data and model start at age 50 and because many of our state variables are, at best, partially observed. Just as in French and Jones, we estimate unobserved heterogeneity in consumption preferences by including an equation for latent types. In addition, over and above French and Jones, we deal with the problem that initial beliefs about retirement are measured with substantial error. Similarly, in the estimation we allow for measurement error in initial wages.

Before using the model, we validate it by displaying the fit against moments that have not been used in the estimation. Specifically, we examine how well the model can explain beliefs about future behaviours. We focus on the belief about the chance of being retired at age 63. This belief clearly depends not only on beliefs about the state system, but also about the individual's endogenous labour supply response. This belief is elicited in the survey, but, in contrast to beliefs about the state system, is not explicitly

[2] We use the term 'measurement error' loosely. Our approach is to assume that that responses capture some underlying belief, which is persistent and rationally considered, and some random disturbance which captures the 'mood' of the interviewee, and potentially a rounding to focal interview points. In keeping with the literature we use the term measurement error to capture this second component. It could also be termed 'response error'.

included in our model as a state variable. To perform our validation we therefore compute this belief for each simulated individual, based on expectations about the distribution of future shocks. We then regress this computed belief on the beliefs about the state pension system and other relevant covariates, both in the simulated population and for the observed sample in SHARE. The regression results line up closely, indicating that the model realistically captures the relationship between beliefs and behaviour.

Our results are as follows. We find that heterogeneity in beliefs drives a non-negligible fraction of heterogeneity in outcomes. For example, conditional on other factors, pessimists about reform save a considerable amount more than optimists. To quantify effects, we focus on heterogeneity for two specific outcomes: consumption growth and wealth levels. We also contrast the effect of belief heterogeneity with the effect of heterogeneity coming from preferences, specifically discount rates. Beliefs and preferences are both potentially important, and traditionally unobserved, factors driving economic outcomes. Of course, both consumption growth and wealth levels are driven overwhelmingly by labour-market shocks. In terms of our 'unobserved' factors, however, we find that belief heterogeneity is important for consumption growth as individuals approach retirement. However, it is generally less important than heterogeneity in preferences. This result is even stronger when looking at wealth levels. Here the effect of preference heterogeneity accumulates over the whole lifetime, and particularly in early working life, when beliefs about the pension system are presumably less important.

Next we examine the effects of pension policy risks and focus more explicitly on welfare costs. Similarly to Caliendo, Gorry, and Slavov (2015), we find that the welfare costs are generally small for the population we study. There are several reasons for this result. First, state pension wealth is only one component of life-time wealth: On average, our population has high private saving, and, even for those close to retirement, non-negligeable expected future earnings. Second, even an aggressive cut to state pension income is comparatively small compared to the loss in earnings from bad career outcomes: pension risks therefore seem less important than earnings risks, studied, for example by Low et al. (2010). Third, for those in work, pension risk sits behind more immediate income risks, such as the chance of being made unemployed. For those in work, pension risk is also reduced by flexibility about retirement choice. Overall, our results on the welfare costs of risk are similar in magnitude to the welfare costs of mistaken beliefs.[3]

Having said this, for some subgroups the welfare costs of risk are higher. In particular, welfare costs are highest for those approaching retirement in unemployment, and particularly for those with low private saving. This is because this subgroup no longer faces such down-side labour market risk, and its pension wealth dominates other forms of expected lifetime wealth. Overall, however, we find that pension risk has moderate effects, even during (or after) the financial crisis, and even allowing for high levels of risk aversion.

As discussed, an important component of our approach is that we allow for reform to happen along a variety of paths. Overall, we find that uncertainty surrounding pension generosity is more important

[3]This result comes from the subjective expected utility framework chosen; in this case preferences for allocation across time match preferences for the allocation across states. With Epstein-Zin preferences, in contrast, the welfare costs of risks would be higher than the welfare costs of mistaken beliefs.

than uncertainty surrounding changes to standard retirement age. Of these two paths, pension generosity affects individuals more obviously: a reduction in pensions, holding tax payments fixed, lowers life-time resources.[4] The effects of changes to standard retirement age are more subtle. In much of the world now, the connection between standard retirement age and pension entitlement is weakening. Several papers, however, find that extending standard retirement age increases labour supply substantially.[5] In our model, age at retirement does indeed affect pension entitlements, as in most systems across Europe. Nevertheless the effects for this path of reform are quantitatively smaller.

As discussed, our paper relates to a several literatures. First, several papers have looked at policy uncertainty about state pensions using a more reduced-form approach. For example Van Der Wiel (2008) assesses how such uncertainty, as measured by subjective beliefs, affects savings, also data from the Netherlands. Luttmer and Samwick (2012) take a slightly different approach by eliciting subjective certainty equivalents of possible social security arrangements in the US; they therefore compute welfare costs of uncertainty directly. Giavazzi and McMahon (2012) assess precautionary saving using a quasi-experimental design provided by a sequence of legislative announcements in Germany. We contribute to this literature by providing a structural model. In this respect, our paper is perhaps most similar to Caliendo et al. (2015), although they do not model or fit a process for belief formation. In addition our model is more detailed than theirs in several other ways and is matched to the data on saving and labour supply more closely.

Second, as discussed, our paper contributes to the literature on risks over the life-cycle: Meghir and Pistaferri (2011), for example, survey the literature that quantifies earnings risks and assesses how they affect behaviour and welfare. More widely, our work relates to a macroeconomic literature. For example, Fernández-Villaverde, Guerrón-Quintana, Kuester, and Rubio-Ramírez (2011) assess the effect of volatility about future fiscal policies using a New Keynesian model. Like us, they address how uncertainty is exacerbated by the multiplicity of paths of reform. They use data from, amongst other sources, Baker, Bloom, and Davis (2015), who provide a new measure of uncertainty using data from newspaper reports.

This paper proceeds as follows. Section 2 describes the model in detail. Section 3 discusses how we use the data from SHARE. We describe how we estimate the dynamic model of policy beliefs in section 4, while section 5 presents results from the second-stage estimation of the retirement model. Section 6 provides a model validation, while full results are discussed in section 7. Finally section 8 concludes. An extensive appendix provides further details.

[4] Of course, the incidence of any changes to the pension system can be complex. In this paper we assume that current tax rates are held constant and that any cut to pension generosity is effectively a transfer away from the cohorts being studied towards younger and future cohorts (who presumably eventually then need to be taxed less).

[5] See Mastrobuoni (2009) and Cribb, Emmerson, and Tetlow (2013).

2 Model of Saving and Retirement

2.1 Model overview

We build a model of the latter life-cycle, from age 50 onwards, based on a subjective expected utility framework. Our model features heterogeneous beliefs about future pension policy together with learning. More conventionally, our model also features private saving, pension wealth accumulation, and several channels of labour supply adjustment including retirement. In addition, we model other risks faced by the individual, specifically those to wages, employment and mortality. Our focus, therefore, is on capturing factors that might influence, attenuate or amplify uncertainties about the social security system.[6]

The model aims to capture the institutional details that are common across most European social security schemes. Accordingly, therefore, we model a entitlement-based defined-benefit pension system. We also include some of the publicly-provided insurance schemes designed to mollify the other risks, specifically a broad entitlement-based unemployment insurance system. Nevertheless, the individual's main protection against uncertainty is self-insurance in the form of precautionary savings.

As a useful tool for reference, table 1 gives definitions of the parameters and variables used in the model.

2.2 Policy beliefs

A central feature of our model is the individual's beliefs about policy and policy reform. The data show that, even in late career, individuals perceive substantial uncertainty about whether and to what extent state pensions will be reformed. In addition, as Kitao (2014) emphasizes, reform can occur along many dimensions. The main path of reform could be, for example, to, the contribution rate, the income tax rate (if social security is partially funded out of the general taxation), the age of entitlement, the level of generosity of benefits, or even how benefits uprate with inflation.

With this in mind, we construct our model of beliefs as follows. At the heart of the individual's planning is a set of potential policies. The set of policies is:

$$\mathcal{P} = \{P_{jk}\} \tag{1}$$

where P_{jk} is a particular policy for social security. In our model we index the policy by two subscripts to capture what we see as the two principal paths of reform. j indexes the standard retirement age (SRA), at which the head can draw his pension. k indexes the pensions replacement rate, expressed as a percentage of some normalized pension entitlement, based on accumulated rights.

[6]We use the term 'social security' and 'pension' somewhat interchangeably. In fact the model does not include private pensions of any form.

Table 1: Variable Definitions

	Policy-related Parameters		**Income-related Parameters**
P_{jk}	Government policy	ρ	Autoregression coefficient on wages
$\mathcal{P}$	Set of government policies	ζ_t	Persistent shock to earnings
$\mathcal{B}_{it}$	Beliefs over government policies	σ_x^2	Variance of innovation of type x
$p_{B,it}$	Subjective prob. of future benefit cut reform	W_t^P	Permanent component on wages
$p_{R,it}$	Subj. prob. of future retirement age reform	W_t	Wage level
$\pi_{xy}^B()$	subj. transition probability of $p_{B,it}$	ϵ_t	Transitory shock to earnings
$\tilde{p}_{B,it}$	Subj. prob. of any benefit age reform	Y_t	Post-tax income
$\tilde{\pi}_{xy}^B()$	subj. transition prob. of $\tilde{p}_{B,it}$	YL_t	Labour earnings
$\mathbb{P}_B$	Binary indicator on passing of a reform	UI_t	UI benefits
$p_{B,it}^*$	Subj. prob. of reform in next year	ss_t	Pension income
λ	Subjective correlation of reforms	$\tau\left(\right)$	Tax function
θ	Autoregression coefficient on beliefs	tr_t	Social assistance level
φ	Shock to beliefs	sp_t	Spousal income
	Preference Parameters		**Labour Market Variables**
β	Time preference	δ	Job seperation probability
ξ	Strength of bequest motive	Λ	re-entry probability
$u\left(\right)$	Intratemporal felicity function	E_{it}	Beginning-of-period employment status
$b\left(\right)$	Bequest function	Θ	Fluctuations on unemployment income
γ	Coefficient of rel. risk aversion	D_{it}	Duration of unemployment spell
η_{F_t}	Consumption cost of work for empl state F_t		
$\eta_{F,0}$	Consumption cost of employment age 50		**Pension Variables**
η_t	Increase in consumption costs of work.	$\mathbb{I}_{it}$	Pension entitlement
η_{REt}	Frictional costs of changing labour supply	adj	Age adjustment for pension rights
κ	Total costs of work	$rr_{it}\left(\right)$	Replacement rate
		α_t	Pension accumulation factor
	Decision Variables/Values		
C_{it}	Consumption		**Miscellaneous**
F_{it}	Labour-force choice	R	Real Interest rate
M_{it}	Beginning-of-period (net) assets	s_t	Survival probability
$V^F\left(\right)$	Value functions conditional on work choice		
$\mathcal{V}^E\left(\right)$	Value functions for each initial l.f. status		

In our model, the relevant policy is that which applies from the start of retirement onwards. This final policy is uncertain. In our application the range of policies is quite simple. The individual might face the status quo or a (less generous) alternative. At age 50, the initial age in the model, the status quo is in place. From this age onwards, the policy might be reformed, or stay in place. We allow for at most one reform along each margin, and any reform is final: it cannot be revoked. We let $\mathbb{P}_B$ and $\mathbb{P}_R$ be binary indicators for whether or not a reform has taken place. Therefore, reforms could happen along either of the two margins, to neither, or to both.

When uncertainty about future policies has not yet been resolved, individuals form beliefs. Formally, at time t, individual i forms beliefs about future reforms, given the set of possible policies, and according to subjective expected utility theory. Their beliefs are captured as:

$$\mathcal{B}_{it} : \mathcal{P} \Rightarrow [0, 1] \tag{2}$$

In line with the data that are available, and for computational tractability, we assume that beliefs are captured by the marginal distributions. In other words, $\mathcal{B}_{it}$ is determined by $p_{B,it}$, which reflects the perceived chance of a cut to benefits, and $p_{R,it}$ which reflects the chance of an extension to the retirement age. We assume that correlation of outcomes, given by λ, is common across individuals and time.

A central feature of our model is that we allow for beliefs to be heterogeneous. Belief heterogeneity has received attention recently, most noticeably in research on asset prices (see Cao, 2013 and Burnside, Eichenbaum, and Rebelo, 2011). To the best of our knowledge, we are the first to consider the implications in a quantitative life-cycle model of saving and labour supply. One aim of our paper is to characterize what happens when beliefs about reform are mistaken. We therefore assume there is an 'objective' probability of reform. An objective probability is natural in this setting because individuals are reporting beliefs not about themselves, but about an external agent, the government. Although this objective probability is not observable, it captures the idea that, when beliefs are heterogeneous, at least some individuals must be wrong. In the application, we assume the objective probability is the average belief. This assumption has the virtue of implying that the mean squared error of beliefs is minimized. Of course, when quantifying this heterogeneity, we also need to consider that beliefs are likely measured with substantial error, or misreported, in our data. We take this into account in the econometric specification discussed further in sections 4 and 5.

2.2.1 The evolution of beliefs

Both the individual and the econometrician know that beliefs will evolve in future. The manner in which beliefs evolve is restricted both by theory and by modeling choice. When the individual is solving his dynamic choice problem, he assumes that his future beliefs will evolve according to transition matrices

Π_B and Π_R, with entries π^B_{xy} and π^R_{xy}.[7] Formally, for example in the case of a reform to benefit levels:

$$Pr\left(p_{B,it+1} = \mathscr{P}_x | p_{B,it} = \mathscr{P}_y \& \mathbb{P}_{B,t+1} = 0\right) = \pi_{xy}$$

where x and y index possible values of beliefs at times $t+1$ and t. Beliefs about reforms to retirement age are specified in exactly the same way.

Future beliefs must be consistent with the present. Therefore, when a reform has not yet happened:

$$\begin{aligned} p_{B,it} &= Pr\left(\mathbb{P}_{B,t+1} = 0\right) \times \mathbb{E}_t\left(p_{B,it+1} | \mathbb{P}_{B,t+1} = 0\right) + Pr\left(\mathbb{P}_{B,t+1} = 1\right) \\ \implies p_{B,it} &= \left(1 - p^*_{B,it}\right) \sum_x \pi^B_{xy} \mathscr{P}_x + p^*_{B,it} \end{aligned} \tag{3}$$

where $\mathbb{E}_t$ is the expectations operator at time t and $p^*_{B,it} = Pr\left(\mathbb{P}_{B,t+1} = 1\right)$, the chance of a reform in the next period. Clearly $p_{B,it}$ and $p^*_{B,it}$ are related. We discuss further their connection, and how they are calculated from the data, in Appendix B.

The restriction in equation 3 is akin to imposing that beliefs evolve over time in line with rational expectations. (See Benitez-Silva and Dwyer (2005), for a test of rational expectations in this context.) In short, it implies that beliefs must form a martingale. This rationality condition effectively places restrictions on the form of π. In fact the belief of a reform happening in the future is just the accumulation of beliefs for a succession of future single periods. The number of future periods is declining as the individual ages. Therefore, even though beliefs are martingales, the reported belief of a reform happening before retirement will drift down over time.

We allow for some uncertainty to be resolved completely before the earliest retirement date. Specifically, and naturally, we assume that uncertainty about standard retirement age, and retirement age rules, is resolved before the individual can retire. The effect of this choice is that we rule out the possibility that individuals might retire early to 'lock-in' existing pension rules, as discussed in Benitez-Silva et al. (2009).[8] We discuss how we match this restriction to the data in section 4.

Importantly, and, as discussed above, we allow for individuals to make mistakes about their beliefs. These mistakes affect how beliefs evolve. Specifically, going forward, we allow for individuals to learn about the 'correct' probability of reform. Therefore, the transition matrix Π which features in the individuals' decision problem doesn't necessarily correspond with the observed distribution of changes. Intuitively, individuals with beliefs that are far away from the correct beliefs should converge to those correct beliefs over time, as long as some public information is available. At the same time, individuals still make idiosyncratic adjustments to their beliefs, based on private signals. This latter force causes the belief heterogeneity to perpetuate. The model of learning, and its estimation, is discussed in more detail in section 4.

[7] We model beliefs to be given according to a discrete distribution. In truth, of course, beliefs are given continuously. We use discrete distributions for notational convenience and because they correspond more closely with the empirical application later. This choice has no effect on results.

[8] As discussed later, we allow for early retirement, and we model the actuarial adjustments that apply to early benefit withdrawal. We model these rules as being tied to the SRA. In our model, therefore, uncertainty about early retirement penalties is also resolved while the individual is still of working age.

2.3 Preferences

Following Low et al. (2010), the household head has standard within-period preferences over consumption with proportional costs of work.

$$U\left(C_t, F_t\right) = \frac{\left(C_t \kappa\right)^{1-\gamma}}{1-\gamma} \tag{4}$$

where $U\left(\right)$ is the intratemporal felicity function, and C_t is consumption at time t. γ is the coefficient of relative risk aversion, constant over working life and over consumption levels. κ captures effective of consumption net of costs of work, determined by labour market choices. It is determined as follows:

$$\kappa = \left(1 - \eta_{Ft}\right)\left(1 - \eta_{REt}\right)$$

where η_{Ft} is the cost of work for labour-market outcome F_t, when employment continues from the previous period. In addition, whenever the individual re-enters the labour market, he incurs an additional cost η_{REt}, discussed further below. Overall, therefore, the disutility effect of work is equivalent to a proportional drop in consumption. In this model, the individual can work full-time (FT), part-time (PT), be unemployed (U) or be retired (Ret). We impose that $\eta_{FT,t} > \eta_{PT,t} > \eta_U = \eta_R = 0$. Moreover, following several authors (French and Jones, 2011 and Van der Klaauw and Wolpin, 2008), we impose that preferences for leisure, reflected in the costs of work, increase with age at a constant rate. In our model, therefore $\eta_{FT,t} = \eta_{FT,0} \times \eta_t$ and $\eta_{PT,t} = \eta_{PT,0} \times \eta_t$, where $\eta_{F,0}$ represents the initial cost of work (at age 50) and η_t is the proportional increase, common across time and across all levels of work. Our specification ensures that consumption and leisure are Frisch substitutes and that work lowers utility.[9] As discussed below, we allow for some further unobserved heterogeneity in costs of work.

We model the individual as having subjective expected utility preferences over uncertainties. An alternative is to consider changes to policy as being ambiguous and to consider individuals with ambiguity aversion. However, our data provide subjective probabilities over policy changes, and our approach therefore imposes that individuals form probabilistic beliefs.

The individual alive at age t dies before entering time $t+1$ with probability $1 - s_{t+1}$. The individual values bequests according to the following specification:

$$b\left(M\right) = \xi \frac{M^{1-\gamma}}{1-\gamma} \tag{5}$$

Due to the curvature on this function (given by γ), the individual's preferences over current consumption and bequests are therefore homothetic. The strength of the bequest motive versus current consumption is given by ξ.

Because our analysis focuses on those in late career rather than older retirees, the bequest function might also capture other motives to save into old age. These motives might be: an intrinsic value to holding

[9] In other words, as wages decrease towards retirement, workers enjoy more leisure (holding the marginal utility of wealth constant) and require less consumption.

wealth, or precautionary saving for uncertain medical expenses, for example. Our bequest motive is 'warm glow', rather than dynastic. Because of this, our model does not allow for Ricardian equivalence in cuts to state pensions. As such, however, we follow most similar analyses which ignore such effects when considering the effects of state pension reforms.

2.4 Labour market choices

As discussed above, the model features a variety of labour market states. When the individual is of working age, he can choose to work full-time or part-time. He can also choose to leave the labour market either by becoming unemployed, or if he is old enough, by retiring. Additionally, while choosing to work, the individual may be hit by an unemployment shock. We model a period as a year. This choice contrasts with say, Low et al. (2010), who use a quarterly time frame both in order to pick up the fast dynamics of unemployment and because unemployment insurance in the US has a short duration. However, a yearly time frame is defensible because we are concentrating on older workers for whom labour market dynamics are much slower, and because, in Europe, individuals receive unemployment insurance much longer.

We treat retirement as being irreversible and is associated with receiving social security payments.[10] Moreover, we assume that everyone must retire by age 70 and that all uncertainties (except for mortality) are resolved by then. When the individual is unemployed he must wait for a job offer before being able to rejoin the labour market. When the unemployed individual gets an job offer, he must pay the one-off cost of work, described above, to rejoin the labour market. The probability of finding a job depends on age, but is otherwise common across types. In our model unemployment is a catch-all for the many reasons to be out of the labour market, which would include disability or care provision, and our treatment, of course, neglects this heterogeneity in job finding rates.

We allow for part-time work because, even though part-time workers are a small fraction of the total labour force (amongst men), such work is potentially an important route into retirement. The choice of working hours is state dependent; we assume that the individual can switch costlessly from full-time to part-time work but the reverse is costly. While working, the individual is hit by productivity shocks that affect wages. We assume that earnings are simply the product of hours worked and the individual's wage rate. This contrasts with, say, French and Jones (2011), who allow for part-time work to incur a wage penalty.

2.5 Budget constraint

The asset accumulation equation is:

$$M_{t+1} = R\left(M_t + Y_t - C_t\right) \tag{6}$$

[10] See, for example, French and Jones (2011) who distinguish decisions over work from decisions to claim pension benefits.

where M_t is time-t assets, C_t is consumption and Y_t is total post-tax income. Income, discussed below, comes from four sources: earnings, social security, unemployment insurance and the spouse. R is the return on end-of-period assets. For computational reasons, our treatment requires that asset income is not included in the individual's taxable income. This modeling does not reflect the tax system in most EU countries.(Citation.) Our approach effectively implies that everybody pays the same tax rate on asset income and we therefore calibrate the return R to the post-tax rate.

The budget set for the individual is:

$$C_t \in [0, M_t] \tag{7}$$

Consumption is bounded below by Inada conditions. We assume that consumption is bounded above by beginning-of-period wealth, i.e. we assume strict non-zero borrowing limits. We do this because, among those with different housing tenures, renters have little access to debt at this stage in their lifecycle, and during the period we are covering. Moreover, even homeowners have mostly paid down their mortgages approaching retirement and find it hard to obtain new mortgage financing. (Citations.)[11]

2.6 Income

2.6.1 Earnings

As is standard in the literature, pre-tax wages evolve as a slow-moving persistent process:

$$\begin{aligned} \ln W_t^P &= \rho \ln W_{t-1}^P + \zeta_t && (8) \\ \ln W_t &= X_t \mu + \ln W_t^P && (9) \end{aligned}$$

where W_t^P is the idiosyncratic, stochastic component and W_t is the observed hourly wage. The idiosyncratic component is subject to a mean-zero shock ζ_t each period, and reverts towards the mean according to the autoregressive parameter ρ. X_t reflects observable and predictable characteristics of the individual; in our case age. μ is the return to these characteristics.

Yearly pre-tax earnings YL_t are realized as follows:

$$YL_t = \begin{cases} 2000 \times W_t \times \exp \epsilon_t & \text{if } F_t = FT \\ 800 \times W_t \times \exp \epsilon_t & \text{if } F_t = PT \\ 0 & \text{if } F_t = Ret, Inv, Vol \end{cases} \tag{10}$$

[11] In the way we've constructed the model, individuals' access to resources is counter-factually too restricted. Their cash on hand could be interpreted as $M_t + Y_t > M_t$ which should sustain higher consumption than our model implies. Our tighter model is constructed for computational simplicity. One way to rationalize our model choice is that individuals make their expenditures before their earnings turn up: effectively they are paid at the end of the period. Therefore, in our approach, individuals know their current-period income Y_t but cannot spend against it. The only substantive implication of our approach is that if individuals get a very high shock to income in the current period then they may be constrained from spending this. Alternatively, if they get a low shock to income then, in our model, this is not reflected in current constraints.

such that individuals can work full time, for 2000 hours per year, or part time at a 40% full-time equivalent (800 hours). Individuals receive no labour earnings if they are retired, or either voluntarily or involuntarily unemployed. Earnings are also composed of a mean-one transitory shock, $\exp \epsilon_t$, which reflects short-run bonuses or expenditures taken as salary sacrifice. In terms of practicalities, transitory shocks are important for the computation of the model solution because they smooth out kinks in the policy and value functions.

We assume that individuals only receive shocks to their wages while working. This is partly to make computation simpler and partly because workers rarely return to the labour market with higher wages. We do not include wage penalties for those re-entering the labour market; instead our model implies that individuals hold out for a job at their last-received wage, but that these jobs only rarely turn up.

2.6.2 Social security system

We model a stylized pension wealth accumulation process that matches many systems across Europe. In particular, it follows closely the Germany system, as discussed by Haan and Prowse (2014). We assume middle earners accumulate pension rights (or 'points') according to the average of their 20 highest years of earnings. Similar to the system in Germany we also assume that these pension points are capped at a certain fraction of average earnings, in our case twice the average. The effect of this cap is that an individual can't obtain high pension entitlements purely by having a single year of extraordinarily high earnings.

In the model, as in all countries, the final pension received depends on age at retirement. We assume that the standard retirement age under the status quo policy is 65. This is the case for all countries in the data except for France where it is 60. We allow for early or late retirement: individuals can retire after 57 but on reduced benefits; individuals can also retire later than 65 on benefits that are enhanced but not actuarially fairly. Retiring at the standard retirement age is therefore the most financially rewarding path, assuming that the individual has participated for sufficiently long in the labour market. Individuals cannot claim state pension before 57 and must claim by 70, as discussed above. Finally, as in Germany, we assume that individuals who are involuntarily unemployed accumulate pension rights according to the wage in their last-held job.

Formally, pension rights evolve as follows:

$$\mathbb{I}_{it+1} = \begin{cases} \mathbb{I}_{it} + \max\left(\frac{\left(\min\left(Y_t, 2\times\bar{Y}_t\right) - \alpha_t \mathbb{I}_{it}\right)}{25}, 0\right) & \text{if working} \\ \mathbb{I}_{it} + \max\left(\frac{\left(\min\left(800\times W_t, 2\times\bar{Y}_t\right) - \mathbb{I}_{it}\right)}{25}, 0\right) & \text{if eligible for UI} \\ \mathbb{I}_{it} \times adj\left(P_{jk}\right) & \text{if choosing retirement} \\ \mathbb{I}_{it} & \text{if already retired} \end{cases} \tag{11}$$

where $\mathbb{I}_{it}$ captures pension rights, $\bar{Y}$ is average income within the age group, $adj\left(P_{jk}\right)$ is the actuarial adjustment for retiring either early or late, which depends on policy P_{jk}.

Social security benefits are then determined as a function of accumulated pension points and the policy in place.

$$ss_t \begin{cases} = rr\left(P_{jk}\right) \times \mathbb{I}_{it} & \text{if retiring at the SRA} \\ = rr\left(P_{jk}\right) \times \mathbb{I}_{it} \times adj\left(P_{jk}\right) & \text{otherwise} \end{cases} \tag{12}$$

where $rr\left(P_{jk}\right)$ is the replacement rate, again conditional on current policy.

2.6.3 Other income and taxes

Income comes from several sources. Aside from earnings and pensions, individuals may receive income from unemployment insurance (when unemployed), asset income, social assistance and from the spouse.

We model a stylized entitlement-based unemployment insurance system. This type of system is common across Europe. The German system, discussed by Haan and Prowse (2010), is typical. A summary of the German system is as follows: individuals who are eligible for UI receive income of around 50-60% of the net income from their last job for an entitlement period. The entitlement period is determined by the amount of time that the individual has previously spent working capped at an overall maximum depending on age. The individual builds up entitlement periods at the rate of half the length of their previous employment (unless entitlement has previously been exhausted by an unemployment spell). Alongside the concept of entitlement period comes eligibility status. Individuals are eligible for UI when entering unemployment on three conditions. First, they have built up some entitlement. Second, they have worked one period during the last three years. In practice this means that individuals are not eligible if they have transitioned quickly from unemployment to employment and back to unemployment. Third, and finally, they must have left their job involuntarily.

Accordingly, we model the stylized UI insurance system in the following way: the individual's benefit rights accumulate in line with the state pension, reflecting past employment-related contributions. Only individuals who are involuntarily separated from their job receive UI. Everybody who is eligible for UI receives a full entitlement for a limited period (in our case 3 years). The benefit level is based on the full-time income at the last wage rate. (We do not keep track of whether the individual was working full-time or part-time.) Thereafter anyone who is eligible receives entitlement indefinitely based on their accumulated work history, but at a much reduced rate. The difference in this modeling versus standard rules is that we vary the *level* of the benefit according to the individual's entitlements rather than the *length* of entitlement spell. This simplifies the computation by allowing us not to keep track of long unemployment spells. Our treatment would be equivalent to monitoring the length of spell if individuals knew the duration of unemployment and smoothed income fully. However, by allowing for the full entitlement for the limited period we ensure that individuals are fully compensated during quick unemployment spells. Finally, similarly to our treatment of earnings, we assume some (small) transitory fluctuations to UI payouts. Again this helps smooth the value function. Formally, UI is specified as follows:

$$UI_t = \begin{cases} 0.5 \times W_t \times 2000 & \text{if } D_t \leq 2 \\ \mathbb{I}_t \times W_t \times 2000 \times 0.1 & \text{if } D_t > 2 \end{cases} \tag{13}$$

where D_{it} indicates the duration of the current spell. While the individual has been unemployed for less than 3 periods, he receives 50% of the full-time earnings (wage W_t multiplied by hours, 2000). When the individual is long-term unemployed, he receives a reduced benefit equal to the last full-time wage multiplied by 10% of his normalized, pension entitlement $\mathbb{I}_t$.

Spousal income is modelled as being a fixed proportion of the household head's earnings. (***** State Exactly how.)

The individual receives net income after gross income from all sources (apart from asset income) has been taxed. The tax function τ guarantees a level of subsistence to the household: tr. We set a minimum income floor equal to 10% of average full-time earnings.

$$Y_t = \tau\left(YL_t + UI_t + ss_t + sp_t\right) \tag{14}$$

$$\tau\left(x\right) = \max\left(tr, 0.7x\right) \tag{15}$$

2.7 Recursive formulation and model solution

When expressing the model recursively, it is helpful to distinguish between the labour market state when entering period t and the labour market activity after making optimizing choices. For an individual entering period t in the labour force, working full time, the value function $\mathcal{V}_t^{WFT}()$ is given as follows:

$$\mathcal{V}_t^{WFT}\left(M_{it}, W_{it}, \mathcal{B}_{it}, \mathbb{I}_{it}\right) = \max\left[\mathbb{1}_{a<a_{Ret}}\left(V_t^{wft}\right), \mathbb{1}_{a<a_{Ret}}\left(V_t^{wpt}\right), \mathbb{1}_{a<a_{Ret}}\left(V_t^{unm}\right), \mathbb{1}_{a\geq a_{ER}}\left(V_t^{ret}\right)\right] \tag{16}$$

where V_t^{wft} is the conditional value function when choosing to remain working full-time, V_t^{wpt} is the value of stepping down to working part-time, V_t^{unm}, is the value of choosing to exit the labour force to unemployment and V_t^{ret} is the value of retiring. The state variables for such an individual are private savings, M_t, the wage, W_t, beliefs, $\mathcal{B}_{it}$, and social security entitlements, $\mathbb{I}_{it}$.

At each point on the state grid, the individual chooses the best of the labour-market options, subject to age restrictions. The restriction that the individual can't work or be unemployed (without retiring) beyond compulsory retirement age is given by $\mathbb{1}_{a<a_{Ret}}()$. The indicator function $\mathbb{1}_{a\geq a_{ER}}()$ is used because the individual can't retire before the minimum eligible age.[12]

Similarly the value functions for individuals entering the labour market as working part-time $\left(\mathcal{V}_t^{WPT}\right)$,

[12] The function $\mathbb{1}()$ here indicates whether or not a choice can be included in the maximization, rather than taking values 0 and 1 strictly speaking.

unemployed $\left(\mathcal{V}_t^{UNM}\right)$ or retired $\left(\mathcal{V}_t^{RET}\right)$ are:

$$\begin{aligned}
\mathcal{V}_t^{WPT}\left(M_{it}, W_{it}, \mathcal{B}_{it}, \mathbb{I}_{it}\right) &= \max\left[\mathbb{1}_{a<a_{Ret}}\left(V_t^{wft} s.t. \eta_{REt} \neq 0\right), \mathbb{1}_{a<a_{Ret}}\left(V_t^{wpt}\right), \ldots \right. && (17)\\
&\qquad \left. \mathbb{1}_{a<a_{Ret}}\left(V_t^{unm}\right), \mathbb{1}_{a\geq a_{ER}}\left(V_t^{ret}\right)\right] \\
\mathcal{V}_t^{UNM}\left(M_{it}, W_{it}, \mathcal{B}_{it}, \mathbb{I}_{it}, D_t\right) &= \max\left[\mathbb{1}_{a<a_{Ret}}\left(V_t^{unm}\right), \mathbb{1}_{a\geq a_{ER}}\left(V_t^{ret}\right)\right] && (18)\\
\mathcal{V}_t^{RET}\left(M_{it}, \mathcal{B}_{it}, \mathbb{I}_{it}\right) &= V_t^{ret}\left(M_{it}, E_{it} = RET, \mathcal{B}_{it}, \mathbb{I}_{it}\right) \text{ if } a \gneq a_{ER} && (19)
\end{aligned}$$

where V_t^{unm} is the value function conditional on being unemployed at the end of period t, and D_t is the duration of receipt of UI payments. Equations 17 and 18 imply that those unemployed may have the option to retire, depending on age restrictions. E_{it} is the state variable for labour market status on entering period t.

The conditional value functions on the right hand side of, say, equation 16 capture optimizing choices about consumption. They are given as follows:

$$\begin{aligned}
V_t^{wft}\left(M_{it}, E_{it}, W_{it}, \mathcal{B}_{it}, \mathbb{I}_{it}\right) &= \max_{C_{it}} U(C_{it}, .) + \beta\mathbb{E}_t\left\{s_t\left[(1-\delta)\mathcal{V}_{it+1}^{WFT} + \delta\mathcal{V}_{it+1}^{UNM}\right] + (1-s_t)b(M_{it+1})\right\} && (20)\\
V_t^{wpt}\left(M_{it}, W_{it}, \mathcal{B}_{it}, \mathbb{I}_{it}\right) &= \max_{C_{it}} U(C_{it}, .) + \beta\mathbb{E}_t\left\{s_t\left[(1-\delta)\mathcal{V}_{it+1}^{WPT} + \delta\mathcal{V}_{it+1}^{UNM}\right] + (1-s_t)b(M_{it+1})\right\} && (21)\\
V_t^{unm}\left(M_{it}, W_{it}, \mathcal{B}_{it}, \mathbb{I}_{it}, D_t\right) &= \max_{C_{it}} U(C_{it}, .) + \beta\mathbb{E}_t\left\{s_t\left[\Lambda\mathcal{V}_{it+1}^{WPT} + (1-\Lambda)\mathcal{V}_{it+1}^{UNM}\right] + (1-s_t)b(M_{it+1})\right\} && (22)\\
V_t^{ret}\left(M_{it}, E_{it}, \mathcal{B}_{it}, \mathbb{I}_{it}\right) &= \max_{C_{it}} U(C_{it}, .) + \beta\mathbb{E}_t\left[s_t V_{it+1}^{RET} + (1-s_t)b(M_{it+1})\right] && (23)
\end{aligned}$$

subject to laws of motion given in all the equations above. In these equations δ is the job destruction probability, Λ the job finding probability when unemployed. The individual discounts the future at rate β. When choosing to stay in the labour market the individual chooses between full-time and part-time work. The entering labour market status E_{it} matters for those ending up retired, because the timing of entry into retirement affects social security entitlement. In all these equations $\mathbb{E}_t$ is the expectation over productivity risk and policy uncertainty: The mortality and employment risks are specified in more detail within the equations.

We simplify the belief state objects in the following way. In the description above, the individual has beliefs about a reform happening before a certain age. We construct the model in that way to correspond to the data. When solving the model, however, we transform this variable to represent the belief that a certain policy will be in place for a certain age. If $\tilde{p}_{B,t}$ captures the latter belief of pension generosity, then

$$\tilde{p}_{B,t} = \begin{cases} p_{B,t} & \text{if } \mathbb{P}_{B,t} = 0 \\ 1 & \text{if } \mathbb{P}_{B,t} = 1 \end{cases}$$

whereas, by definition,

$$p_{B,t} = 0 \text{ if } \mathbb{P}_{B,t} = 1$$

The effect of this transformation is to eliminate $\mathbb{P}_{B,t}$, the presence or not of the reform, as a state variable in the choice problem. In effect, in the individual's choice problem there is no difference between a world in which the individual knows with 100% certainty that there will be a reform in the future, and a world in which that reform has already happened.

We, likewise, work with a transformed versions, $\tilde{\Pi}_B$ and $\tilde{\Pi}_R$, of the transition matrices. The transformed version of $\tilde{\Pi}_B$ has entries $\tilde{\pi}_{xy}^B$ given as follows:

$$\tilde{\pi}_{xy}^B = \begin{cases} \pi_{xy}^B \left(1 - p_{B,it}^*\right) & \text{if } \mathscr{P}_x \neq 1 \\ \pi_{xy}^B + p_{B,it}^* & \text{if } \mathscr{P}_x = 1 \& \mathscr{P}_y \neq 1 \\ 1 & \text{if } \mathscr{P}_x = 1 \& \mathscr{P}_y = 1 \end{cases}$$

Notice, however, that when we simulate forwards to match the data, we must work with the original variable and track the presence of reforms. This is because some of those who believe that a reform will happen with near certainty will, in reality, be mistaken. Whereas, when a reform has actually happened, the individual is assumed to know the full details. When simulating forwards, therefore, only some individuals will experience a reform. This approach corresponds well with our pan-European dataset where reforms hit some, but not all, individuals. Again, we discuss how we fit the model to the data in more detail in sections 4 and 5.

The model has no analytical solution. It is solved numerically backwards from the final period at age 105, when we assume that all individuals die. We solve the model using the method of endogenous grid points extended to discrete choices as discussed in Fella (2014), Iskhakov et al. (2012) and Clausen and Strub (2013). In short, in each period we solve the conditional value functions given in equations 20, 21, 22 and 23, then obtain the full value functions by taking the maximum. It is important to emphasize that there is no resolution of uncertainty between labour supply and consumption decisions, and so we can treat these as consecutive or simultaneous without loss of generality. The numerical solution is extremely accurate. We discuss it further in Appendix A.

3 Data

3.1 Selection of dataset for profiles and definition of variables

We start from the complete SHARE sample for waves 1-2-4-5, taking countries Austria, Germany, Sweden, Netherlands, Spain, Italy, France, Denmark, Switzerland, Belgium. From this sample, we keep only men aged [50,65] (30,970 observations), and we exclude one observation for which there is a coding error on assets, 1853 observations whose employment status is permanently sick/disabled, 252 observations with missing employment status, 239 with missing hours of work for which we cannot distinguish full or part time. The final sample size is 28,625.

For earnings and wealth, we use values imputed by the developer of SHARE (SHARE, 2013). We average across the five replications of imputed values distributed to users and we convert all income and assets data to German 2005 euro using PPP exchange rates. Wealth includes financial assets (bank accounts, government and corporate bonds, stock and shares, mutual funds, retirement accounts, contractual saving, life insurance holdings) and real assets (value of main residence and other real estate, value of own share of businesses, value of cars), net of mortgages or loans on main residence and of other debts. Earnings are the sum of annual earnings from employment (last year) and annual earnings from self-employment. Earnings are gross of income and social contributions in wave 1, net in the following waves. We therefore exclude wave 1 when we use income data.

Our reference person lives in a couple. We therefore multiply all household income and assets by 1.5 (OECD modified equivalence scale) for those who are single (no partner in the household). We ignore the fact that there may be adult children, assuming that all wealth and income is mostly coming from the couple.[13]

Weekly hours of work are obtained from the question on usual hours per week, excluding meal breaks but including paid or unpaid overtime. We sum over the main and, where present, secondary job. We define part time as those with weekly hours 30 or less.

We calculate total annual hours as usual week hours, times 4.33 weeks, times usual months worked in a year (including payed holidays), multiplied by 47/52 to account for 5 weeks of paid holidays. Hourly earnings are obtained by dividing last year earnings by total annual hours. There are 25 cases in which usual months per year are missing, 131 cases who report hours of work, 719 who report zero earnings (which is possible given the presence of self-employed) and 9 who miss both. These cases, which account for around 5% of the employed in the selected sample, are excluded in estimating the earnings profile.

3.2 Initial distribution

Our methodology requires a complete set of observations of all variables for the population of 50 year olds (see section 5), corresponding to the first year of the model solution. For this initial sample, we have to drop wave 1 from the main sample, because earnings are before taxes, as discussed above. We also drop wave 4, because expectations data are not asked to the longitudinal respondents (who are the entire sample for Sweden and Germany). We are left with 381 observations. There are few retired, less than 1% of the sample, that we re-assign to being unemployed, because our model does not allow for retirement earlier than 57. We also replace negative assets with 0. This affects 5% of the sample.

To get a full sample for the initial distributions, we need to impute some missing observations. We have two cases. First, among those who are employed, some report zero earnings or zero hours of work, and some do not report beliefs about retirement reforms. Among these employed individuals, we use a mean matching regression approach, as in French and Jones (2011). We first predict the value (for all the employed sample) of the variables (earnings, hourly wage, expectations) based on a regression on log household income, log assets, being in Southern European countries, being a citizen, being part-time

[13] For financial wealth, which is collected only for responding persons (aged 50+) and for their partners, there may be some issues if there are other individuals aged 50+ in the household. However, in the selected sample only 0.5% of individuals live in households where there are three or more adults aged 50+. For household income, also income for any other non-responding household member is collected.

and educational level (dummies).[14] We predict these values for all the employed individuals, both with missing and non missing values. Separately for each variable, we then split the entire sample of employed individuals (both with missing and non missing values) in deciles of the predicted variable. For each individual with a missing value we randomly choose a donor for the residuals from those with no missing, but who belong to the same decile of the predicted value. We sum up the predicted value and the residual.

Next, for the unemployed, we observe neither beliefs nor hourly wage for the unemployed. These unemployed account for around 5% of the sample of 50 year olds. In order to preserve the joint distribution of hourly wage and beliefs, we follow a different procedure, choosing the "closest" employed individual to use as a donor for these values. To identify a proper donor, we first estimate a propensity score for the probability of being unemployed on log household income (and its square), log assets (and its square), being in Southern European countries, being a citizen, and educational level (dummies). We then split the sample of unemployed in three groups, at the .33 and .66 percentiles of the propensity score. For each unemployed we randomly extract (with replacement) a donor from the employed individuals whose propensity lies in the same interval (defined by the min and max of the tertile in the unemployed group). From the donor we directly get the value of hourly wage and expectations about reforms (the imputation is, therefore, joint for the three variables). Their earnings are still set to zero.

4 A Process for the Evolution of Beliefs

4.1 Empirical implementation

From Section 2 we know that ex-ante individuals should expect their future beliefs to be consistent with their current beliefs, so that

$$E(p_{j,it+1}|p_{j,it}\&\mathbb{P}_{j,t+1}=0)=\sum_{x}\pi^{j}_{xy}\mathscr{P}_{x}.$$

This ex-ante expected transition matrix is the one that the individuals should take into account when taking their choice at time t. However, we allow for the possibility that individuals make mistake and that, when moving forward, they adjust their beliefs through learning. We assume that individuals may learn about the "true" chances of reform by observing some private or public information. We do not model this learning component explicitly, but we allow individuals to revise their beliefs proportionally with respect to the current distance to the "true" ones, assumed to be captured by average beliefs. Over time, therefore, we should expect individual beliefs to converge to "objective" beliefs. Beliefs may converge more or less quickly depending on how much information individuals receive. In order to account for this adjustment in the simulations, we estimate an "ex-post" transition matrix from SHARE data. To this purpose we work with a simplified model:

$$E(p_{j,it+1}|p_{j,it}\&\mathbb{P}_{j,t+1}=0)=\pi_0+\pi_1 p_{j,it}+\gamma\left(p_{j,it}-\bar{p}_{j,t}\right).$$

[14] For earnings and hourly wage we actually use their logs and we then simply get the exponential of the prediction plus imputed residual.

where the γ coefficient captures the speed of adjustments to the true chances of reform ($\bar{p}_{j,t}$). In line with previous discussions, beliefs should decrease over time, so it should be that $\pi_1 < 1$. Rationality further impose that $\pi_0 = 0$. The reason is that, when $p_{j,it} = 0$ and there is no learning, then it should be that $\pi(p_{j.it}) = 0$. See also Bernheim (1990) and Benitez-Silva and Dwyer (2005) for empirical evidence on the absence of an intercept in a related context.[15]

Furthermore, as discussed in Bernheim (1987), expectations also react to the arrival of new information, although they do not completely adjust to it, implying again that individuals make, from an external point of view, mistakes. Our empirical model for the evolution of beliefs is as follows:

$$p_{j,it+1} = \pi_1 p_{j,it} + \gamma\left(p_{j,it} - \bar{p}_{j,t}\right) + X_{it+1}\beta_X + C_{t+1}\beta_C + \varphi_{j,it+1} \tag{24}$$

$$E(\varphi_{j,it+1}|p_{j,it}, \bar{p}_{j,t}, X_{it+1}, C_{t+1}) = 0 \tag{25}$$

where $\bar{p}_t$ is the objective probability of reform, which is unknown and needs to be empirically modeled. γ captures the speed of learning. X_{it} is a $1 \times K_X$ vector of observables for the individual at $t+1$, which address unexpected changes in their conditions, such as health deterioration, which may induce a change in their beliefs. C_{t+1} is a vector of $1 \times K_C$ country observations, such as announcements of reforms, or innovations to the government's budgetary position. $\varphi_{j,it+1}$ is the individual-specific innovation, also assumed to be mean zero with respect to the information set in the previous period.

The main issue in estimating an empirical counterpart for equation 24 is that $\bar{p}_t$ is not observed. Nevertheless, one could simply rewrite the model as

$$\begin{aligned} p_{j,it+1} &= \gamma\bar{p}_{j,t} + (\pi_1 + \gamma)p_{j,it} + X_{it+1}\beta_{X+1} + C_{t+1}\beta_{C+1} + \varphi_{i,t+1} \\ &= \alpha_t + \theta p_{j,it} + X_{it+1}\beta_{X+1} + C_{t+1}\beta_{C+1} + \varphi_{j,it+1} \end{aligned}$$

which would still allow us to obtain the quantities required to estimate the two transition matrices. In this model, the correct chances of a reform enter the model as a time drift. However, in our sample and in the empirical application the situation is more complicate. First of all, we have multiple countries, that have different "true" chances of a reform. Secondly, reforms may not be universal, and therefore they may not affect all categories of individuals. For this purpose, our empirical model also includes country dummies, plus country and individual covariates at t, which are used to proxy the true chances of reform:

$$p_{j,it+1} = \alpha_t + \alpha_c + X_{it}\beta_X + C_t\beta_C + \theta p_{j,it} + X_{it+1}\beta_{X+1} + C_{t+1}\beta_{C+1} + \varphi_{j,it+1}. \tag{26}$$

For these regressions we selected, from the main sample (Section 3.1) only employed individuals, with no missing data in the belief questions, aged 50-64. We can use transitions only across two couples of waves: waves 1-2 and waves 4-5. The reason is that only longitudinal respondents are asked about beliefs in wave 4. We assume that there are always two years between waves (this is not true for all individuals in the

[15]Although in their case the expectation is not related to an event which may take place in the following year, and therefore rationality imposes $\pi_1 = 1$.

first two waves), so we can estimate only the relation across two years. Therefore, the actual model being estimated is:

$$p_{j,it+2} = \tilde{\alpha}_t + \tilde{\alpha}_c + X_{it}\tilde{\beta}_X + C_t\tilde{\beta}_C + \tilde{\theta}p_{j,it} + X_{it+2}\tilde{\beta}_{X+2} + C_{t+2}\tilde{\beta}_{C+2} + \tilde{\varphi}_{j,it+2}. \tag{27}$$

In principle, there is a problem of omitted variables, because we miss X_{it+1} and C_{t+1}. We assume that the same variables at t and $t+2$ are good enough proxies, so that the approximation error is (i) uncorrelated with p_{it} and (ii) with negligible variance. It follows that $\tilde{\theta} = \theta^2$ and $\tilde{\varphi}_{j,it+2} \approx \theta\varphi_{j,it+1} + \varphi_{j,it+2}$.

Another problem is that expectations are measured with error. To solve this, under the assumption that current beliefs over p_{it} capture the available information, we use the expectation over the other margin of reform as an instrument for current beliefs. More explicitly, if $p^e_{j,it} = p_{j,it} + e_{j,it}$ is the observed belief, we assume classical measurement error so that $E[e_{j,it}|p_{j,it}] = 0$ and $E[e_{j,it}\tilde{\varphi}_{j,it+2}] = 0$. We further assume that the measurement error is uncorrelated with all the other covariates, so that we can treat them as exogenous. The equation becomes

$$p^e_{j,it+2} = \tilde{\alpha}_t + \tilde{\alpha}_c + X_{it}\tilde{\beta}_X + C_t\tilde{\beta}_C + \tilde{\theta}p^e_{j,it} + X_{it+2}\tilde{\beta}_{X+2} + C_{t+2}\tilde{\beta}_{C+2} + \tilde{\varphi}_{it+2} - \tilde{\theta}e_{it} + e_{it+2} \tag{28}$$

Note that we need the other margin of reform $p_{-j,it}$ to be (i) excluded from the structural equation ($E[\tilde{\varphi}_{j,it+2}|p_{-j,it}] = 0$); (ii) mean independent from current and future measurement error on the other margin.

A final concern is related to the fact that we are imposing a linear model. The linear restriction seems to be supported by the data: Figure 1 shows that a 5th degree polynomial fit does not differ much from a simple linear function. We also experimented with IV estimates (see below) that include higher order polynomials, estimated both using standard 2SLS (and using higher order terms of the instrument as additional excluded variables) and using a control function approach. Predicted values are still showing an almost linear pattern.

For the chances of a reform that increases retirement age before the individual goes into retirement, results (Table 2) show a range of $\tilde{\theta}$ from .425 to .515, which would imply a θ from .652 to .718. Note that estimates do not change much across specifications. Covariates are statistically significant, but their introduction does not change the beliefs persistence much.

For the chances of a reform that reduces pension benefits before the individual goes into retirement, results (Table 2) show a range of $\tilde{\theta}$ from .438 to .620, which would imply a θ from .662 to .787.

In the simplest specification, with no covariates, we know that the attenuation bias of OLS is equal to the signal-to-total variance ratio if we also assume that the measurement error is uncorrelated over time:

$$\tilde{\theta}_{OLS} = \frac{cov(p^e_{j,it}, p^e_{j,it+2})}{var(p^e_{j,it})} = \frac{cov(p_{j,it} + e_{j,it}, \tilde{\alpha} + \tilde{\theta}p_{j,it} + \tilde{\varphi}_{j,it+2} + e_{j,it+2})}{var(p_{j,it} + e_{j,it})} = \frac{\sigma^2_p}{\sigma^2_p + \sigma^2_e}\tilde{\theta}.$$

Figure 1: $p_{R,it+2}$ as a function of $p_{R,it}$ (chances of a reform that increases retirement age); polynomial fit (degree 5, OLS estimation, men only)

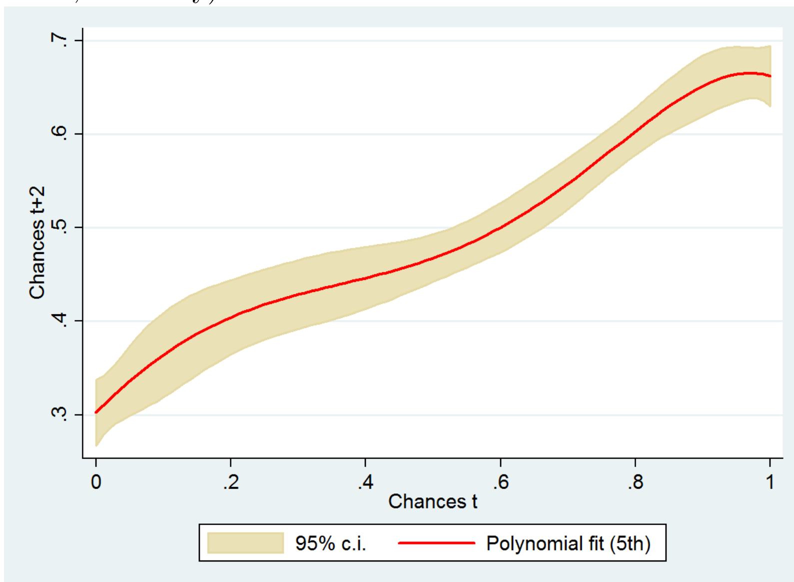

From comparing IV and OLS estimates, we get that the proportion of total variance in beliefs due to measurement error is around 30-40%. As mentioned above, we need to get the variance $\sigma_\varphi^2$of innovations $\varphi_{j,it+1}$. First, we can get the variance of the measurement error as:

$$\sigma_e^2 = \left(1 - \frac{\sigma_p^2}{\sigma_p^2 + \sigma_e^2}\right) \sigma_{p^*}^2.$$

Secondly, we know that the variance of the residuals in equation 28 is equal to $(1+\theta^2)\sigma_\varphi^2 + (1+\theta^4)\sigma_e^2$, from which we can obtain σ_φ^2. The standard deviation for the innovation is approximately 0.15-0.20.

The assumption of classical measurement error is not likely to hold in this context, because the variable $p_{j,it}$ is bounded at 0 and 1. In this case, as discussed by Black et al. (2000), the measurement error is likely to be negatively related with the true variable, because it can only be positive when $p_{j,it} = 0$, while it can only be negative when $p_{j,it}$=1. In practice, the measurement error in this case is likely to bias the observed variable toward the mean, so that the (latent) simple regression of the observed value on the true one has a positive intercept and a slope smaller than 1. This implies that also the instrument is necessarily related to the measurement error, because its relation with the true value makes it by construction negatively correlated with e_{it}. As a result, Black et al. (2000) prove that $\tilde{\theta}_{IV}$ is upward biased, so that the true $\tilde{\theta}$lies between $\tilde{\theta}_{OLS}$ and $\tilde{\theta}_{IV}$. In our case bounds are quite informative. Taking the minimum and maximum estimates from the tables (across different specifications and estimation methods) we get a range for θ equal to $[.514, 718]$ in the case of chances of an increase in NRA and equal to $[.536, 787]$ for chances of a decrease in benefits. Further standard algebra also shows that $\tilde{\theta}_{OLS}/\tilde{\theta}_{IV}$ provides a downward estimate for the fraction of variance in the observed $p^e_{j,it}$ that can be attributed to the true variance of $p_{j,it}$, and therefore we should be attributing to much variance to measurement error. Although we cannot improve these bounds, they are informative on the possible distortions related to our choice and can be used as a guideline for the sensitivity analysis.

4.2 Belief transitions and policies in the main retirement model

- We estimate no-reform and reform policies.
  - For pension generosity, we estimate a no-reform replacement rate of 0.8 and a reformed generosity of 0.5.
  - For retirement age, our model assumes that individuals can retire early with some adjustment to their pension entitlement.
  - The post-reform rules therefore modify these adjustments. For example whereas in the no-reform world individuals can retire from 65 with full entitlement, after a reform, they would only receive xx% of their full replacement rate. To get a full pension, they need to wait until 67 to retire.
  - Table [TO DO] captures the policies in each scenario.
- The belief variable in our model matches that in the data with a slight exception. In the model uncertainty is resolved by a fixed date, whereas individuals in the survey are asked about reforms before they retire, which is endogenous. Adapting the model variable to take account of endogenous retirement makes the dynamic programming problem significantly harder. Nevertheless we are careful to match the variables in the estimation is closely as possible. For example, when computing the yearly subjective chance of reform we use each individuals estimated retirement age, taked from the survey.
- When solving the dynamic programming problem, the state variables are beliefs about which of these states will apply in future.
  - We discretize this distribution, for both the beliefs about generosity and the beliefs about reform.
  - The state variable can therefore take one of n values between 0 and 1.
  - As discussed later, we assign to each individual at age 50 a (discretized) belief, coming from the data, and taking into account measurement error.
- A key input into the model is the transition matrices of these beliefs.
  - Here there are two relevant inputs: first, there are observed transitions of beliefs, taking into account learning about the true measure.
  - Second there are the ex-ante expected transitions of beliefs reflecting the distribution of shocks that individuals expect to receive.
  - These second transitions respect the martingale property of expectations, assuming rational decision makers.

Table 2: Regressions of the following wave belief about chances [0-1] of a reform on current beliefs and covariates

	Reform increasing NRA ($p_{R,it+2}$)		Reform decreasing benefits ($p_{B,it+2}$)	
	OLS	IV	OLS	IV
No covariates (weak test)				
$p_{j,it}$ [0-1]	.359***	.515***	.372***	.62***
	(.019)	(.04)	(.019)	(.04)
Constant	.306***	.217***	.315***	.177***
	(.013)	(.024)	(.013)	(.024)
N	2766	2766	2766	2766
H0: coeff on $p_{j,it} = 1$ (p-value)	.000	.000	0.000	0.000
First-stage F		734		752
With individual and country covariates at t, plus country and interview year dummies (strong test)				
$p_{j,it}$	.264***	.425***	.288***	.438***
	(.02)	(.045)	(.02)	(.043)
Constant	2.02	1.28	4.00***	3.61***
	(1.46)	(1.45)	(1.34)	(1.23)
N	2738	2738	2738	2738
H0: coeff on $p_{j,it} = 1$ (p-value)	.000	.000	0.000	0.000
H0: coeff on individual covar = 0 (p-value)	.000	.000	0.000	0.000
H0: coeff on country covariates = 0 (p-value)	.000	.000	0.000	0.000
H0: coeff on optimism controls = 0 (p-value)	.468	.385	0.000	0.000
H0: country&year dummies=0 (p-value)	.000	.000	0.000	0.000
First-stage F		547		560
With individual and country covariates at t and t+2, plus country and interview year dummies (strong test)				
$p_{j,it}$	.273***	.432***	.287***	.453***
	(.02)	(.045)	(.02)	(.043)
Constant	1.9	1.44	6.09**	5.25**
	(2.79)	(2.68)	(2.59)	(2.51)
N	2717	2717	2717	2717
H0: coef on $p_{j,it} = 1$ (p-value)	.000	.000	0.000	0.000
H0: coef individual covar (at t) = 0 (p-value)	.061	.1	.161	.167
H0: coeff country covar (at t) = 0 (p-value)	.503	.55	.011	7.8e-03
H0: coeff on optimism controls = 0 (p-value)	.034	.014	0.000	0.000
H0: country&year dummies=0 (p-value)	.000	.000	.013	.064
H0: coef individual cov (at t+2) = 0 (p-value)	.249	.165	.352	.309
H0: coeff country covar (at t+2) = 0 (p-value)	.363	.354	.537	.508
H0: coef optimism controls (t+2)=0 (p-value)	.011	.000	.029	.019
First-stage F		539		556

Note: * p<.10 ** p<.05 *** p<.01. S.e. clustered by household in parentheses. Individual covariates: log household income, log household net assets, age, dummies for: public sector, self-employed, part-time, education, citizenship, general health, marital status, household size, interview year. Country variables: long-term interest rate on state bonds, government debt over gdp, government deficit, per-capita gdp, dummies for reform and for reform announcement. Optimism controls: chances tomorrow is sunny (residuals after a regression on dummies for interview year, month and country), difference between reported chances to be alive at 75 minus survival rate derived from age-specific mortality tables (by country and gender, in year 2015), dummy for reported chances of survival missing (for these obs, the difference is imposed to be zero). In 2SLS regressions the instrument is $p_{-j,it}$, that is the chances of reform on the other margin.

  - An example may usefully highlight the constrast between these two transition matrices.
  - Suppose an individual believes with 99% certainty that there will be a reform. Then,...
  - ...
  - Estimation of these transition matrices is discussed above in section 4.

- When simulating the model forwards, we take the observed transitions of beliefs, which includes learning about the process.
  - One difficulty is matching the model to a world in which many different reforms may happen.
  - As discussed, our model is simplified, in that we only allow for 1 reform.
  - We do this as follows. When simulating forwards and matching to the data, we allow reforms to happen to a subset of our many simulated individuals.
  - One way to think of this is that our model mixes individuals from across the EU and so a reform needn't happen to everyone.
  - When a reform hits these individuals, they then report a belief of 0 for future reforms.
  - In this way, the average reported chance of reform drifts down over time.
  - This drift matches that seen in the data, due because the time gap to retirement is decreasing with age.
  - This average drift is consistent with the beliefs about future reforms being correct on average.

- The belief concept in the model therefore matches that in the data with one exception:
  - In those countries where there has recently been a reform, in the model individuals believe that the post-reform policy will happen with probability 1.
  - In the data, on the other hand, these individuals will report a belief of 0 of future reform.
  - This is, of course, assuming that reforms cannot be reversed and there are not multiple reforms.
  - When there might be multiple reforms or reversals, the relevant probabilities might be close to 1 and close to 0, but the same principal carries through.

- Figure **??** shows paths for beliefs and the evolution of reforms.

5 Estimation of the Retirement Model

5.1 Overview

We estimate the model by simulated method of moments (SMM) using two steps. As has been standard since at least Gourinchas and Parker (2002), we first pre-estimate several items from the data and fit further key parameters using the full SMM algorithm. The main component of the first stage is the estimation of beliefs, discussed in Section 4. In this section we give an overview of the remaining first-stage inputs and the second stage. Precise details of the second stage, in particular, are left to Appendix F.

5.2 First-stage inputs (incomplete)

5.2.1 Wage process

We estimate a wage process using the SHARE data. We estimate an age profile for the average hourly wage:

$$\ln W_{it} = \mu_0 + \mu_c + \mu_y + \sum_{k=50}^{65} \mu_{tk} 1[t=k] + \ln W_{it}^P + \epsilon_{it} \tag{29}$$

where μ_c are country effects, μ_y are wave effects and μ_{tk} captures the age profile of interest. $\ln W_{it}^P$ captures the idiosyncratic (persistent) wage of individual i at time t, and ϵ_{it} captures transitory fluctuations, assumed to be measurement error. To estimate this we select from the main sample (Section 3.1) all employed individuals, waves 2-4-5, truncating the sample at the 5th and 95th percentiles. We predict mean wages over 50-65, referencing to Germany and wave 2 (2006/07). We then extrapolate the growth rate from 66 to 70. All individuals are forced to retire at 70, so wage profiles thereafter are irrelevant. To take into account non-random selection into employment (vs unemployment or retirement), we follow the procedure in French (2005), discussed further in the Appendix.

As final inputs, we need to estimate dynamic components of the wage process. These components are the persistence of shocks, ρ, underlying dispersion in persistent productivities, $\sigma^2_{W^P}$, and the variance of shocks, σ^2_ζ. Due to the limited time series in the SHARE data and because we only observe individuals every 2 years, we do not estimate persistence of shocks, but rather impose it to be 0.975 per year. This estimate is in line with most studies (French and Jones (2011)). With this persistence imposed we can identify the dispersion in productivities and the variance of shocks from simple covariance restrictions for the longitudinal dependence of $\tilde{w}_{it} = \ln W_{it}^P + \epsilon_{it}$, which is the observed residual wage, net of common components. In practice, data limitations allow us to use only waves 4-5 to estimate, which are at two years distance. We exploit the fact that

$$\sigma^2_{W^P} = \frac{1}{\rho^2} cov\left(\tilde{w}_t, \tilde{w}_{t+2}\right).$$

These inputs imply a variance of underlying shocks, because $\sigma^2_\zeta = \left(1-\rho^2\right)\sigma^2_{W^P}$. From the data, we get a s.d. of the shocks $\sigma_\zeta = 0.0965$, in line with Meghir and Pistaferri (2010), and $\sigma_{M^P} = 0.4345$. Further details on all aspects of the wage process are in appendix D.

5.2.2 Mortality data

Mortality data come from Eurostat. The data are cohort mortality rates for an individual aged 50, from 2013 (currently available data). These are projections of future death rates for current individuals, from age 50 to 101. We use male mortality and take a raw average over the 20 countries included which overlap with the wider SHARE dataset. The mortality rate at 101 is around 40% per annum. We extrapolate mortality rates from 101 to 105 linearly and assuming certain death in the final year. The life expectancy at 50 is 85.6. The chance of living to 100 at age 50 is around 2%, and doesn't reach 5% until the individual is 87 years of age.

5.3 Choice of moments for second-stage SMM

We use data from across Europe, and compute life-cycle moments accounting for country and wave effects. To account for different standard retirement ages across countries, we normalize retirement ages to being 65. Our age concept is therefore years before retirement age. We use the sample from waves 1-2-4-5, excluding those with missing employment status, but including those with missing information on hours worked.

We compute the 1/3rd and 2/3rd asset quantiles by age, giving us $2T$ moments. We further compute the proportion working full-time, the proportion working part-time, and the proportion unemployed, all by age and wealth tertile. This gives us a further $3 \times T \times 2 = 6T$ moments, giving us 8T moments in total. Complete details of these computations are given in appendix E.

5.4 Initial conditions

Because we only solve the dynamic programme back to age 50 we face an initial conditions problem. In particular we have to explain the wealth accumulation of households up to age 50. Similarly, household labour supply and pension accrual depends on behaviours before 50. Rather than solving the model back to the beginning of life, we tackle this problem in a similar way to French and Jones (2011). Our strategy is based on two main features: first, initiating the simulations of the model using real data; and second allowing for, and estimating, unobserved heterogeneity in preferences.

To incorporate the micro-data for our initial distribution we take the following steps. We replicate the initial sample 100 times. The population of 50 year olds in the main sample selected from SHARE is 381 (Section 3.2); we therefore end up with 38,100 simulated individuals. We do this for wealth, for wages, and for work status.

For those who are unemployed, we amend this procedure slightly. Because we do not observe the length of time the unemployed have been out of the labour market, we impute durations. We do this by taking the job-destruction and job-finding parameters. We compute the steady-state proportions of unemployed in each duration assuming no additional voluntary separations, and make random assignments among the simulated individuals. Notice that because each real observations is re-created 100 times, we can roughly match the distribution of unemployment durations within each real observation. Similarly, we do not observe pension wealth entitlement. In this case we... (**what do we do? We'll have to do something better).

When assigning the initial distribution of beliefs, we face the problem that these beliefs are measured with substantial error. Ignoring this problem would bias our estimates of the relationships between beliefs and other observables, especially in the first few years of our simulations. Rather than attempting to back out the underlying distribution using a full deconvolution, we 'subtract' measurement error in a simpler way. We do this by simulating a possible 'true' underlying value for each of the 100 copies of each observation. To simulate these underlying values we use conditional distributions from the multivariate gaussian. We implement the procedure using estimates of the proportion of dispersion in beliefs due to measurement error obtained in section 4. Specifically, for each observed belief b^* we simulate a distribution of true beliefs according to the formula

$$b = \bar{b^*} + \left(b^* - \bar{b^*}\right)\frac{\sigma_b}{\sigma_{b^*}}\omega + \sqrt{(1-\rho^2)}\sigma_b Z$$

where: $\bar{b^*}$ is the mean of observed beliefs, assumed equal to the mean of true beliefs (because the average measured belief is assumed correct); σ_{b^*} is the standard deviation of observed beliefs; σ_b is the standard deviation of true beliefs, easily obtained once we have quantified the measurement error; ω is the correlation between observed and true beliefs, with $\omega = \frac{\sigma_e}{\sigma_{b^*}}$, and where σ_e is the standard deviation of measurement error, and Z is the standard normal variable. Intuitively, we add an uncentred random noise to each observation, moving it closer to the average belief. This simulated distribution therefore has a much smaller variance than the observed distribution. Finally, we assign these underlying beliefs to our discrete grid. Note that this formulation assumes 'classical' measurement error, a strong assumption given that beliefs are bounded between 0 and 1.

Second, we allow for unobserved heterogeneity in discount rates. Such unobserved heterogeneity is important in explaining the large heterogeneity in wealth accumulation and in labour supply patterns, particularly towards the end of working life. Our method is the standard approach, developed in Heckman and Singer (1984) (see also, for example Van der Klaauw and Wolpin (2008)). We allow for a finite number of unobserved types (in this case 2). The probability of belonging to each type, with respect to the econometrician's information set, is a logistic function of variables observed at the initial age, 50. We use, as regressors, beliefs (both of benefit reform and retirement age), wage level, pension entitlement, wealth level, and wealth/wage ratio. We estimate parameters of this logistic function at the same time as estimating the main preference parameters.

5.5 Imposed parameters

We use an interest rate of 2%. Our bequest parameter is 80% which implies a marginal propensity to consume in the final period of life of.... A part-time job is 40% of full-time hours. The job separation probability is 0.05. The job finding probability is 0.2. We impose a variance of yearly transitory shocks to earnings of 0.04. For computational reasons we also impose a small transitory shocks on unemployment income. This is also 0.04. We impose a UI replacement rate during entitlement of 50%, a UI replacement rate upon exhaustion of 10%. Spousal income is set at 50% of average individual income. As for the tax function: the (linear) tax rate is 30% above the minimum income threshold. The minimum income is set at 10% of average full-time gross earnings. (**This all needs to be discussed much more.)

Table 3: Final-Stage Parameter Estimates

Description	Label	Estimate	
Discount rate, type 1	β_1	0.960	(0.004)
Discount rate, type 2	β_2	0.982	(0.006)
Cost of FT work, age 50	η_{FT}	0.348	(0.009)
Cost of PT work, age 50	η_{PT}	0.174	(0.074)
Yearly incr. in work cost	η_t	0.024	(0.002)
Cost of lf re-entry	η_{RE}	0.101	(0.066)

Perhaps most importantly, our results depend crucially on risk preferences. These preferences determine, for example, the welfare costs of the risk associated with policy uncertainty. Given that these preferences have been extensively investigated using a variety of approaches, our approach is to impose commonly-used values, and to rely on estimates of other preference parameters, such as the time discount rate and the consumption costs of work, to fit the data. We perform two estimations. We first use a comparatively high average coefficient of relative risk aversion of 5, the approximate value estimated by French and Jones (2011) when evaluating the benefits of Medicare medical insurance. Second we use a more conservative CRRA parameter of risk aversion of 1.5 as estimated by Attanasio and Weber (1995). In both cases, the value of γ required to generate these preferences depends on the costs of work, which are changing over the life-cycle We therefore fit γ to match these risk preferences on average over all our simulated individuals over their entire life-cycles. As a final point, our approach is further justified by discussions in, for example, Gourinchas and Parker (2002) and Guvenen and Smith Jr (2014). They discuss how difficult it is to estimate risk aversion cleanly and separately from the discount rate using life-cycle wealth formation as the main piece of evidence.

5.6 Estimated parameters

- Table 3 shows the parameters used to fit the model.

5.7 Model fit

- Figure 2 shows the empirical and simulated patterns of labour supply.
- Figure 3 shows the empirical and simulated patterns of total wealth and wealth dispersion.

Figure 2: Labour Force Status: Model and Data

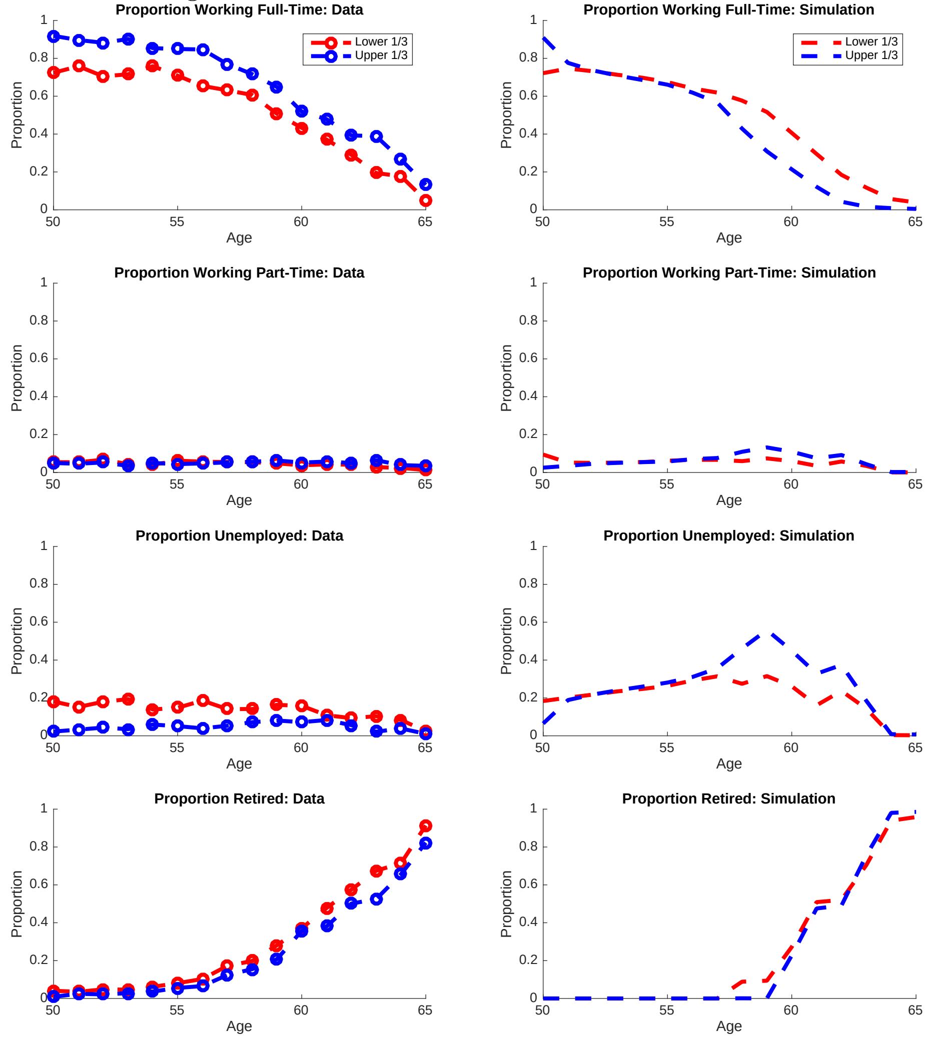

Figure 3: Average Wealth: Model and Data

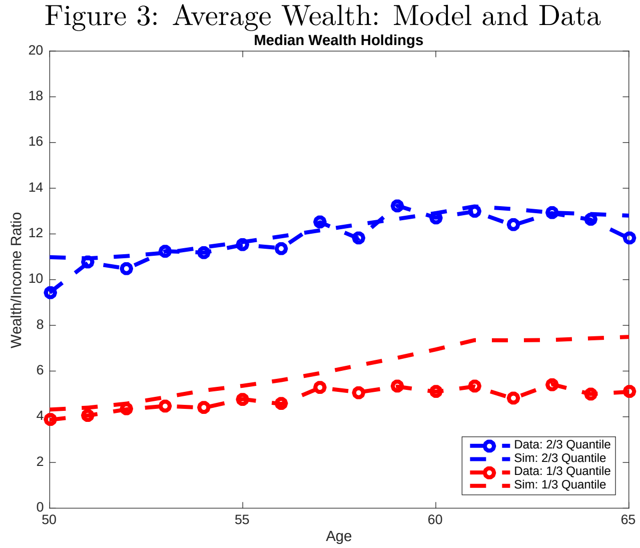

6 Model Validation

Before moving on to the implications of the model we validate of the model by displaying the fit against moments that have not been used in the estimation. The core of the model is a well-used specificaton for explaining labour supply behaviour at the end of the life. Therefore we validate our specific model by focusing on the beliefs of reform and the effects on behaviour. We focus on the expectation of retiring at age 63. Remember that these expectations are asked of all the survey respondants. However, these expectations are not explicitly computed in the model. To this end we compute each individual's expectation of working at age 63 as follows: We compute for each simulated individual whether or not they do indeed end up working full-time at age 63 (remember to correct for France). We then use a linear probability model to regress this indicator in each year on all information available to individuals: wealth, labour income, pension wealth as well as interactions and higher-order terms of these. The R^2 of this regression in later years is around 0.8. We take the predicted value of this regression to be the probability of working. We then regress this predicted value on a the full information set of the individual just as we do in the data. The results are shown in table 4.

7 Model Results

7.1 The role of belief heterogeneity in explaining heterogeneity of outcomes

Our model includes several standard features which drive heterogeneity in outcomes within a cohort, such as shocks to productivity and working status. Our model is novel in that we also allow for heterogeneous beliefs. We now explore how important this belief heterogeneity actually is, using various measures.

We first examine how belief heterogeneity affects the level and evolution of consumption. This evolution is shown in figure 4. It shows the average consumption path for extreme 'optimists' compared to extreme 'pessimists' coming from the simulations. Specifically, each line in the figure shows results from regressions of consumption at future ages up to 60 on a full set of indicators in three base years: 50, 53 and 56.[16] We plot the predicted coefficient on having belief of 0 that a reform will happen (the optimists) versus a belief of 1 (the pessimists). The figure can be thought of as showing impulse response functions: it plots the dynamic response to initial differences.

The bottom line in the figure shows that, conditional on other characteristics, optimists about the level of future benefits consume more than pessimists on average by 7% at age 50. Thereafter, the effect of initial heterogeneity quickly dies away. By the age of 55, optimists consume only 3% more than initial pessimists, and by the age of 60, the initial difference has died away completely. This declining effect comes about because beliefs are autoregressive. As both optimists and pessimists learn on average about the objective chance of reform, their mistakes dissipate and their behaviours converge. This is notwithstanding the fact

[16] These indicators are a 4-th order polynomial in assets, earnings, productivity level, pension entitlement, employment status and indicators for beliefs. The R^2 from this regression in the base year is 90%, declining to 70% in the final year.

Figure 4: Belief heterogeneity and the evolution of consumption: optimists vs pessimists

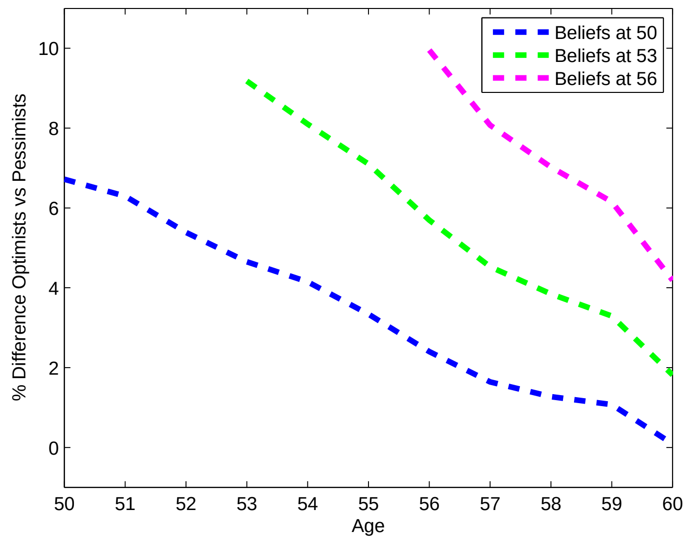

Notes: Plots show predicted average consumption after regressions of consumption on characteristics in base year, taken from simulations. These indicators are: a 4-th order polynomial in assets, earnings, productivity level, pension entitlement, employment status and indicators for preferences and for beliefs. The R^2 from this regression in the base year is 90%, declining to 70% in the final year.

that households receive idiosyncratic shocks to their beliefs, and reforms are continually occurring at a uniform rate to everyone.

The other lines in the figure show that the profiles are similar if we choose different base years. The effect on consumption levels starts high and quickly dissipates. The figure also implies, however, that the effect of belief heterogeneity is larger at older ages, for example reaching 11% at age 56. Finally, note that this figure only shows the effect from beliefs about benefit levels. In fact, the effect from heterogeneous beliefs about retirement age is far lower.

Another important implication of our model is that we can contrast the effect of belief heterogeneity with other 'un-observed' heterogeneities, principally differences in preferences. Contrasting the role of beliefs versus preferences in explaining outcomes is important: it has been a major motivation in driving the research agenda on eliciting beliefs (see Manski (2004)). Our model provides a good framework to assess these contributions, even though we capture only a small subset of all the possible heterogeneities in beliefs, preferences and other circumstances.

Here we concentrate on two outcomes: consumption growth and total wealth holdings. We use consumption growth, because, although we presented the consumption level in figure 4, the Euler equation for consumption growth includes the discount rate directly. In a simple life-cycle model, those with more patience will have a higher consumption growth, whether or not we condition on assets. In contrast, for example, the effect of the discount rate on the level of consumption comes mainly through endogenous effects on wealth accumulation and the sign on preferences depends on whether or not we condition on assets.

For each of these outcomes, we therefore calculate the contribution of the various components of heterogeneity as follows. We non-parametrically regress the outcome variable on beliefs and the preference indicator.[17] For consumption growth, we regress on the change in beliefs, and for wealth level we use the belief level.[18] The change in beliefs, rather than the level, is a more natural driver of changes to consumption. We compute the contribution by taking the R^2 from these regressions, as well as the R^2 when we remove each of the regressors in turn.

Table 5 shows the results. The results for consumption growth are shown in the table's first column. The first row shows the total contribution of both the variance of (changes to) beliefs and the variance of discount rates to the variance of consumption innovations. Unsurprisingly, perhaps, the total contribution is very small. This is because the vast majority of shocks to consumption come through labour-market innovations. That being said, the next three rows give the breakdown of interest. These rows show that innovations to consumption are caused far more by innovations to beliefs about pension generosity than by changes to beliefs about retirement age. However, innovations to consumption are explained overall more by preference heterogeneity.

The second column shows the results for consumption growth when we only look at households who are not borrowing constrained. The contribution of belief and preference heterogeneity should be higher in this case because consumption for households at the constraint is driven entirely by current circumstances i.e. by labour market outcomes. To compute the contribution, we now condition on households who ended the previous period with wealth at more than half the median income, thus removing the bottom 9% of the wealth distribution. The first row shows that belief and preference heterogeneity do indeed explain more of the total variance, although the relative contribution from each factor is similar to that shown in the first column.

The final column of the table shows the results for the dispersion in wealth. Within this column, the first row shows that the total contribution from heterogeneity in preferences and beliefs is again small.[19] Moving down, the next three rows show that, in this case, preferences account for almost all of total wealth heterogeneity. Again, this is perhaps not too surprising: a major reason for including heterogeneity in discount rates in this type of model is to explain the large dispersion in wealth accumulation that seemingly cannot be explained entirely by labour market outcomes. On the other hand, the preliminary discussion of the data in section 3 showed that policy beliefs and household wealth have little correlation.

As a final point it is worth comparing results across the columns of table 5. The contribution from preference heterogeneity is larger for total wealth because, even though its effect on consumption innovations in any year is small, these effects are perfectly correlated over time, and, more importantly, they persist from the beginning of the planning horizon right to the end. On the other hand, shocks to policy beliefs, for example, are likely only relevant later in the households' working life: their effect on younger households is much smaller because pension wealth is a much smaller component of life-time wealth.

[17] Because beliefs are discretized we can run the regression non-parametrically by having a full range of dummies.

[18] We have 4 different levels for beliefs over benefit reform, 3 different levels for beliefs over retirement age, and 2 levels of preferences. This implies requiring 15 dummies for the beliefs of changes to benefit reforms, and 8 dummies for the retirement age beliefs. We do not use interactions between these beliefs nor interactions with other characteristics.

[19] In this case this contribution is no larger if we only examine those households who are not constrained.

Table 4: Model Validation

	Chances work FT after age 63 ([0,1] scale)		Expected age at collection of benefits (years)	
	Data	Simulation	Data	Simulation
Chances gov raises ret age before s/he retires (scale [0,1])	.188***	.093	.704***	1.346
	(.019)	(.001)	(.222)	(.008)
Chances gov reduces pension before s/he retires (scale [0,1])	.084***	.123	1.059***	1.593
	(.019)	(.001)	(.229)	(.011)
Obs	4,552		5,772	
Selection	employed	employed	employed	employed
			entitled pension	
	aged [50,60]	aged [50,60]	aged [50,65]	aged [50,60]
	waves 2-4		waves 1-2-4	

Note: *** p-value < 0.01; OLS regression; s.e. clustered at the household level in brackets. The regression includes income, wealth, food consumption, rent related expenditure, dummies for: age, health, marital status, type of job, education, wave, country, wave $\times$ country.
Results from simulations are OLS regressions on simulated data from baseline scenario. Regressors include, income, wealth and year dummies. Variable "Chance work FT after age 63" is calculated as a regression of observed outcomes (at 63) on variables observed at time t: beliefs of reforms, income, wealth, pension wealth and higher order terms of these. Also includes heterogeneous preferences. "Expected age at collection of benefits" is generated similarly. These expectations variables are generated from simulations when reported beliefs are, on average, correct. See text for more details.

Table 5: Relative Importance of Preference Heterogeneity Versus Belief Heterogeneity

	Cons growth	Cons growth for unconstrained	Wealth
% contr. From Beliefs and Preferences	2.0%	9.0%	1.7%
of which, % from			
(Revisions to) beliefs ret. Age	1%	1%	1%
(Revisions to) beliefs pens. Gen	27%	23%	0%
Preference heterogeneity	72%	76%	98%

Notes: First row shows R^2 from regression of relevant observable on indicators for beliefs and preferences. Each subsequent row shows R^2 from regression on a subset of the regressors, shown as a percentage of the first row. In the first and second columns the regressors are changes to beliefs and indicators for preferences. In the final column the regressors are belief levels and preferences. All regressions are performed nonparametrically with a full set of dummies for beliefs (levels or changes) and preference types. The data are pooled observations of households between 50 and 55.

Figure 5: Welfare costs of uncertainty

Figure xxx

7.2 The welfare effect of uncertainty (to do)

Welfare effects are much larger for those with no labour earnings. Point made recently by xxx.

Figure 5

8 Conclusions (to do)

References

Attanasio, O. and G. Weber (1995): “Is consumption growth consistent with intertemporal optimization? Evidence from the consumer expenditure survey,” *Journal of Political Economy*, 103, 1121–1157.

Baker, S., N. Bloom, and S. Davis (2015): “Measuring economic policy uncertainty,” *NBER Working Paper*, wp21633.

Benitez-Silva, H., S. Brook, D. Dwyer, F. Heiland, and W. Sanderson (2009): “Retirement and Social Security Reform Expectations: A Solution to the New Early Retirement Puzzle,” *Mimeo*.

Benitez-Silva, H. and D. Dwyer (2005): “The rationality of retirement expectations and the role of new information,” *Review of Economics and Statistics*, 87, 587–592.

Bernheim, B. (1987): “Social Security benefits: An empirical study of expectations and realizations,” *NBER Working Paper*, wp2257.

——— (1990): “How do the elderly form expectations? An analysis of responses to new information,” *Issues in the economics of aging*.

Black, D., M. Berger, and F. Scott (2000): “Bounding parameter estimates with nonclassical measurement error,” *Journal of the American Statistical Association*, 95, 739–748.

Burnside, C., M. Eichenbaum, and S. Rebelo (2011): “Understanding booms and busts in housing markets,” *NBER working paper*, wp16734.

Caliendo, F., A. Gorry, and S. Slavov (2015): “The Cost of Uncertainty about the Timing of Social Security Reform,” *NBER working paper*, wp21585.

Cao, D. (2013): “Speculation and financial wealth distribution under belief heterogeneity,” *Mimeo*.

Clausen, A. and C. Strub (2013): “Envelope Theorems for Non-Smooth and Non-Concave Optimization,” *Mimeo*.

Cribb, J., C. Emmerson, and G. Tetlow (2013): “Incentives, shocks or signals: labour supply effects of increasing the female state pension age in the UK,” *IFS Working Paper*, W13/03.

Fella, G. (2014): “A generalized endogenous grid method for non-smooth and non-concave problems,” *Review of Economic Dynamics*, 17, 329–344.

Fernández-Villaverde, J., P. Guerrón-Quintana, K. Kuester, and J. Rubio-Ramírez (2011): “Fiscal volatility shocks and economic activity,” *NBER working paper*, wp17317.

French, E. (2005): “The effects of health, wealth, and wages on labour supply and retirement behaviour,” *The Review of Economic Studies*, 72, 395–427.

French, E. and J. Jones (2011): “The Effects of Health Insurance and Self-Insurance on Retirement Behavior,” *Econometrica*, 79, 693–732.

Giavazzi, F. and M. McMahon (2012): “Policy uncertainty and household savings,” *Review of Economics and Statistics*, 94, 517–531.

Gourinchas, P. and J. Parker (2002): “Consumption over the life cycle,” *Econometrica*, 70, 47–89.

Guvenen, F. and A. Smith Jr (2014): "Inferring labor income risk from economic choices: An indirect inference approach," *Econometrica*, 82, 2085–2129.

Haan, P. and V. Prowse (2010): "The Design of Unemployment Transfers: Evidence from a Dynamic Structural Life-Cycle Model," *mimeo.*

——— (2014): "Longevity, life-cycle behavior and pension reform," *Journal of Econometrics*, 178, 582–601.

Heckman, J. and B. Singer (1984): "A method for minimizing the impact of distributional assumptions in econometric models for duration data," *Econometrica*, 52, 271–320.

Iskhakov, F., J. Rust, and B. Schjerning (2012): "A Generalized Endogenous Grid Method For Discrete-Continuous Choice," *mimeo.*

Kitao, S. (2014): "Sustainable social security: Four options," *Review of Economic Dynamics*, 17, 756–779.

Low, H., C. Meghir, and L. Pistaferri (2010): "Wage risk and employment risk over the life cycle," *American Economic Review*, 100, 1432–67.

Luttmer, E. and A. Samwick (2012): "The Welfare Cost of Perceived Policy Uncertainty: Evidence from Social Security," *Mimeo.*

Manski, C. (2004): "Measuring expectations," *Econometrica*, 72, 1329–1376.

Mastrobuoni, G. (2009): "Labor supply effects of the recent social security benefit cuts: Empirical estimates using cohort discontinuities," *Journal of Public Economics*, 93, 1224–1233.

Meghir, C. and L. Pistaferri (2011): "Earnings, consumption and lifecycle choices," in *Handbook of labor economics*, vol. 4, 773–854.

Van der Klaauw, W. and K. Wolpin (2008): "Social security and the retirement and savings behavior of low-income households," *Journal of Econometrics*, 145, 21–42.

Van Der Wiel, K. (2008): "Preparing for policy changes: social security expectations and pension scheme participation," *IZA Discussion Papers*, wp3623.

A Model solution: to be updated

- Use the method of endogenous grid points extended to discrete choices as discussed in Fella (2014)Fella (2014), Iskhakov et al. (2012) and Clausen and Strub (2013)

- Put simply, we solve endogenously the consumption rule at each end-of-period asset point conditional on each labour force decision.

  – We work with the conditional value functions given above in equations 20, 21, 22 and 23. More specifically, in each period we solve for each choice. For example

$$\begin{aligned}
V_t^{wft}(X_t) &= \max_{C_t} U(C_t, FT_t) + \beta \mathbb{E}_t \left(s_t \left(\delta \mathcal{V}_{t+1}^{WRK}(X_{t+1}) + (1-\delta) \mathcal{V}_{t+1}^{INV}(X_{t+1}) \right) + (1-s_t) b(A_{t+1}) \right) \\
V_t^{wft}(X_t) &= \max_{C_t} U(C_t, FT_t) + \beta \mathbb{E}_t \left(s_t \left(\delta \mathcal{V}_{t+1}^{WRK}(X_{t+1}) + (1-\delta) \mathcal{V}_{t+1}^{INV}(X_{t+1}) \right) + (1-s_t) b(A_{t+1}) \right)
\end{aligned}$$

  – This gives us 4 conditional value functions $V_t^{wft}(X_t)$, $V_t^{wpt}(X_t)$, $V_t^{unm}(X_t)$,$V_t^{ret}(X_t)$. From these we can work out the unconditional value functions according to equations 16, 18, 17 and 19. Note that our only friction is the choice that is available (people do not have to pay a cost to re-enter the labour market for example).

  – i.e. for our conditional value functions our state variables are

$$\left\langle \underbrace{M_{it}}_{\text{assets}}, \underbrace{W_{it}}_{\text{wage}}, \underbrace{\mathcal{B}_{it}}_{\text{beliefs}}, \underbrace{\mathbb{I}_{it}}_{\text{pension rights}}, \underbrace{D_t}_{\text{duration}}, \underbrace{F_{it}}_{\text{l.f. status}} \right\rangle$$

  where, if F is working or voluntarily unemployed, then D_{it} is redundent.

- We use 100 grid points for assets.

  – We use 200 grid points for end-of-period assets.

- The overall problem is non-concave. Therefore consumption can be non-linear in end-of-period assets. For example it may be the case that those with high wealth at the beginning of the period choose low consumption and high savings because they want to retire earlier in future, whereas those with lower wealth at the beginning of the period choose higher consumption and even lower savings. This can be the case even conditional on current labour force choices. Therefore given an ordered grid of end-of-period assets, when we construct the beginning of period grid by $M_t = A_t + C_t$ we can find that the beginning of period grid is non-monotone. We therefore then re-arrange the beginning of period grid, knowing that the consumption rule for each point satisfies an Euler Equation.

  – For example, for the choice of working full-time in the current period:

$$\begin{aligned}
U_c(C_t, FT_t) &= \beta R \mathbb{E}_t \left(s_t \left(\delta \frac{\partial \mathcal{V}_{t+1}^{WORK}(X_{t+1})}{\partial A_{t+1}} + (1-\delta) \frac{\partial \mathcal{V}_{t+1}^{UNEM}(X_{t+1})}{\partial A_{t+1}} \right) + (1-s_t) b'(A_{t+1}) \right) \\
&= \beta R \mathbb{E}_t \left(s_t \left(\delta U_c(C_{t+1}, F_{t+1} | E_{t+1} = WORK) + (1-\delta) U_c(C_{t+1}, F_{t+1} | E_{t+1} = UNEM) \right) + (1
\end{aligned}$$

– where F_{t+1} is choice of work in period $t+1$ conditional on labour force status E_{t+1}

- Finally given all the grids M_t conditional on each labour force choice, and the final grid M_t taken as the maximum of this, we make the consumption rule more precise by moving grid points such that the consumption 'drop' is almost vertical rather than gently sloped.
- We then compute the best of the discrete choices at a beginning-of-period grid point.
- This method focuses on accuracy of the consumption-saving rule.
- This is important given that the welfare costs of uncertainty here are proportional to the precautionary saving motive (Kimball ??)

$$\mathcal{V}_t^{WORK}\left(M_{it}, W_{it}, \mathcal{B}_{it}, \mathbb{I}_{it}\right) = \begin{cases} \max\left[V_t^{work}\left(\right), V_t^{volu}\left(\right), V_t^{ret}\left(\right)\right] & \text{if } a \geq early-retirement \\ \max\left[V_t^{work}\left(\right), V_t^{volu}\left(\right)\right] & \text{if } a < early-retirement \\ V_t^{ret}\left(M_{it}, E_{it} = WRK, \mathcal{B}_{it}, \mathbb{I}_{it}\right) & \text{if } a = end-work \end{cases}$$

$$\mathcal{V}_t^{VOLU}\left(M_{it}, W_{it}, \mathcal{B}_{it}, \mathbb{I}_{it}\right) = \begin{cases} \max\left[V_t^{volu}\left(\right), V_t^{ret}\left(\right)\right] & \text{if } a \geq early-retirement \\ V_t^{volu}\left(\right) & \text{if } a < early-retirement \\ V_t^{ret}\left(M_{it}, E_{it} = WRK, \mathcal{B}_{it}, \mathbb{I}_{it}\right) & \text{if } a = end-work \end{cases}$$

$$\mathcal{V}_t^{INVU}\left(M_{it}, W_{it}, \mathcal{B}_{it}, \mathbb{I}_{it}, D_t\right) = \begin{cases} \max\left[V_t^{invu}\left(\right), V_t^{ret}\left(\right)\right] & \text{if } a \geq early-retirement \\ V_t^{invu}\left(\right) & \text{if } a < early-retirement \\ V_t^{ret}\left(M_{it}, E_{it} = WRK, \mathcal{B}_{it}, \mathbb{I}_{it}\right) & \text{if } a = end-work \end{cases}$$

$$\mathcal{V}_t^{RET}\left(M_{it}, \mathcal{B}_{it}, \mathbb{I}_{it}\right) = V_t^{ret}\left(M_{it}, E_{it} = RET, \mathcal{B}_{it}, \mathbb{I}_{it}\right) \text{ if } a \gneqq early-retirement$$

- X_t is state vector, including A_t, assets (do more on the state variables... these differ by people)

$$X_{it} = \left\langle \underbrace{M_{it}}_{\text{assets}}, \underbrace{E_{it}}_{\text{l.f. status}}, \underbrace{W_{it}}_{\text{wage}}, \underbrace{\mathcal{B}_{it}}_{\text{beliefs}}, \underbrace{\mathbb{I}_{it}}_{\text{pension rights}}, \underbrace{D_t}_{\text{duration}} \right\rangle \tag{30}$$

- Computation:

– For the estimation we use 60 nodes for the asset grid, 12 nodes for the pension entitlement grid, 4 nodes for the beliefs about pension benefits, 3 nodes for state retirement age, 5 nodes for productivity, and 2 types of unobserved preference heterogeneity. Together with the 4 employment nodes and the time periods, we solve the Euler equation and optimal work choices at just under $10,000,000$ grid points.

A.1 Time Periods

We have 2 divisions of life in the model. The first division is based around the labour market. In the first part of the model, men can only work or choose to be unemployed. In the second part of the model, they can also choose to retire, under varying incentives. In the final part of the model, men must retire.

The variables used to code this are iAGE_EAR_RET (currently set to 57) and iAGE_END_WOR (currently set to 70).

The second division relates to the timing of policy uncertainty. We have that uncertainty about retirement age resolves before the agent can retire. But uncertainty about cuts to the benefit system can resolve later. The variables used to code this are iAGE_RES_R (currently set to 57) and iAGE_RES_B (currently set to 70)

B Details on the Treatment of Beliefs

- we can relate p^* to $\tilde{p}$ by

$$p^*_{B,it} = \begin{cases} 1 - (1 - \tilde{p}_{B,it})^{\frac{1}{T}} & \text{if } \mathbb{P}_{B,t} = 0 \\ 0 & \text{if } \mathbb{P}_{B,t} = 1 \end{cases}$$

- We discretize the beliefs in the following way
  - (Done in SCRP_Init_dibn.m, a hack at the moment)
  - This discretization is designed to preserve the mean and the variance of the belief distribution.
- These belief variables are related to each other.
  - Specifically, if there are no announcement effects, and a reform may arrive in each year with constant probability, then clearly

$$p_{B,it} = 1 - \left(1 - p^*_{B,it}\right)^T$$

    where T is the number of years left until retirement.
  - Therefore, given data on $p_{B,it}$ and T then

$$p^*_{B,it} = 1 - (1 - p_{B,it})^{\frac{1}{T}}$$

Figure 6: Average Pension Accrual By Age

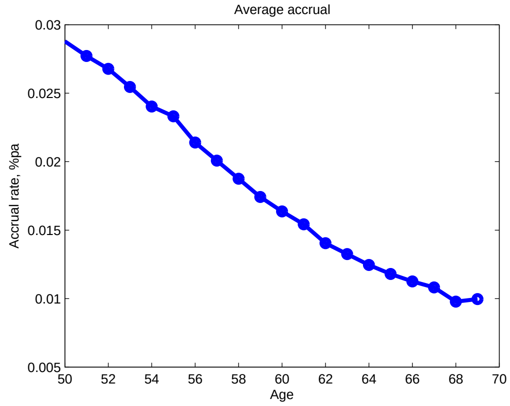

- For retirement age, we do something slightly different. We assume that, for most of working life, reforms arrive in accordance with a retirement age of 65. However, recall that we assume that uncertainty is resolved before the age of early retirement (here 57). Therefore we allow for a one-off increase in the resolution of uncertainty at this age. This is operated as follows: if p_R^* is the yearly probability of a reform to retirement age (calculated as above, according to a retirement age of 65), then this implies $1-(1-p_R^*)^{\hat{T}}\left(1-p_{R,RES}^*\right)=p$, which implies that $p_{R,RES}^*=1-\frac{(1-p)}{(1-p^*)^{\hat{T}}}$.
- Therefore, if we want to allow for policy uncertainty stretching beyond retirement, we might allow that the yearly chance of reform stays constant, but that the new uncertainty is

$$p_{longer,B,it}=1-\left(1-p_{B,it}^*\right)^{T_{longer}}$$

where, for example, T_{longer} might last until age 70 or even further.

C Details of the pension system

Figure 6

D Estimation of the wage process

If

$$\begin{aligned}
\tilde{w}_{it} &= \ln W_{it}^P + \epsilon_{it} \\
\ln W_{it+1}^P &= \rho \ln W_{it}^P + \zeta_{it+1} \\
\implies \tilde{w}_{it+1} &= \rho \ln W_{it}^P + \zeta_{it+1} + \epsilon_{it+1} \\
cov\left(\tilde{w}_{it+1}, \tilde{w}_{it}\right) &= cov\left(\rho \ln W_{it}^P + \zeta_{it+1} + \epsilon_{it+1}, \ln W_{it}^P + \epsilon_{it}\right) \\
&= cov\left(\rho \ln W_{it}^P, \ln W_{it}^P\right) \\
&= \rho Var\left(\ln W_{it}^P\right) \\
&= \rho \sigma_{W^P}^2
\end{aligned}$$

and

$$\begin{aligned}
\ln W_{it+1}^P &= \rho \ln W_{it}^P + \zeta_{it+1} \\
\implies Var\left(\ln W_{it+1}^P\right) &= Var\left(\rho \ln W_{it}^P + \zeta_{it+1}\right) \\
\implies \sigma_{W^P}^2 &= \rho^2 \sigma_{W^P}^2 + \sigma_\zeta^2 \\
\implies \sigma_\zeta^2 &= \left(1-\rho^2\right) \sigma_{W^P}^2
\end{aligned}$$

E Details on SMM moments

Asset quantiles

To estimate on data the .33 and .66 quantiles of the asset distribution by age, we ues the main sample (Section 3.1) and we first remove country and wave effects. We assume that these effects are (i) proportional and (ii) are constant across the distribution of assets, so that we can calculate them at the average:

$$M_t = exp(\gamma_0 + \gamma_y + \gamma_c + \sum_{k=50}^{65} \gamma_{tk} 1[t=k])\eta_t \tag{31}$$

$$E(\eta_t | \gamma_y, \gamma_c, t) = 1$$

This model is estimated using Poisson Quasi-Maximum Likelihood (truncating at the 5th and 95th percentiles of M_t). We then remove from all observations (also below the 5th and above the 95th) the country and wave effects, by fixing all individuals to Germany in 2006/07 (wave 2) and then taking

$$exp(\hat{\gamma}_0 + \hat{\gamma}_{wave2} + \hat{\gamma}_{Germany} + \sum_{k=50}^{65} \hat{\gamma}_{tk} 1[t = k])\hat{\eta}_t$$

The .33 and .66 quantiles by age are estimated on these corrected values in the sample. In the model all values are normalized to the median gross earnings of a 50-year old full time worker. Assets are usually expressed as multiples of this value.

Employment moments

We use the proportion unemployed, by age and wealth quantile (2T moments). To obtain these 6T moments we fit on the main sample, for each y_t, a linear regressions on age and log assets

$$E(y_t|\beta_y, \beta_c, t) = \beta_0 + \beta_c + \beta_y + \beta_t t + \beta_{t2} t^2 + \beta_M ln(M_t) + \beta_{M2} ln^2(M_t) + \beta_{t\times M} t \times ln(M_t) + \epsilon_t \qquad (32)$$

and we predict the values by age at the .33 and .66 asset quantiles, fixing Germany in 2006/07 (wave 2). There are only few cases, for young individuals, where the sum of predicted full-time, part-time and unemployed exceeds one. As the difference is very small, we simply impose that the proportion retired is zero, and we rescale the other three proportions by their sum (so that they sum up to 1).

F Details on second-stage SMM estimation procedure