30 November 2011

Do up-front tax incentives affect private pension saving in the United Kingdom?

Rowena Crawford, Institute for Fiscal Studies

Richard Disney, Institute for Fiscal Studies and University of Nottingham*

Carl Emmerson, Institute for Fiscal Studies

Abstract:

The paper examines how individuals respond to complex decision-making environments – in particular, whether up-front financial incentives are an effective policy lever to change behaviour. The paper argues that incentives differ in their transparency and in their complexity; individuals are more likely to respond to incentives that are both transparent and imply a large pay-off in terms of net income.

The paper focuses on household 'tax planning' in the context of tax reliefs for retirement saving in the United Kingdom. It examines whether take-up of retirement saving instruments increases at the higher rate threshold for income tax, since tax relief is given at the marginal tax rate and should be more attractive to those just above this threshold than to those just below it. It then examines a more complex case where the tax system provides an incentive for pension saving to do be done by one member of a couple. Econometric results are obtained from the Family Resources Survey on these two tests of household responses to complex incentives.

Key words: Tax incentives Retirement saving Regression discontinuity

JEL Classification: E21 H24

Acknowledgments:

This research is funded by the HMRC/HMT/ESRC under project 'Household responses to complex incentives', RES 194-23-0013. Co-funding was received by the IFS Retirement Saving Consortium which comprises Age UK, Association of British Insurers, Department for Work and Pensions, Financial Services Authority, HM Treasury, Investment Management Association, Money Advice Service, National Association of Pension Funds, Partnership Pensions and the Pensions Corporation. We are grateful for comments from Magali Beffy, Mike Brewer, Jonathan Lepper, Gareth Myles and Darren Warren.

*Corresponding author: School of Economics, University of Nottingham. Nottingham, NG7 2RD; tel: 01159515619 email: richard.disney@nottingham.ac.uk

Do up-front tax incentives affect private pension saving in the United Kingdom?

1. Introduction

Tax schedules give strong incentives for individual and household agents to arrange their financial affairs in such a way as to minimise their tax burden. At the same time, complexity in tax schedules makes it difficult for agents to do so. As a general principle, the tax structure is organised in the UK such that higher income individuals face a higher income tax burden (i.e. a rising average tax rate) with step changes in their marginal income tax rate at specified intervals (i.e. the marginal rate structure). Nevertheless, in the UK neither average effective tax burdens nor marginal tax structures rise monotonically with income once the whole range of National Insurance contributions, ceilings on tax reliefs and tax credits are incorporated. Moreover, in tax planning, individuals may focus on the more salient features of their total income (for example, their annual earnings) in computing tax liability, rather than computing their whole taxable income including taxable interest on savings and so on.

The traditional public finance textbook model focuses on individuals who have an exact understanding of their tax liability, generating a theoretical and empirical literature which spans a range of issues from the determinants of spending behaviour through decisions on asset allocation to models of tax avoidance and evasion. Such models underpin policy questions such as the optimal marginal rate structure, the balance to be struck between direct and indirect taxation, the case for separating (or integrating) tax and social security schedules, and so on (see, for example, Meade, 1978; Mirrlees *et al*, 2011). However, all reviews of the tax system of this type emphasise the importance of simplicity and transparency in the design of tax systems.

Where tax systems are complex, and not always transparent (for example, where effective marginal rate schedules require computation of the combined schedules from several components of the tax system, tax credits and welfare benefits), the scope for miscalculation is heightened. In such circumstances, individual taxpayers may focus on certain salient features of the tax system or of their perceived taxable income. For example, Chetty, Looney and Croft (2009) show that taxpayers respond to the posted

tax price on purchases rather than the true tax price. In similar vein, Liebman and Zechauser (2004) refer to 'spotlighting', whereby consumers (or taxpayers) respond to certain features of the price or tax structure, whilst ignoring other features. A second common response is to conflate or confuse marginal and average tax rate structures (Bartolome, 1995), which Liebman and Zechauser (2204) refer to as 'ironing'. A natural response to all this, described by Chetty (2011), is for individual taxpayers to approximate in constructing their tax liabilities, changing their behaviour in response to changes in their income or to exogenous changes in the tax structure only when the anticipated gains to doing so exceed a certain threshold. This last argument harks back to an old argument in macroeconomics concerning first order and second order welfare losses in the face of agent inertia (Akerlof and Yellen, 1985).

In this paper, we focus on the extent to which individuals, and households respond to up-front financial incentives in a complex decision-making environment, using tax reliefs for private pension saving in the UK as our case study. We illustrate the analysis with various examples of tax planning in the context of retirement saving which require differing degrees of understanding concerning the tax system.

We first examine the hypothesis that the level of take-up of private pensions increases at the higher rate threshold for income tax, since tax relief on own income for pension contributions is given at the marginal tax rate and should be more attractive to those with higher incomes. By doing this, we are not testing simply whether higher income individuals are more likely to take out pensions (they are) but whether there is a discontinuity in take-up at the higher rate threshold, either because marginal incentives increase at that point or because, more generally, becoming a higher rate taxpayer has an effect on the individual's perception of the need for tax planning. We regard this threshold as a transparent and simple example of a financial incentive.

We then examine a more complex case. Tax relief for pension saving on the partner's income is given at his or her marginal tax rate. Consequently, this might induce households with differential marginal tax rates to adjust joint saving as a form of tax planning. For example a taxpayer facing the basic rate of income tax with a partner facing the higher rate of income tax might, other things being equal, reduce their pension saving as it is optimal for the partner to take a greater share of the

saving. For a number of reasons, we believe that this effect is less likely to show up in the data. In an analysis in a companion paper, we investigate another more complex case, whereby pension saving may be affected by entitlement to tax credits.

2. Existing Literature on tax incentives and retirement saving

The role of tax incentives in encouraging retirement saving has generated a long debate in the United States and a smaller literature in the UK and elsewhere. The debate concerning retirement saving incentives in the United States has focused on two issues: first, the extent to which individuals responded to the introduction of specific retirement saving instruments such as Individual Retirement Accounts (IRAs) and to changes in the thresholds and tax incentives contained therein, and second whether the observed responses were compatible with a 'standard' optimising model of individual life-cycle saving or with models where individual responses had other 'behavioural' interpretations (see, for example, Bernheim and Scholz, 1993; Poterba, 1994; Journal of Economic Perspectives, 1996; Engen, Gales and Uccello, 1999; Attanasio and DeLeire, 2002; Benjamin, 2003; and Chernozhukov and Hansen, 2004).

In understanding this debate, two important factors should be considered. First, different data sets and alternative estimation techniques tend to give different empirical results. Second, and more fundamentally, small responses of *net* private saving to, for example, the introduction of new retirement saving instruments could be compatible with two quite contrary 'stories'. In the first, when there are closely substitutable saving instruments, optimising near-rational households would switch portfolios between close substitutes inducing little overall response in total saving. In the alternative 'story', a lack of response of total saving to changes in the tax treatment of retirement saving assets could be treated as evidence of inertia and myopia (and often has been, by those who do not believe in near-rationality). Obviously testing between these alternative extreme hypotheses would require data with complete household asset portfolios in order to examine substitution behaviour.

Evidence from the UK and elsewhere (and indeed much of the US evidence) has not had sufficiently rich data as to test the 'asset substitution' hypothesis at the level of the individual or household. Evidence from a randomised trial of the Saving Gateway accounts in England by Harvey *et al* (2007) found evidence that contributions to savings accounts tended to represent reshuffling of other financial

assets, in particular for higher income account holders. In the US, the findings of Attanasio and DeLeire (2002) concerning consumer spending and the holding of IRA accounts are consistent with the hypothesis that most funds placed into tax favoured vehicles represented either a reshuffling of existing assets or the diversion of a flow of funds that would have been saved in the absence of such accounts.

Other studies of changing tax treatments and household saving exploit the differential changes in incentives facing households that arise during a typical tax reform. Milligan (2003) utilised changes in tax ceilings in Canada to examine retirement saving decisions (and also discussed the implications of a setting where tax reliefs could be 'rolled over' from one period to another). In the UK, the evidence on up-front tax incentives on retirement saving comes from two policy 'experiments': first, the introduction of Personal Pensions in the late 1980s and changes to the structure of tax reliefs and contribution rebates in subsequent years and, second, the introduction of Stakeholder Pensions in 2002 along with changes in the structure of tax reliefs, especially concerning the introduction of a minimum amount of tax relieved pension saving for low and zero earners.

Taking the latter reform first; Disney, Emmerson and Wakefield (2010) showed that, while the 'targeting' of Stakeholder Pensions on middle-income earners had no effect on take-up of private pensions, the associated changes to the structure of tax reliefs for low and zero earners had significant effects on take-up of private pensions among low earners. Chung *et al* (2008) summarise a series of analyses by researchers at the Institute for Fiscal Studies concerning several changes to the retirement saving regime from the introduction of Personal Pensions in the late 1980s, and show that in all cases there were significant changes in behaviour in the 'predicted' direction, casting doubt on several analyses cited, for example, by the Pensions Commission (2004), that questioned whether individuals responded to retirement saving incentives at all. Some of these changes – such as the change in the terms of the contracted-out rebate for private pensions in 1997 – were quite complex, and yet there was evidence of a changing pattern of take-up in response to reforms even among young employees who might not have been thought particularly responsive to changes in financial incentives between alternative retirement saving vehicles. What these analyses were not able to show, however, was that responses were in any sense 'optimal' given the preferences and lifetime budget constraints facing individuals and households.

Answering such a question would require more 'structural' models with explicit assumptions concerning these parameters as in, for example, Scholz, Seshadri and Khitatrakun (2006).

No analyses, to our knowledge, have concentrated on estimating the effects of marginal and average tax rates in the UK on the propensity of households to take-up private pensions, or to increase the amount of retirement saving (at least, since Barrientos, 1998). This is perhaps a surprising omission, since tax structures and especially marginal rate structures are both a key policy instrument, and have been a central component of the analysis of several other facets of individual and household behaviour in the UK, such as the overall taxable income elasticity (Brewer, Saez and Shephard, 2010), labour supply (Blundell, Duncan and Meghir, 1998) and the extent of charitable giving (Jones and Posnett, 1991a, 1991b; Scharf and Smith, 2009*).* This omission suggests that an extended analysis of the issue of up-front tax incentives and retirement saving is warranted.

3. Incentives arising from tax treatment of pensions

3.1 General principles

In thinking about incentives to save in various assets that are created by the tax system, it is useful to have a framework. There are broadly three stages at which saved funds can be taxed: first, funds may be taxed when income is received (i.e. before or at the point that they are paid into an asset); second, returns (interest, capital gains or dividends) may be taxed as they accrue; and finally, funds may be taxed when they are withdrawn from an asset.[1]

There are two broad possibilities for minimising the distortion to intertemporal spending choices of an agent. First, funds could be taxed when the income is received, the returns should be exempt from tax, and the resulting payoff should be exempt from tax – known as 'TEE' treatment. This is broadly the tax treatment of individual savings accounts (ISAs) in the UK. Second, funds should be exempt from tax when the income is received, the returns should be exempt from tax, but the

[1] Saving(s) could also be subject to other taxes. For example a wealth tax could tax the funds held in a savings vehicle, stamp duties might be levied whenever particular investments are bought and/or sold, and an estate tax might be levied on funds held at death.
[2] A comprehensive summary of the incentives to save in different forms created by the tax

resulting payoff could be taxed at the rate at which income is ordinarily taxed when received − known as 'EET' treatment. The return on the asset is in principle the same in both cases, and is simply equal to the compounded effect of accruing interest; hence, the agent's decision as to whether to consume today or in future (by investing in the asset) is not affected by the tax system. However, the analysis must be modified in practice if the marginal tax rate facing the individual differs between the two periods either through a more complex rate structure or through incomplete indexation to inflation.

The tax treatment of pension saving in the UK is relatively complicated but is most akin to 'EET'. Contributions to pension plans are exempt from income tax (up to a ceiling) and the interest that accrues to pension funds is exempt from tax, while funds that are withdrawn from the pension in retirement are subject to income tax. One complication that arises is that upon retirement individuals are allowed to take up to one-quarter of their total pension savings as a tax-free lump sum.[2] A second complication, mentioned above, is that individuals do not necessarily face the same tax rate on retirement income as they did on the income from which they made their pension contributions (we denote these tax rates as t_r and t_w respectively). Defining as the annual rate of interest on the asset, and as the number of years invested, then the return to £1 of income invested in a pension in the UK is given by:

$$\textbf{Return on pension saving} = \frac{\mathbf{1}(\mathbf{1}+\boldsymbol{r})^{\boldsymbol{x}}(\mathbf{0.25}+\mathbf{0.75}(\mathbf{1}-\boldsymbol{t_r}))}{(\mathbf{1}-\boldsymbol{t_w})}$$

Pension saving is therefore typically tax favoured relative to TEE assets – primarily due to the lump sum, but also particularly if the tax rate faced in retirement is lower than that during working life. Such tax favoured treatment is designed to entice individuals to save in pensions despite the illiquidity of this form of saving and (at least until recently) the commitment to buy an annuity in future with any accrued funds. One way of quantifying the relative tax treatment of pension saving is to consider what percentage greater return an individual would get from saving in a 'neutral' TEE asset such as an ISA:.

[2] A comprehensive summary of the incentives to save in different forms created by the tax and benefit system can be found in Wakefield (2009).

$$Pension\ premium = \left(\frac{Return\ on\ pension\ saving}{Return\ on\ TEE} - 1\right) \times 100$$

Table 1 describes the gain from pension saving relative to investing in an ISA, according to the working age and retirement tax situation of individuals, which we term the 'pension premium'. The situation is slightly more complicated than that described above because it is not just the tax system that matters but also the tax credit system, since pension contributions are deducted before income is assessed for some tax credits in the UK. The marginal tax rates t_r and t_w should therefore be interpreted as *effective* marginal tax rates, taking account not just the proportion of each additional £1 that has to be paid in tax, but also the proportion of that £1 that is lost from the reduction of means-tested tax credits.

Table 1: Pension Premium created by the tax and tax credit system, by working age and retirement situation

Working age situation		**Retirement situation**	
		Basic rate income taxpayer	Higher rate income taxpayer
Tax / tax credit situation	t_w	t_r = 0.20	t_r = 0.40
WTC taper plus basic rate income tax	0.590	107.3	70.7
Basic rate income tax	0.200	6.3	-12.5
CTC taper plus basic rate income tax	0.267	15.9	-4.5
Higher rate income tax	0.400	41.7	16.7
CTC taper plus higher rate income tax	0.467	59.4	31.3

Notes: For the 2008–09 tax system. Figures are for individual (rather than employer) contributions, since the latter also receive relief from employee and employer National Insurance Contributions.

By way of an example, a working age basic rate tax payer who receives the Working Tax Credit (WTC), but has an income high enough such that he or she is on the taper where WTC starts to be withdrawn as income increases would have to pay 20p in income tax, and would lose 39p from reduced WTC, for every £1 earned, thereby creating an effective marginal tax rate of 0.59. Since pension contributions are exempt from income tax and are not counted as income for the purposes of assessing eligibility to tax credits, an individual could contribute £1 to a pension and only lose 41p of his or her disposable income. In addition individuals should consider the

impact of the benefit system on their marginal tax rate in retirement. In particular basic rate income tax payers who expect to be on the taper of the Pension Credit in retirement could face a similar disincentive to save in a private pension as those basic rate income taxpayers who expect to be higher rate income taxpayers in retirement (since the Pension Credit is, for those in receipt of the savings credit, withdrawn at an effective rate of 40% which is the same as the higher rate of income tax).

For the purposes of this work we focus on the 'up-front' incentive to save, the effective tax relief given to pension contributions (t_w), as shown in the 'Working age situation' component of Table 1, rather than the full premium from investing in pensions as opposed to a 'neutral' TEE asset. There are two reasons for so doing. First, it is a useful simplification in this context since, for any given retirement tax situation, an increase in the up-front incentive to save also implies an increase in the pension premium. Comparing individuals in terms of their upfront incentive is therefore equivalent to comparing people based on their pension premium so long as it is just the ordering of who has a greater or lesser incentive to save that is of interest (as opposed to the precise magnitude of the difference in the incentives between individuals). Second, from a behavioural standpoint it might be argued that individuals are better able to understand the upfront incentives to save, and are therefore more likely to respond to these than the full pension premium which requires much greater understanding of the tax system and computational ability. In this paper, we also only focus on the impact of the higher rate threshold (HRT) of the income tax system, but in different household settings.

3.2 Impact of tax structure on retirement saving

In the standard optimisation model of retirement saving, tax regimes alter the incentives for individuals to shift spending between periods, and individuals are assumed to understand fully these incentives. If this is the case, then individuals with a higher up-front incentive to save than otherwise similar individuals would be expected both to be more likely to save in a pension, and on average save larger amounts as income increases. It is however well-established that individuals' propensities to engage in retirement saving at all, as well as levels of saving increase, as incomes increase. What we are looking for therefore is a discontinuity in saving behaviour that occurs at rate thresholds in the tax structure: specifically here at the

higher rate threshold for income tax, where the effective value of tax relief on pension contributions increases from the standard rate of between 20% and 22% over our period to 40%. We should expect to find an increase in the probability and/or level of retirement saving at this threshold over and above the effect of higher incomes.

The discontinuity argument allows us to handle the fact that the propensity to engage in retirement saving increases with income so long as the discontinuity in the marginal structure of tax rates does not coincide with any discontinuity in the underlying relationship between income and retirement saving behaviour. This seems a fairly safe assumption – it is unlikely that subjective discount rates or credit constraints, for example, happen to be discontinuous at the same point as the tax thresholds.

On the other hand, evidence of a jump in the probability of retirement saving at the point where the individual becomes a higher rate taxpayer does not of itself confirm that individuals were simply responding to the change in the upfront incentives to save arising from the tax system. Becoming a higher rate taxpayer may have an effect on the individual's perception of the need for tax planning, and therefore make them more likely to make use of tax advantaged methods of saving even if the relative tax advantage of saving in these assets has increased only marginally at the point of discontinuity. Alternatively becoming a higher rate taxpayer may act as a signal that an individual has reached a stage in their life when they should be saving more, or saving specifically for retirement. Again this could act to make individuals increase their pension saving, even if the financial incentive to do so had not changed.

We should also consider the possibility that the individual's perception of exactly when they reach this threshold of the higher marginal tax rate may be masked if they do not have full information on their income position. Individuals may only perceive that they face a higher marginal rate either when it is obvious from their pay packet (for example, where earnings form by far the bulk of their income, and their annual earnings now exceed the threshold), or at the point at which they (or their accountant, if one is employed by the individual) are calculating their overall tax liability (for example, just in advance of submitting their tax return), or *ex post* once their tax coding and tax calculation reveals that they are now liable for a higher rate of tax over part of their income. All of these 'triggers' may induce an individual to engage in tax

planning, but they may occur sometime after the individual actually faces a higher marginal rate. Alternatively, an 'approximating' individual may not have done the full tax liability calculation but realise that, combining their earnings and an estimate of their allowances and other taxable income (for example from other forms of saving), they are coming close to the overall income level at which they will be liable for a higher rate of tax on the margin. This will heighten the individual's perception of the need for tax planning as their income approaches the threshold, before they are actually liable for the higher marginal rate of tax.

Finally, it is important to distinguish between the upfront incentive to make a particular £ of pension contributions and the average upfront incentive across the pension contributions made. For example, someone who has an income £1 over the higher rate threshold for income tax can get tax relief on their first £1 of pension contributions at 40%, but on all subsequent £s the rate of relief would only be 20%. The average upfront incentive across all the individual's pension contributions would be between 20% and 40% (depending on how many £s are contributed), and would increase as income increased further above the higher rate threshold. Here, we calculate for each individual the tax relief that they would get on their first £1 of pension contributions. The implication of this is that the discontinuities in pensions saving behaviour that one might initially expect at rate thresholds in the tax structure could be muted. A marked discontinuity in behaviour at the rate threshold might be most expected for individuals who are not yet members of a pension and react to their upfront incentive to save their first £1 in a pension. In terms of the *level* of pension contributions, these are probably less likely to respond to a change in the upfront incentive to contribute the first £1, but are more likely to respond to an increase in the average incentive to save across pension contributions.

The test of the hypothesis that individuals with a higher upfront incentive to save in a pension are more likely to save, and on average save larger amounts, therefore has two aspects: first, whether there is a discontinuity in pension saving behaviour at the point at which the upfront incentive changes – this would capture effects associated with a change in the upfront incentive to make the first £ of contributions. Second, whether there is any kink in the relationship between pensions saving behaviour and income that occurs at the point at which the upfront incentive changes – this would

capture effects associated with the more gradual change in the average upfront incentive to save.

To investigate the potential difference between responses to simple tax incentives and responses to more complex tax incentives we examine two potential instances of where upfront incentives to save in a pension differ across incomes. The first is a straightforward enhanced incentive to increase pension saving when an individual hits the higher marginal rate of tax on income. The second, potentially more complex, example is pension saving within couples. Taxation in the UK occurs at the individual level, and so couples may adjust joint saving as a form of tax planning. For example, a basic rate income taxpayer with a higher rate taxpayer partner might, other things being equal, reduce their retirement saving as it is optimal for the partner to take a greater share of the pension saving.

There are, however, a number of reasons why we might not expect households in the second case to respond to these incentives, aside from simply the additional complexity involved. First, the availability of particular types of pensions to the household will matter. If an individual in the household has access to a 'good' employer-provided pension, in particular one with a relatively generous treatment of the spouse, then this will affect whether the tax treatment of pension contributions gives their partner a greater incentive to save.[3] Second, within-family separation risk is likely to be important. If the household thinks there is a chance that they will no longer be together in retirement then they may want to save for retirement independently. Both of these reasons might make it less likely that there is a detectable effect of upfront tax incentives on the likelihood of being a member of a pension; however we might still expect an effect on the magnitude of pension saving in terms of the employee contributions made.

3.3. Methods and Data

3.3.1 A regression discontinuity approach

To investigate whether upfront incentives to save in a pension affect behaviour we use a regression discontinuity (RD) approach to analyse the sharp change in

[3] An example would be where one household member has a high paying private sector job but the second member has a lower paying public sector job but with a defined benefit pension.

incentives that occurs when an individual's income puts them over the Higher Rate Threshold (HRT) – at this point the tax relief on pension contributions increases from 22% to 40% (20% to 40% for the last year of our analysis). We first focus on the individual's own income and then on that of their partner. Individuals observed just below and just above the HRT should be very similar in terms of both their observed and unobserved characteristics. Under the assumption that pension saving increases smoothly with taxable income, any unusual jump in pension saving at that point can therefore be reasonably attributed to being an effect of hitting the higher rate threshold. Appendix A shows a range of graphs illustrating how various observed characteristics vary across the income distribution in the range of income £10,000 either side of the HRT. None of the characteristics considered exhibit a discontinuity at the HRT, which is crucial if any discontinuity in pension saving at that point is to be claimed to be an effect of the HRT itself.

To operationalise the RD approach we utilise both non-parametric and parametric methods. The non-parametric method plots lowess curves of the outcome of interest by distance of taxable income from the HRT for individuals above and below the HRT separately. Lowess curves are created using locally weighted scatterplot smoothing – a methodology introduced by Cleveland (1979) and further developed by Cleveland and Devlin (1988).[4] If there were no effect of the HRT on pension coverage then the two curves, one for those above the HRT and one for those below the HRT, would be expected to line up neatly at the HRT.

The parametric method is a standard one, where we estimate the models:

(A) $$Y_i = \alpha + \beta_1(X_i - c) + \tau I_i + \gamma_1(X_i - c)I_i + \varepsilon_i$$

(B) $$Y_i = \alpha + \beta_1(X_i - c) + \tau I_i + \gamma_1(X_i - c)I_i + \beta_2(X_i - c)^2 + \gamma_2(X_i - c)^2 I_i + \varepsilon_i$$

where $(X_i - c)$ is distance of the individual's taxable income from the HRT and I_i is an indicator of being above the HRT. At the HRT $(X_i - c) = 0$, and so τ is the

[4] To create the smooth lowess curve, for each distance of income from the higher rate threshold $(X_i - c)$, a weighted linear regression of the outcome of interest Y_i on $(X_i - c)$ is conducted on those observations j with $(X_j - c)$ close to $(X_i - c)$. The regression estimate is then used to predict Y_i at the point $(X_i - c)$, and the lowess curve is constructed as the plot of these predicted Y_i. In this work a tricube weighting function is used in the local weighted linear regressions.

discontinuity coefficient of interest – the estimate of the effect of hitting the HRT itself, over and above having the level of income associated with it. The coefficients γ_1 and γ_2 (where applicable) are also of interest − they indicate whether the underlying relationship between individual taxable income and the outcome Y_i is different above the higher rate threshold to below the higher rate threshold (which could indicate an effect of changes in the average upfront incentive to save).

These equations are estimated for individuals 'close' to the HRT, in other words, within a distance $\frac{h}{2}$ of the HRT (i.e. $i: |X_i - c| < \frac{h}{2}$). The size of the bandwidth *h* trades-off between improving the similarity of individuals in the sample (as the bandwidth gets smaller) and improving the sample size (as the bandwidth gets bigger). Model (B), which controls for the squared distance of income from the higher rate threshold, is used in preference over model (A) when the bandwidth is large enough such that it is felt that a non-linear relationship between income and the outcome of interest might be pertinent.

In addition to the two equations (A) and (B), we also estimate similar models that also control for individual characteristics:

(C) $$Y_i = \alpha + \beta_1(X_i - c) + \tau I_i + \gamma_1(X_i - c)I_i + \beta_2(X_i - c)^2 + \gamma_2(X_i - c)^2 I_i + \vartheta Z_i + \varepsilon_i$$

(D) $$Y_i = \alpha + \beta_1(X_i - c) + \tau I_i + \gamma_1(X_i - c)I_i + \beta_2(X_i - c)^2 + \gamma_2(X_i - c)^2 I_i + \vartheta Z_i + \varepsilon_i$$

Where Z_i is a vector of covariates, including: sex, age, age squared, education, marital status, the number of children, whether the individual has a child aged under 5, housing tenure, hours worked, industry worked in, whether the individual worked in the 'public sector', region of residence, partner characteristics if applicable (age difference, age difference squared, education, work, hours, industry, 'public sector') and time dummies.[5] Since X_i controls for the individuals' taxable income we also include controls for other gross income of the family (and its square) in order to control for a broader measure of the income of the family. This will include their

[5] The data used in this work does not identify whether each employee works in the public or private sector. A crude indicator is constructed based on industry of work, which counts as working in the 'public sector' those employed in: public administration and defence, compulsory social security, education, health and social work.

partner's income (where relevant) and also any non-taxable income that they receive (such as child benefit).

If the RD approach is fully justified then the individuals in the sample just below the HRT should be similar in their observed and unobserved characteristics to those just above the HRT. The inclusion of these additional controls for individual characteristics should therefore not affect the coefficient τ that measures the effect of hitting the HRT. However since in practice we may be looking at individuals in a relatively large income band around the HRT, including these covariates can help to eliminate some bias in the estimates of the effect of the HRT. It could also improve the precision of the estimate of the HRT if the covariates are correlated with the outcomes of interest.

When we look at our second example, pension saving in couples, the variables in the regression analysis are slightly different. The indicator variable I_i indicates whether the individual's partner has a taxable income greater than the HRT, and $(X_i - c)$ is the distance of the partner's income from the higher rate threshold. The RD approach is still justified when considering individuals with a partner just above or just below the HRT – there is no reason that these individuals should differ in their average characteristics – but there will be more heterogeneity among these individuals than when comparing individuals who are themselves just above or just below the HRT. The inclusion of controls for individual (and partner) characteristics is potentially therefore much more important in the analysis of responses within couples. The covariates included in the analysis of couples are the same as those for the individual analysis, with the exception that 'other' income now relates to gross family income less the partner's taxable income (i.e. it includes the individuals own income).

3.3.2 Data

The data are taken from the Family Resources Survey (FRS). This is an annual cross-sectional survey that records detailed income, socio-economic and demographic information for a sample of nearly 30,000 households in the UK.[6] In order to get a large sample of individuals with an income 'close' to the HRT we pool 7 years of the FRS: 2000–01 to 2008–09. We convert all monetary amounts into real terms by

[6] It covers the UK population from 2002–03 onwards and the Great British population prior to that.

adjusting them using the change in the retail price index (RPI) between the date of interview and December 2009. The sample is restricted to employees aged between 22 and 59 (inclusive), who do not have any self-employment income or any pension income. These latter restrictions are imposed as the individuals excluded as a result are likely to have atypical pensions saving behaviour – those with self employment income may save for retirement through their business rather than through pensions, and those who are already receiving a pension income are unlikely to be still saving for retirement. This results in a sample of 196,249 individuals before any restrictions by income are imposed.

To calculate for each individual the incentive to save in a pension, we calculate taxable income − the income which is compared with the income tax system thresholds to work out an individual's tax liability – for each individual.[7] Importantly, since pension contributions are tax deductible, taxable income excludes from income any pension contributions. To avoid the obvious endogeneity of taxable income generated by existing pension contributions we add back into taxable income any pension contributions that an individual made.[8] Hereafter, taxable income including pension contributions will be referred to simply as 'income', and comparison of this income with the thresholds in the tax system will indicate the individual's up-front incentive to save the first £1 in a pension.

Table 2 describes the average characteristics of individuals below and above the HRT across the whole sample, and across the sample of individuals within £10,000 either side of the HRT. Not surprisingly, individuals calculated to be above the HRT of the income tax system are on average older, better educated, and are more likely to be men than women. Those individuals within £10,000 either side of the HRT are primarily home owners (90% either own their home outright, are buying it with the

[7] Taxable income is calculated as the sum of earnings from employment or self employment, pension income, Job Seekers Allowance, widow's benefits, taxable interest, rental income from property and border or lodger payments (in excess of the annual allowance), less any pension contributions.

[8] We ignore the issue of whether individuals manipulate their taxable income to avoid higher marginal tax rates by changing their wage rate or their hours – in other words, we treat income other than pension contributions as exogenous, which is required for the validity of the RD approach.

help of a mortgage, part own their home or are in a shared ownership scheme), and are therefore unlikely to be credit constrained (which would likely impact on pensions saving behaviour).

The retirement saving outcomes considered here are membership of a private pension scheme (whether provided through the workplace or individually arranged) and employee pension contributions. An individual is counted as being a current member of a private pension if they report having made employee contributions to a personal or employer-provided pension in the past year.

Table 2: Average characteristics of individuals above and below the HRT (Higher Rate Threshold of income tax)

	All			h = £10,000		
	All	<HRT	>=HRT	All	<HRT	>=HRT
Proportion male	48.9%	44.8%	77.7%	68.1%	66.0%	72.9%
Average age	39.9	39.6	42.1	41.4	41.2	41.8
Proportion married	60.2%	58.5%	72.4%	66.5%	65.2%	69.3%
Proportion cohabiting	14.0%	14.4%	11.1%	12.9%	13.2%	12.1%
Average number of kids	0.7	0.7	0.8	0.7	0.7	0.7
Proportion with child <1	4.0%	3.9%	4.5%	4.1%	4.0%	4.4%
Proportion with child <5	16.9%	16.5%	19.8%	17.2%	16.9%	18.0%
Proportion with low education	49.6%	53.2%	23.6%	34.0%	36.7%	27.8%
Proportion with high education	27.6%	23.7%	54.9%	43.3%	40.3%	50.2%
Proportion owner-occupier	79.6%	77.9%	91.9%	90.0%	89.4%	91.3%
Proportion 'public sector'	32.2%	33.2%	24.9%	35.9%	37.5%	32.5%
Mean hours	38.7	37.3	48.2	45.6	45.1	46.8
Mean income tax income (£'000)	25.474	19.507	67.784	39.417	36.368	46.295
Mean household 'other' income (£'000)	18.685	18.260	21.695	19.404	18.874	20.601
Proportion contributing to pension	56.9%	53.9%	77.7%	77.7%	77.1%	79.0%
Sample	*196,249*	*171,994*	*24,255*	*34,697*	*24,040*	*10,657*

Notes: Low education refers to those who left full time education at age 16 or below while high education refers to those who left full time education at age 19 or above. Owner-occupier refers to those who either own their home, part-own their home, are buying it with the help of a mortgage or are in a shared ownership scheme. 'Public sector' refers to those who are employed in: public administration and defence, compulsory social security, education, health and social work.

Source: Family Resources Survey pooled 2000–01 to 2008–09, earlier years' incomes revalued by RPI.

By way of background, we first demonstrate that individuals with higher incomes (and therefore, broadly, higher average income tax rates on income) are more likely to engage in retirement saving through private pensions. This is demonstrated for our sample of employees in the pooled FRS in Figure 1, which shows average pension

membership by income. Only around one quarter of individuals with the lowest incomes are active members of a pension scheme. This fraction increases with income until it plateaus at just under 80% coverage for those with incomes of greater than £35,000 p.a. (2009 prices). Table 2 indicates that across our whole sample of individuals 60% are members of a pension, while among those who have an income within £10,000 either side of the HRT nearly 80% are members of a pension. Figure 2 shows that mean and median employee pension contributions also increase with income.

Figure 1: Pension membership by income

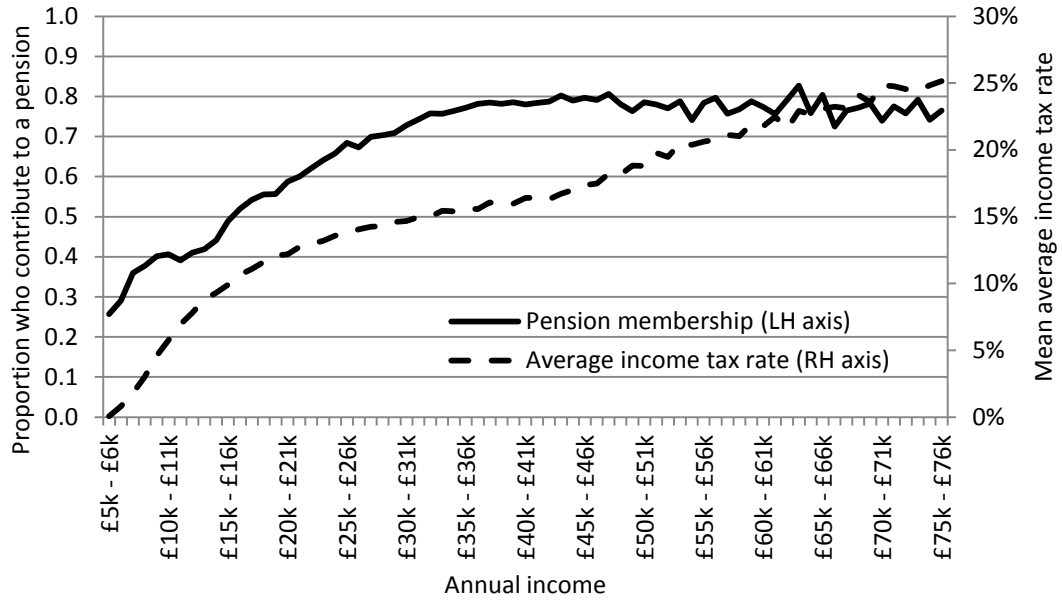

Source: Family Resources Survey pooled 2000–01 to 2008–09, earlier years revalued by RPI.
Notes: Solid line plot shows the proportion of individuals within each £1,000 band of income contributing to a pension. Dashed line plot shows the mean average income tax rate across individuals within each £1,000 band of income.

Figure 2: Employee pension contributions (of pension members) by income

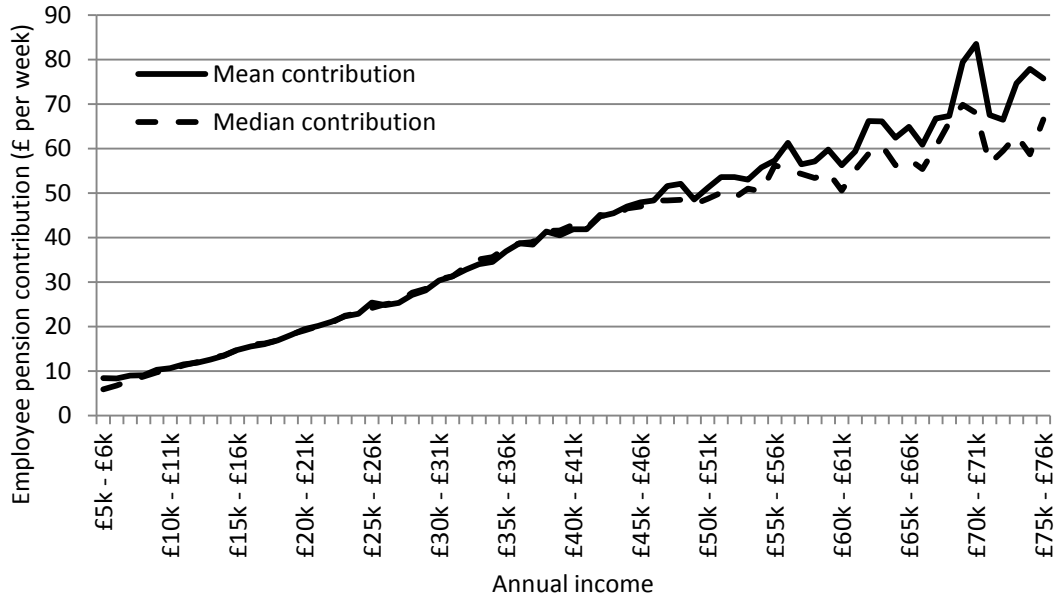

Notes: Line plots show the mean and median pension contribution of pension members within each £1,000 band of income.
Source: Family Resources Survey pooled 2000–01 to 2008–09, earlier years revalued by RPI.

4. Results

4.1. The effect of an individual's own income on the incentive to save

Coverage effects

Figure 3 plots the proportion of individuals contributing to a pension, by £500 bands of income, for individuals who have an income between £10,000 below the HRT and £10,000 above the HRT. The dashed lines are lowess curves, created using locally weighted scatterplot smoothing for individuals above and below the HRT separately. If there were no effect of the HRT on pension coverage then these two curves would be expected to line up neatly at the HRT. In fact there is around a 1.5 percentage point jump up in the smoothed line. This implies that there is a discontinuity, albeit a small one, in the relationship between income and the probability of contributing to a pension that occurs at the HRT.

Figure 3: Proportion of individuals contributing to a pension, by income

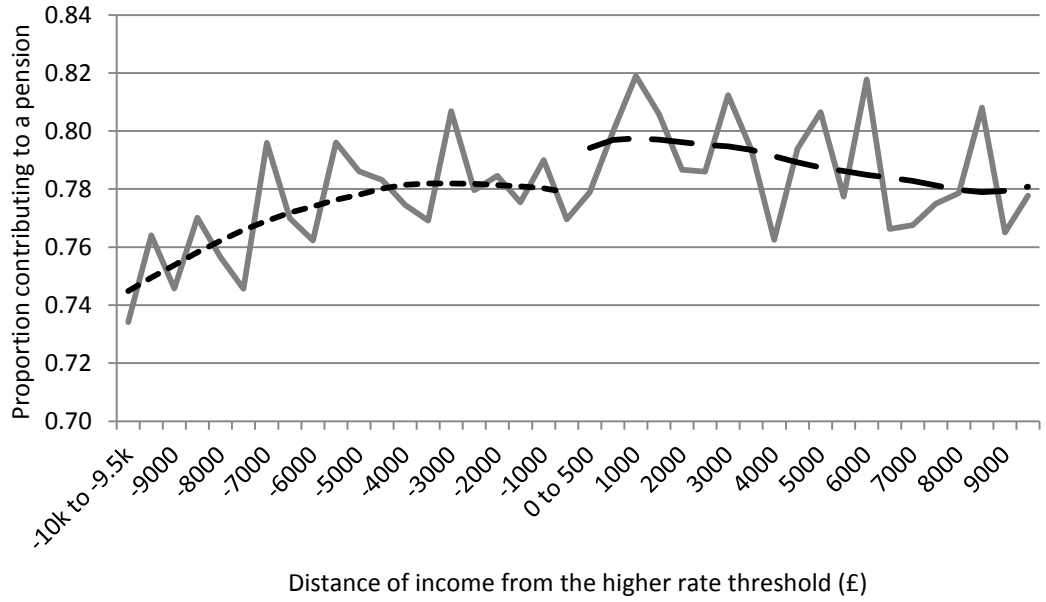

Notes to Figure 3 and subsequent plots: the line plot shows the proportion of individuals within each £500 band of income contributing to a pension. The dashed lines are lowess curves, created using locally weighted scatterplot smoothing for individuals above and below the HRT separately.

The discontinuity in take-up and proximity to the HRT threshold as specified in text equations (A) to (D) in Section 3.3.1 is examined parametrically using Ordinary Least Squares. Table 3 shows the regression estimates of τ, γ_1 and γ_2 (where applicable), the impact of the HRT on private pension coverage, for three different sample bandwidths. In all of the specifications shown hitting the HRT led to a boost in pension membership of about 2 percentage points – a similar magnitude to the effect that is implied by Figure 3 – but the effect is not quite statistically different from zero at standard significance levels. The coefficients γ_1 and γ_2 are also small and insignificant, implying that the HRT did not change the relationship between pension coverage and income in this income range. Controlling for individual observed characteristics did not make an important difference to the size of the coefficients, as would be expected if the RD approach were justified.

As argued in Section 3.2 above, although an individual's 'true' up-front incentive to save in the form of a pension depends on how their total taxable income compares to the thresholds of the tax and tax credit system, people may not fully understand their full taxable income and therefore their exact tax position. Individuals are likely to approximate, but individuals will likely to be more aware of their earnings from employment. Repeating the analysis from Figure 3 and Table 3, therefore, we define

an individual as being above or below the HRT on the basis of their earnings only. Figure 4 and Table 4 then show that having earnings just greater than the HRT as opposed to just below is found to lead to just over a 3 percentage point increase in pension coverage, which is statistically significant from zero at conventional levels. This suggests that total earnings may be more salient for individuals than total taxable income. For the remainder of the results presented in this paper individuals are defined in terms of the distance of their total earnings from the higher rate threshold.[9]

Table 3: Effect on pension coverage of having income just greater than the HRT as opposed to just lower

	h = £20,000		h = £10,000		h = £4,000	
τ	0.023	0.019	0.021	0.020	0.014	0.014
	(0.014)	(0.013)	(0.013)	(0.013)	(0.021)	(0.020)
γ_1 *	0.004	0.001	-0.002	-0.003	0.018	0.015
	(0.007)	(0.006)	(0.005)	(0.004)	(0.018)	(0.018)
γ_2 *	0.001	0.001				
	(0.001)	(0.001)				
β_1 *	-0.005	-0.003	-0.001	-0.001	-0.006	-0.005
	(0.004)	(0.004)	(0.003)	(0.003)	(0.012)	(0.012)
β_2 *	-0.001*	-0.001				
	(0.000)	(0.000)				
Characteristics	×	✓	×	✓	×	✓
Income specification	Quadratic	Quadratic	Linear	Linear	Linear	Linear
Equation	B	D	A	C	A	C
N	34,697	34,108	16,278	16,001	6,339	6,213

Notes: * Distance from the HRT, $(X_i - c)$ in equations A to D, is in £,000s. Standard errors in parentheeses. ***, **, * indicate statistical significance at the 1%, 5% and 10% levels respectively.

[9] For most of the individuals in the sample, income and earnings are very similar – 90% of the sample have annual income and earnings within £1,000 of each other. Only 8.2% of the sample have investment income of more than £1,000, while only 2.5% of the sample have rental income from property of more than £1,000.

Figure 4: Proportion of individuals contributing to a pension, by earnings

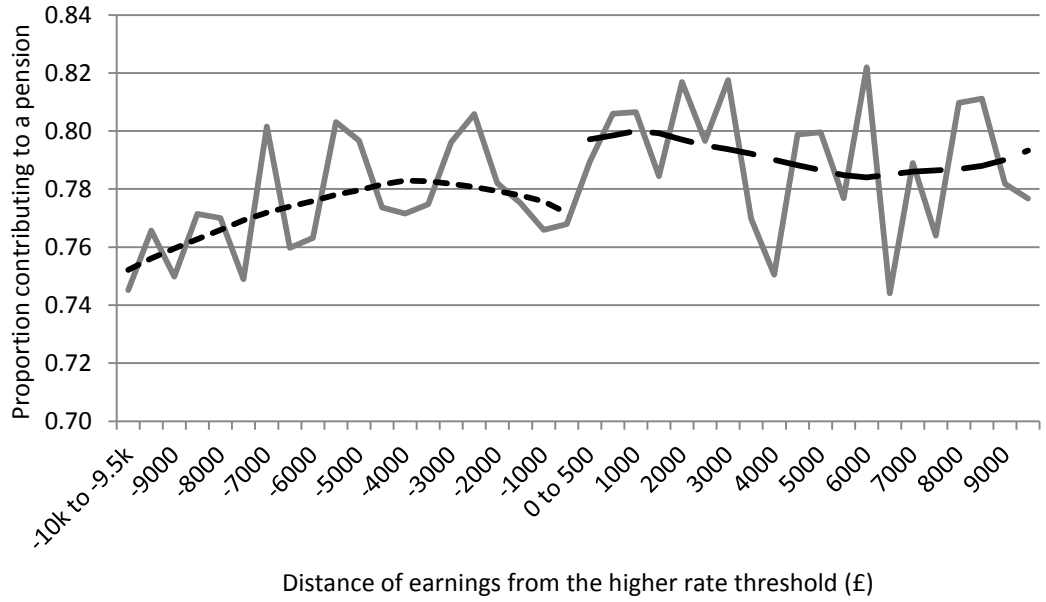

Table 4: Effect on pension coverage of having earnings just greater than the HRT as opposed to just lower

	h = £20,000		h = £10,000		h = £4,000	
τ	0.034**	0.030**	0.032**	0.029**	0.036*	0.033
	(0.014)	(0.014)	(0.013)	(0.013)	(0.021)	(0.020)
γ_1 *	0.001	0.000	-0.002	-0.002	-0.002	-0.008
	(0.007)	(0.007)	(0.005)	(0.005)	(0.019)	(0.018)
γ_2 *	-0.001**	-0.001*				
	(0.000)	(0.000)				
β_1 *	-0.007	-0.005	-0.003	-0.002	-0.005	-0.002
	(0.004)	(0.004)	(0.003)	(0.003)	(0.012)	(0.012)
β_2 *	0.001*	0.001*				
	(0.001)	(0.001)				
Characteristics	×	✓	×	✓	×	✓
Income specification	Quadratic	Quadratic	Linear	Linear	Linear	Linear
Equation	B	D	A	C	A	C
N	34,015	33,432	15,893	15,617	6,185	6,061

Notes: * Distance from the HRT, $(X_i - c)$ in equations A to D, is in £,000s. Standard errors in parentheses. ***, **, * indicate significance at the 1%, 5% and 10% levels respectively.

An important consideration when using regression discontinuity analysis is whether the underlying relationship between the outcome of interest and the characteristic exhibiting the discontinuity is truly smooth. If the relationship between pension coverage and income were not smooth, then a discontinuity in pension coverage that just happened to occur at the HRT may be spuriously associated with the HRT. In order to test the validity of the discontinuity identified at the HRT, the

analysis of Tables 3 and 4 is repeated using a four placebo thresholds instead of the true HRT. The results are reported in Table A1 in the appendix, and indicate that pension coverage does not exhibit a significant discontinuity at other points in the income distribution.

Finally, the RD approach is based on the assumption that individuals just below and just above the HRT are similar in terms of their observed and unobserved characteristics. While this is true on average there is still a large amount of heterogeneity among individuals who are just below the HRT and among individuals who are just above the HRT. In order to investigate whether reaching the HRT affects different types of individuals differently, the analysis is run separately for different subgroups.[10] The results of this are shown in Table 5.

Table 5: Effect on pension coverage of having earnings just greater than the HRT as opposed to just lower, by subgroups

	h = £10,000 (model D)			h = £5,000 (model C)		
	τ	Std. err	N	τ	Std. err	N
All	0.030**	0.014	33,432	0.029**	0.013	15,617
Male	0.020	0.017	22,913	0.023	0.016	10,911
Female	0.048**	0.023	10,519	0.034	0.021	4,706
Single	0.038	0.031	6,953	0.030	0.029	3,135
Couple	0.027*	0.015	26,479	0.029*	0.014	12,482
Aged 22−39	0.032	0.023	14,867	0.038*	0.021	6,728
Aged 40−59	0.026	0.017	18,565	0.021	0.016	8,889
Public sector	0.023	0.018	12,107	0.014	0.016	5,681
Private sector	0.033*	0.019	21,325	0.035**	0.018	9,936
Low education	0.027	0.027	11,282	0.022	0.025	4,915
Mid education	0.031	0.030	7,544	0.023	0.028	3,477
High education	0.028	0.018	14,606	0.030*	0.017	7,225

Notes: Regressions include controls for individual characteristics but the τ coefficients are robust to the exclusion of these. Standard errors in parentheses. ***, **, * indicate significance at the 1%, 5% and 10% levels respectively.

Positive effects were found for all the subgroups considered, and a significant effect is found for women, those in couples, the younger half of the sample (those

[10] The effect of the HRT, and the associated increase in the upfront incentive to save in a pension, on pension coverage might be expected to differ not just according to individual characteristics, but also by the type of pension. This is not considered here as for most of the period in question, pensions types are not suitably distinguished in the FRS.

aged 22-39), those working in the 'private sector' and those with high levels of education[11]. The effect of hitting the HRT could be driven by the responsiveness of those who work in the private sector, for whom the HRT is associated with a 3.5 percentage point increase in pension coverage rates. This could plausibly be explained by the greater availability of defined benefit pension schemes to public sector workers in the UK, which individuals would typically be expected to join as soon as they could afford to do so (since retirement benefits are tenure related). We do not have enough power in the tests however to say that any of these subgroups were significantly more or less responsive than other groups.

As mentioned previously, our approach does not separately identify the effect of the increase in the upfront incentive to save that occurs at the HRT, from other more general signalling effects that might be thought to occur when the individual becomes liable for the higher rate of income tax. That pension membership has a significant discontinuity associated with the HRT does not therefore necessarily imply that individuals are responding to the increase in the upfront incentive to save in a pension that occurs at that point. We therefore consider the proportion of individuals just above the HRT and just below the HRT who hold Individual Saving Accounts (ISAs). ISAs are tax favoured relative to many other savings products which have 'TTE' treatment, but the relative incentive to save in an ISA relative to, for example standard savings accounts, only increases slightly when an individual reaches the HRT.[12] Moreover the relative incentive to save in an ISA compared to a private pension falls at the HRT, which might suggest that ISA coverage would decline at this point. However if reaching the HRT acts as a signal that individuals should be take a greater active interest in their financial choices, then we would expect the proportion holding ISAs to increase in a similar way to the proportion holding pensions.

Figure 5 shows that there is in fact no discontinuity in the proportion of individuals holding ISAs at the HRT. This could imply that the effect of the HRT on pension

[11] Individuals are defined as having low levels of education if they left full-time education at or below age 16, mid levels of education if they left full-time education above age 16 but below age 19, and high levels of education if they left full-time education at or above age 19.

[12] The absolute return to saving in an ISA does not change when an individual crosses the HRT, but the relative incentive to save in an ISA does increase since the return to saving in TTE assets (such as standard savings accounts) decreases (due to the higher tax payable on interest income).

membership is due to the increased upfront incentive to save in a pension, rather than a general signalling effect of attaining the HRT of the need for tax planning more generally. However it still cannot be rejected that the HRT acts as a signal of the specific need to save in a pension for retirement, rather than a more general need to save in a more tax efficient way.

Figure 5: Proportion of individuals with an ISA, by earnings

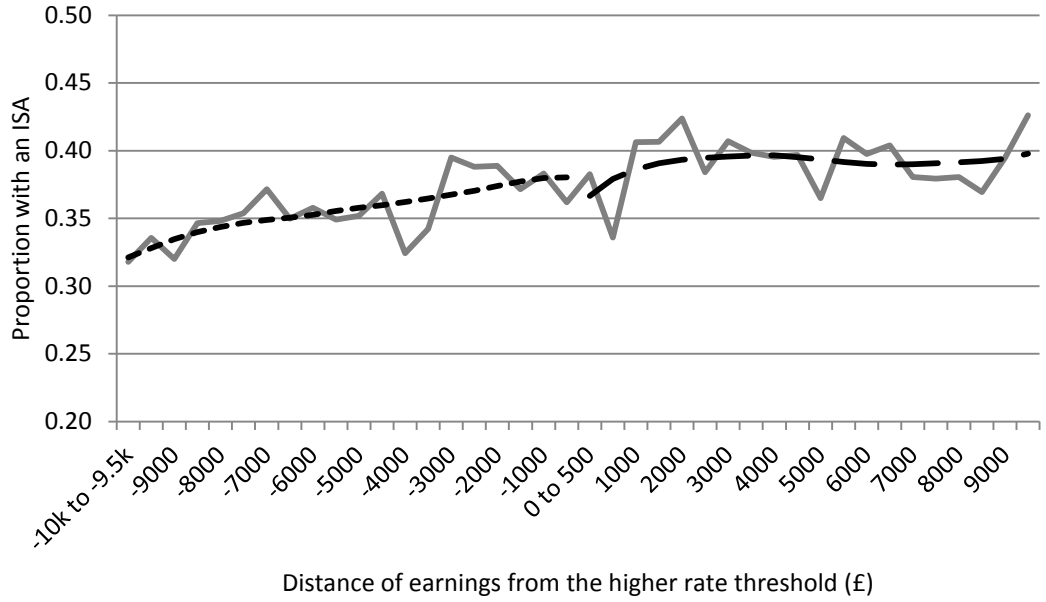

Contribution levels

We now investigate whether the increase in the upfront incentive to save in a pension that occurs when individuals cross the HRT results in an increase in pension contributions. As discussed in section 3.2, the effect on pension contributions is likely to be one that comes through the (more gradual) change in the average incentive to save, and therefore appear as a kink in the relationship between contribution levels and earnings at the HRT, rather than a level discontinuity in pension contributions at that point. However this object of interest cannot be identified from the pooled cross sectional data used in this work. We observe the pension contributions of employees with earnings greater than the HRT, and can compare these to the pension contributions of those with earnings below the HRT, but this confounds two effects: the effect of the HRT on the pension contributions of those who would also have been contributing to a pension if their earnings were below the HRT (the object of interest here), and the contributions behaviour of those who are brought into pension saving as a result of having earnings greater than the HRT (the discontinuity in pension

membership found in the previous section). If these individuals contribute the same amount on average as someone on the same level of earnings who would have contributed to a pension had their earnings been lower than the HRT then there would be no problem.

A priori it is unclear what the expected effect of the HRT on the relationship between pension contributions and earnings should be. Arguments can be made that at or above the HRT threshold, individuals are likely to contribute either lower amounts (they were only induced to save by the higher upfront incentive, which could imply lower preferences for saving) or higher amounts (they may have been expecting to become higher rate taxpayers and therefore have been holding off saving until they could get greater tax relief). Looking at average contributions across existing *pension members*, a positive or negative discontinuity at the HRT and a positive or negative effect on the rate at which contributions change with earnings could be expected. Looking at average contributions across the *whole sample*, counting individuals who are not members of a pension as having zero contributions, a positive discontinuity at the HRT would be expected (since those brought into pension saving will certainly have a positive level of contributions), but the expected effect on the rate of contributions is still unclear.

Figure 6 plots the average weekly contribution to a pension, by £500 bands of earnings, for individuals have earnings between £10,000 below the HRT and £10,000 above the HRT. These averages are computed both for just those who are saving in a private pension, and then also across all individuals. Average weekly pension contributions are remarkably smooth, and increase linearly with earnings at a rate of approximately an extra £0.88 of contributions per £1,000 of additional annual earnings. Looking first at the average level of contributions across just those in a pension, there appears to be no discontinuity at the higher rate threshold. This is confirmed by the parametric results reported in the first three columns of Table 6. There is also no consistent or significant evidence that the HRT is associated in a change in the rate at which pension contributions increase with earnings.

Figure 6: Average weekly pension contributions, by earnings

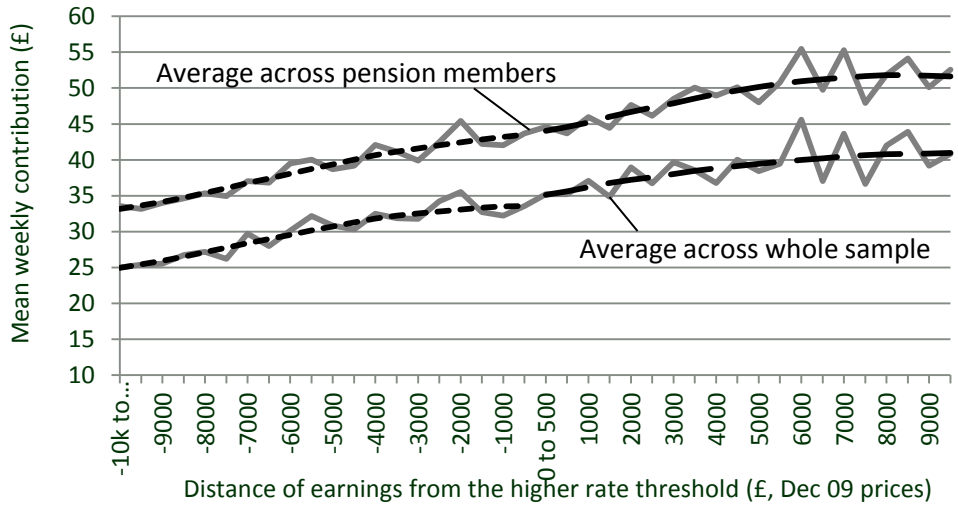

Table 6: Effect on mean pension contributions of having earnings just greater than the HRT as opposed to just lower

	Those in a pension			All individuals		
τ	0.200	-0.365	1.596	0.633	0.819	2.741*
	(0.735)	(1.115)	(1.717)	(0.689)	(1.026)	(1.585)
γ_1	-0.133	0.228	0.274	-0.166	0.203	0.185
	(0.139)	(0.443)	(1.445)	(0.130)	(0.396)	(1.352)
Characteristics	✓	✓	✓	✓	✓	✓
Model	C	C	C	C	C	C
Sample bandwidth	£20,000	£10,000	£4,000	£20,000	£10,000	£4,000
N	26,062	12,299	4,754	33,432	15,617	6,061

Notes: The coefficient τ is the effect of the HRT on pension contributions, while γ_1 is the effect of the HRT on the marginal effect of an additional £1,000 of earnings on contributions. Standard errors in parentheses. ***, **, * indicate significance at the 1%, 5% and 10% levels respectively.

Looking at average contributions across all individuals in the sample, there is some indication that there might be a small positive discontinuity at the HRT - the coefficient τ is at least consistently positive in the last 3 columns of Table 6. As a point of reference, if pension coverage increased by 3 percentage points at the HRT, and if all the individuals brought into pension saving by the HRT contributed the same average amount as those just below the HRT, then a positive discontinuity of about £1.30 would be expected. This is in the range of the coefficient estimates reported in Table 6. Again there appears to be little consistent or significant effect of the HRT on the relationship between pension contributions and earnings. This could be an indication that individuals did not respond to the greater upfront incentive to save (and consequently the increasing average incentive to save) by increasing their

rate of pension contributions. Alternatively, it could be the case that individuals who would have also contributed to a pension if their earnings were below the HRT increased their rate of pension contributions, but that this is offset by those brought into pension saving by the HRT having a lower rate of pension saving. As previously stated, given the nature of our data we are unfortunately unable to identify whether or not the increase in the upfront incentive to save in a pension that occurs when individuals cross the HRT results in an increase in pension contributions.

4.2. The effect of partner's income on the incentive to save

The second example of where individuals may respond to tax incentives that we consider is pension saving within couples. Here we take as our sample all married individuals who are a basic rate taxpayer and who have a partner either just above or just below the HRT. As discussed in section 3.2, it might be expected that those with a partner just above the HRT would be *less* likely to engage in pension saving, and on average save lower amounts, than those with a partner just below the HRT.

Coverage effects

Figure 7 plots the proportion of married individuals contributing to a pension, according to the distance of their partner's earnings from the HRT. As in earlier Figures, the dashed lowess curves would be expected to join at the HRT if there were no effect of the individual's partner hitting the HRT on their pension membership. In fact, it seems as though there is around a 2.5 percentage point *increase* in pension coverage associated with the married individual's partner hitting the HRT. This positive effect is contrary to expectations insofar as, if individuals were responding optimally to the financial incentives to save in pensions then those individuals who have a partner with earnings just above the HRT should be less likely to contribute to a pension than those who have a partner with earnings just below the HRT, since the pension saving should be done by the partner with the greater incentive to save. The parametric analysis which controls for the individual's own characteristics, the results of which are reported in Table 7, indicates that the individual's partner hitting the HRT appears to have a positive effect on pension coverage of the order of around 5 percentage points.

Figure 7: Proportion of individuals contributing to a pension, by partner's earnings

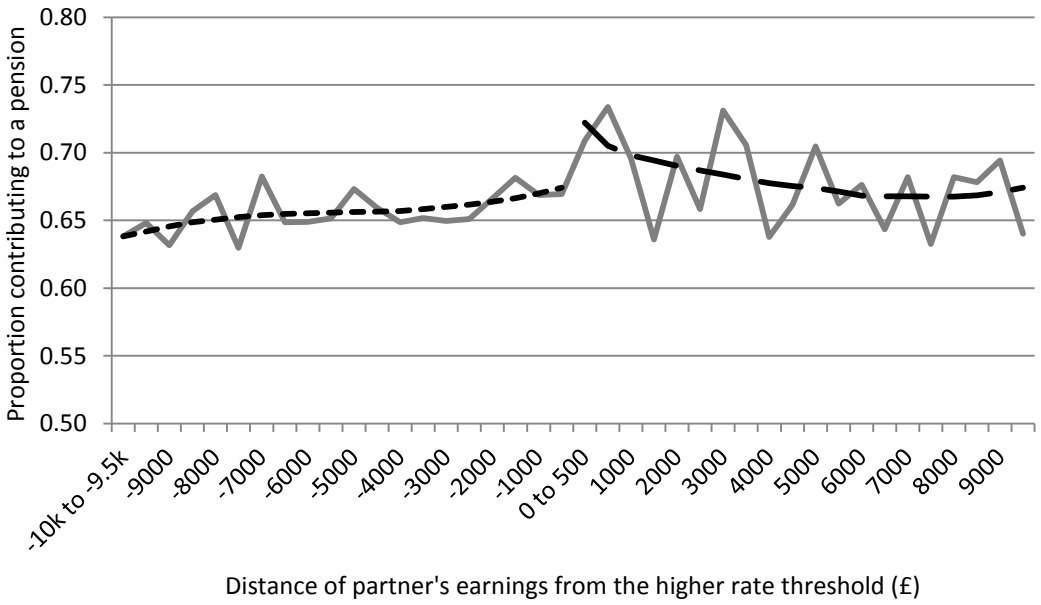

Table 7: Effect on pension coverage of having partner with earnings just above the HRT as opposed to just below

	h = £10,000		h = £5,000		h = £2,000	
τ	0.063**	0.050**	0.055*	0.044**	0.051	0.013
	(0.025)	(0.023)	(0.023)	(0.021)	(0.037)	(0.035)
γ_1	-0.009	-0.000	-0.012	-0.005	-0.026	-0.014
	(0.012)	(0.011)	(0.008)	(0.008)	(0.033)	(0.030)
γ_2	0.002	0.001				
	(0.001)	(0.001)				
N	*14,585*	*14,312*	*6,838*	*6,713*	*2,630*	*2,569*
Characteristics	×	✓	×	✓	×	✓
Income specification	Quadratic	Quadratic	Linear	Linear	Linear	Linear
Model	B	D	A	C	A	C

Notes: Standard errors in parentheses. ***, **, * indicate significance at the 1%, 5% and 10% levels respectively.

This result clearly raises doubts about the tax incentive 'story' for membership of pension schemes in relation to partners' saving decisions. It could however support a 'signalling effect' one, in which one partner reaching the HRT is expected to result in both partners taking a greater interest in retirement saving.

Contribution levels

As discussed in section 3.2, for a number of reasons individuals in couples might be expected to be more likely to react to a change in one individual's upfront

incentive to save by changing the relative magnitude of contributions rather than for one individual to cease pension contributions entirely in favour of the individual with the greater incentive doing all of the saving. We are therefore interested in knowing whether the average contributions of those whose partner earns above the HRT are greater or lower than similar individuals whose partner earns below the HRT.

However as was the case in the individual analysis in section 4.1, this effect is not one that can be easily identified in our data. Comparing the pension contributions of those who have a partner just above the HRT with the contributions of those who have a partner just below the HRT will conflate two effects: the effect of a partner reaching the HRT on the pension contributions of those whose pension membership decision is not affected by the HRT, and a selection effect if those whose membership decision is affected have different average contribution rates from those whose decision is not affected. The first of these would be expected to result in a negative discontinuity in pension contributions at the HRT for the partner, while for the latter of these it is not clear *a priori* what the expected direction of any resulting discontinuity would be (particularly since the results above indicate that the coverage effect appears to run counter to expectations).

Figure 8 shows mean pension contributions of married individuals by their partner's distance from the HRT, and lowess curves for individuals with partners above the HRT and below the HRT separately. Looking first at the average across individuals who are members of a pension, there appears to be a very slight negative effect on individual pension contributions of his or her partner reaching the HRT. However, this visual result is not borne out in the parametric results reported in the first three columns of Table 8. Looking across all individuals, there is a small positive discontinuity in pension contributions associated with the partner reaching the HRT. If all individuals who joined a pension as a result of their partner being above the HRT contributed the same average amount as those in a pension whose partner was just below the HRT, the effect on pension contributions would be about 1.7 – in the ballpark of the coefficients reported in Table 8. This could therefore imply that there is no effect of an individual reaching the HRT on their partner's pension saving. However, it could of course be the case that those individuals brought into pension saving have higher average savings rates, and that this is being offset by those whose

membership decision is not affected decreasing their pension saving when their partner hits the HRT. Unfortunately, similar to the individual analysis, it is not possible with the data available to precisely identify the effect of a partner hitting the HRT on an individual's pension contributions.

Figure 8: Mean pension contribution, by partner's earnings

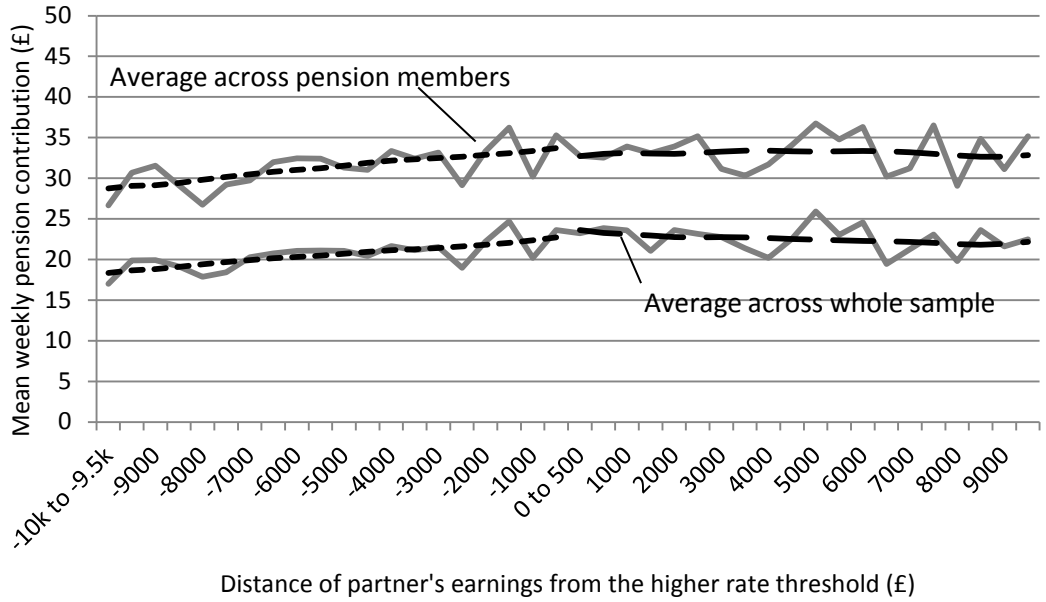

Table 8: Effect on pension contributions of having partner with earnings just above the HRT as opposed to just below

	Those in a pension			All individuals		
τ	0.660	1.003	0.112	1.234**	1.628*	-0.298
	(0.793)	(1.114)	(1.732)	(0.628)	(0.887)	(1.399)
γ_1	-0.154	-0.587	3.273**	-0.183	-0.348	2.691**
	(0.148)	(0.3807)	(1.507)	(0.115)	(0.298)	(1.225)
N	*9,267*	*4,391*	*1,709*	*14,312*	*6,713*	*2,569*
Characteristics	✓	✓	✓	✓	✓	✓
Model	C	C	C	C	C	C
Sample bandwidth	£20,000	£10,000	£4,000	£20,000	£10,000	£4,000

Notes: Standard errors in parentheses. ***, **, * indicate significance at the 1%, 5% and 10% levels respectively.

5. Conclusions

This paper has considered how pension scheme membership and contributions are affected by the incentives to save created by the structure of income tax rates. The incentive to engage in tax-relieved retirement saving increases as individuals cross into higher rates of tax. Consequently we have investigated: first, whether individual membership of, and contributions to, pensions increases in proximity to the higher

rate threshold (HRT) and second, the more complex decision-making case where couples engaged in tax planning would optimally shift their pension saving according to the joint marginal tax rates within the couple. In other analysis we are investigating whether the introduction of the New Tax Credit, which increased incentives to save in a pension for those claiming tax credits by increasing the effective tax relief on pension contributions, resulted in greater pension saving.

The theoretical background to all the analysis is one in which individuals do not have full understanding of the tax regime which is complex and not always transparent; we therefore expect a more robust statistical effect in the case of a 'simple' tax incentive (for example, an individual's own marginal tax rate) as opposed to a more 'complex' incentive (for example, household tax planning across different individuals' tax thresholds, or where the income tax and tax credit systems interact). Moreover individuals may plan relative to simple measures of 'income' such as pay slip annualised earnings rather than full income measures, and may be characterised by inertia in the face of tax changes or income levels where the gains to re-optimising are relatively small.

Given all these facets, we find some evidence that individuals above the higher rate threshold of income tax are more likely to be members of a pension scheme. This could be due to individuals responding to the higher upfront incentive to save created by the greater tax relief on pension contributions at that point, but that is by no means the only plausible 'story'. We cannot reject that this is a 'signalling effect' rather than an optimisation one: while we do not see any such jump in contributing to other standard tax relieved instruments such as ISAs, the signal could be more specifically about the need to be saving in a pension. Such a signalling 'story' would be supported by the evidence from the analysis of couples where, contrary to expectations, we find that individuals with higher marginal rate partners are more likely to save in pensions. There is also no strong evidence of an increase in the *rate* of contribution at the HRT, or of a decrease in the level of contributions when a partner reaches the HRT – though identification is problematic here given the data available. Finally, in all the analysis conducted the effect is most robust when an individual's distance from the HRT is defined in terms of earnings rather than income, giving the interpretation that individuals respond to the salient features of their total income.

References

Akerlof, G. and Yellen, A. (1985) 'A near-rational model of the business cycle with wage and price inertia', *Quarterly Journal of Economics*, 100, 823-838.

Attanasio, O. and DeLeire, T. (2002) 'The effect of individual retirement accounts on household consumption and national saving', *Economic Journal*, 112, July, 504-538.

Barrientos, A. (1998) 'Supplementary pension coverage in Britain', *Fiscal Studies*, 19, November, 429-446.

Bartolome, C. (1995) 'Which tax rate do people use: Average or marginal?' *Journal of Public Economics*, 56, 79-96.

Benjamin, D. J. (2003),, "Does 401(k) Eligibility Increase Saving? Evidence from Propensity Score Subclassification," *Journal of Public Economics*, 87, 1259–1290.

Bernheim, B.D. and Scholz, J.K. (1993) 'Private saving and public policy' in Poterba, J. (ed.) *Tax Policy and the Economy*, Volume 7, NBER, Cambridge: MIT Press.

Blundell, R., Duncan, A. and Meghir C. (1998) 'Eastimating labor supply responses using tax reforms', *Econometrica*, 66, July, 827-861.

Brewer, M., Saez, E., and Shephard, A. (2010), 'Means-Testing and Tax Rates on Earnings', in J. Mirrlees, S. Adam, T. Besley, R. Blundell, S. Bond, R. Chote, M. Gammie, P. Johnson, G. Myles, and J. Poterba (eds*), Dimensions of Tax Design: The Mirrlees Review*, Oxford: Oxford University Press for Institute for Fiscal Studies.

Chernozhukov, V. and P. Hansen (2004): "The effect of 401k participation on the wealth distribution: An Instrumental Quantile Regression Analysis", *Review of Economics and Statistics*, 86, 735-751.

Chetty, R. (2011) 'Bounds on elasticities with optimisation frictions: A synthesis of micro and macro evidence on labour supply', *forthcoming, Econometrica.*

Chetty, R, Looney, A. and Kroft, K. (2009) 'Salience and taxation: Theory and evidence.' *American Economic Review*, 99(4): 1145–77.

Chung, W., Disney. R, Emmerson, C. and Wakefield, M. (2008) 'Public policy and retirement saving incentives in the UK', 169-209 in De Menil, G., Pestieau, P. and Fenge, R. (eds), *Strategies for Pension Reform*, Cambridge: MIT Press.

Cleveland, W. (1979) 'Robust Locally Weighted Regression and Smoothing Scatterplots', *Journal of the American Statistical Association*, 74(368):829-386

Cleveland, W. and Devlin, S. (1988) 'Locally Weighted Regression: An approach to regression analysis by local fitting', *Journal of the American Statistical Association*, 83(403):596-610

Disney, R. And Emmerson, C. (2002) Choice of pension scheme and job mobility in Britain', Institute for Fiscal Studies *Working Paper* 09/02, London.

Disney, R. Emmerson, C. and Wakefield, M. (2010) Tax reform and retirement saving incentives: take-up of Stakeholder Pensions in the UK', *Economica*, 77, April. 213-233.

Engen, E., Gale, G., and Uccello, C. (1999) **'**The adequacy of household saving' *Brookings Papers on Economic Activity*, 2, 65–187.

Gale, W. and Scholz, J.K. (1994) 'IRAs and household saving', *American Economic Review*, 84, December, 1233–60.

Jones, A. and Posnett, J. (1991a) 'Charitable donations by UK households: evidence from the Family Expenditure Survey', *Applied Economics*, 23, 2, 343-351.

Jones, A. and Posnett, J. (1991b) 'The impact of tax deductibility on charitable giving by covenant in the UK', *Economic Journal*, 101, September, 1117-1129.

Journal of Economic Perspectives (1996) Symposium on 'Government incentives for saving', vol. 10, Fall, pp.73–138.

Liebman, J. and Zechauser, R. (2004) 'Schmeduling', *mimeo*, National Bureau of Economic Research.

Meade, J. (1978) *The Structure and reform of Direct Taxation*, London: Allen & Unwin, for Institute for Fiscal Studies.

Milligan, K. (2003) 'How do contribution limits affect contributions to tax-preferred saving accounts?' *Journal of Public Economics*, 67, 253–281.

Mirrlees, J., S. Adam, T. Besley, R. Blundell, S. Bond, R. Chote, M. Gammie, P. Johnson, G. Myles, and J. Poterba (2011), *Reforming the Tax System for the 21st Century: Tax by Design*, Oxford: Oxford University Press for Institute for Fiscal Studies.

Pensions Commission (2004), *Pensions: Challenges and Choices. The First Report of the Pensions Commission*, London: The Stationery Office (http://www.pensionscommission.org.uk/publications/2004/annrep/index.asp).

Poterba, J. (1994) (ed.) *Public policies and household saving*, Chicago: Chicago UP for National Bureau of Economic Research.

Scholz, J.K., Seshadri, A., and Khitatrakun, S. (2006) 'Are Americans saving 'optimally' for retirement?' *Journal of Political Economy*, 114, 4, 607–643.

Scharf, K. and Smith, S. (2009) *Gift aid donor aid research: exploring options for reforming higher rate relief*, Report to HMRC and HMT, London: (http://webarchive.nationalarchives.gov.uk/+/http://www.hm-treasury.gov.uk/d/gift_aid_reseach_report_091208.pdf)

Venti, S. and Wise, D. (1987) 'The determinants of IRA contributions and the effect of limit changes', in Bodie, Z., Shoven, J. and Wise, D. (eds) *Pensions in the US Economy*: Chicago: Chicago UP for National Bureau of Economic Research.

Wakefield, M. (2009), How Much Do We Tax the Return to Saving?, IFS *Briefing Note* no. 82, London: Institute for Fiscal Studies (http://www.ifs.org.uk/bns/bn82.pdf).

Appendix A

A. Change in observed characteristics with income

Notes for Figure A1 to Figure A6: The solid grey line shows the proportion/average of each characteristic for individuals grouped into £500 bands of income by distance from the higher rate threshold. The dashed lines are lowess curves, created using locally weighted scatterplot smoothing for individuals above and below the HRT separately.

Figure A1: Percentage of sample male, by income

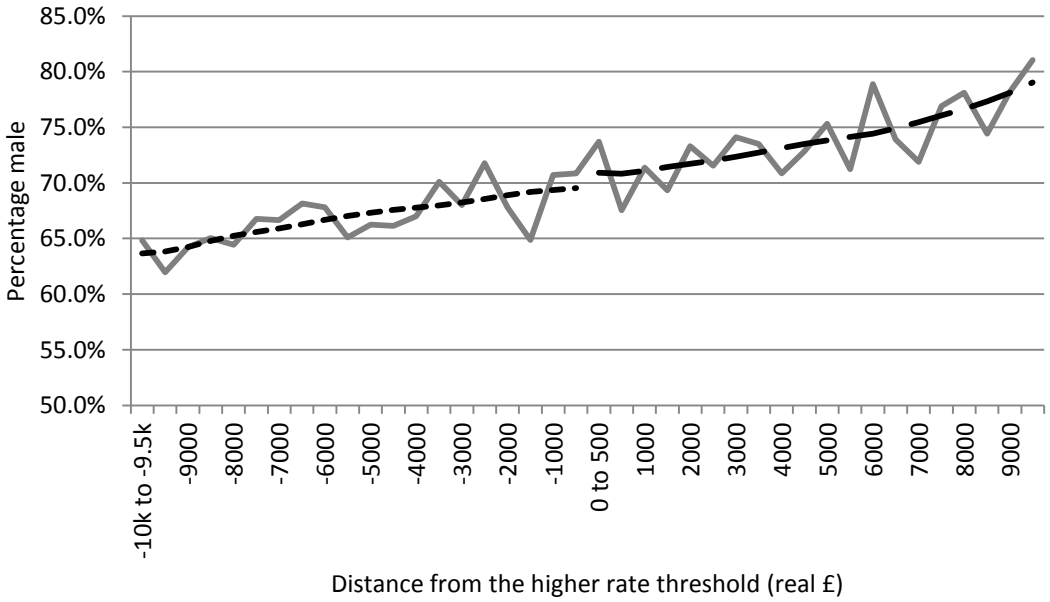

Figure A2: Average age, by income

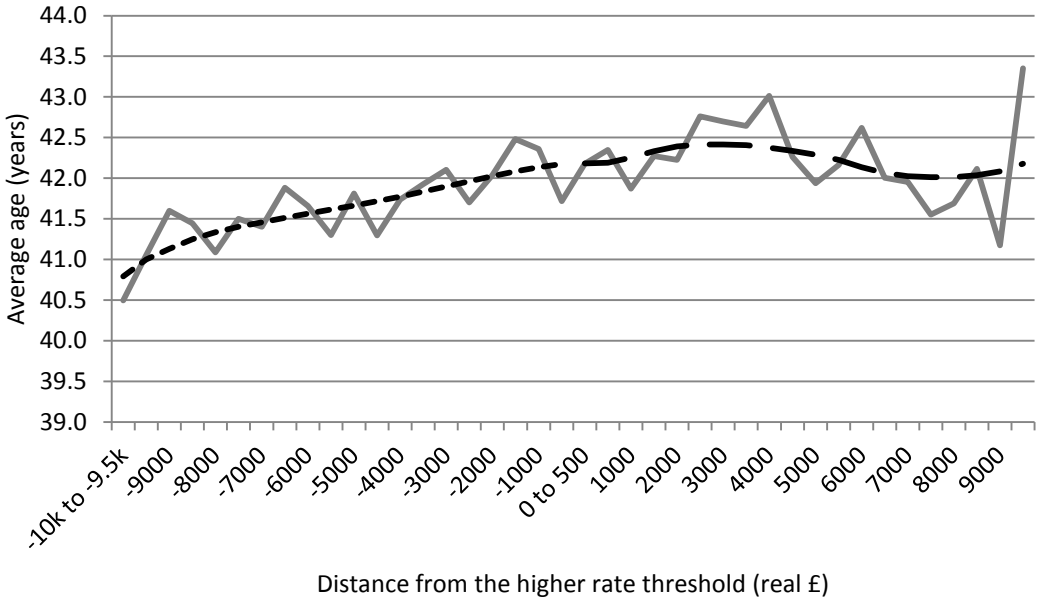

Figure A3: Percentage of sample married, by income

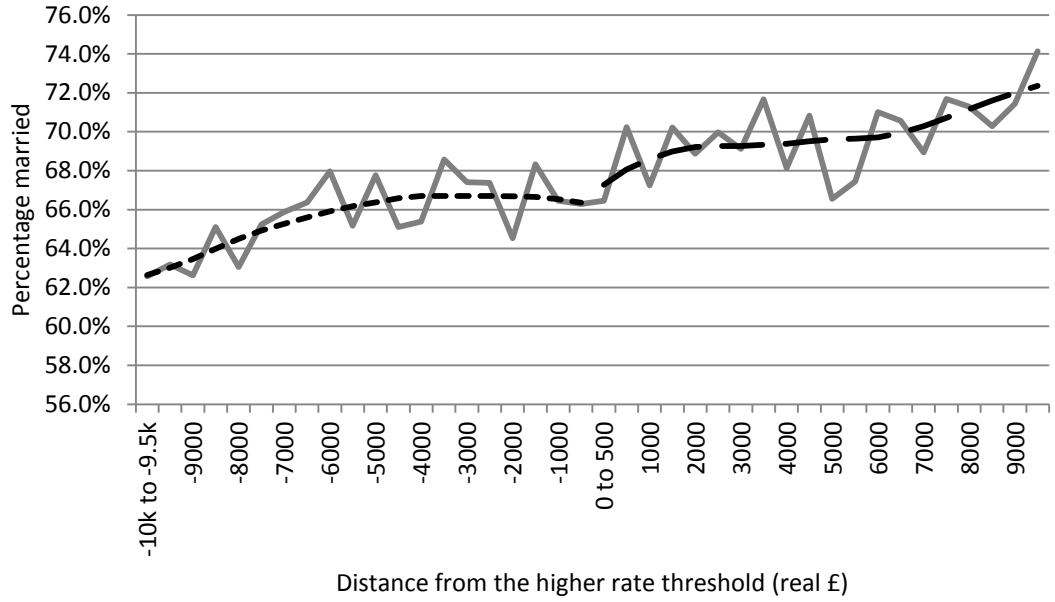

Figure A4: Percentage of sample cohabiting, by income

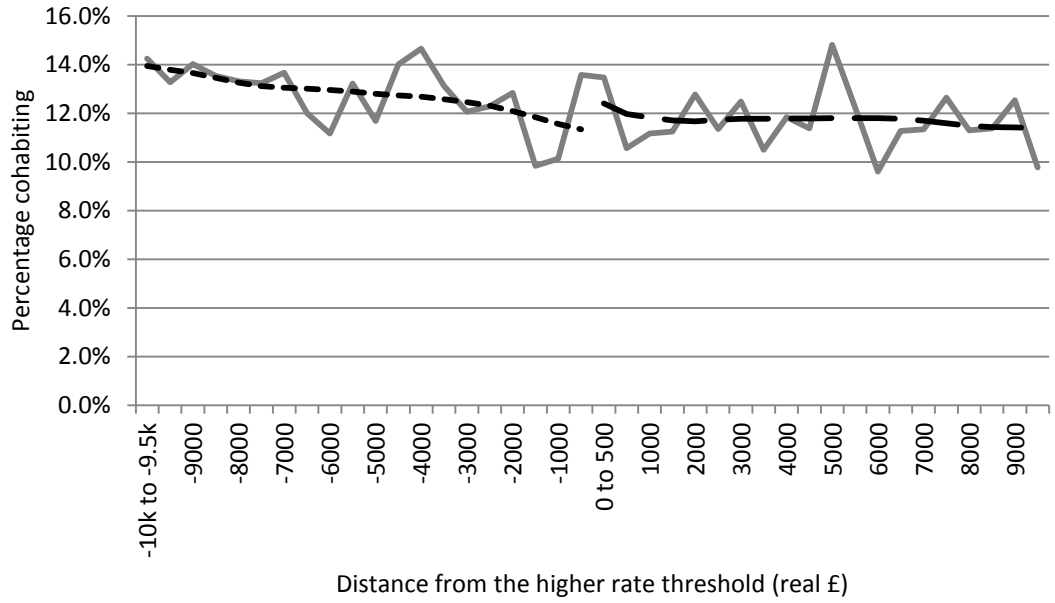

Figure A5: Percentage of sample with high/low levels of education, by income

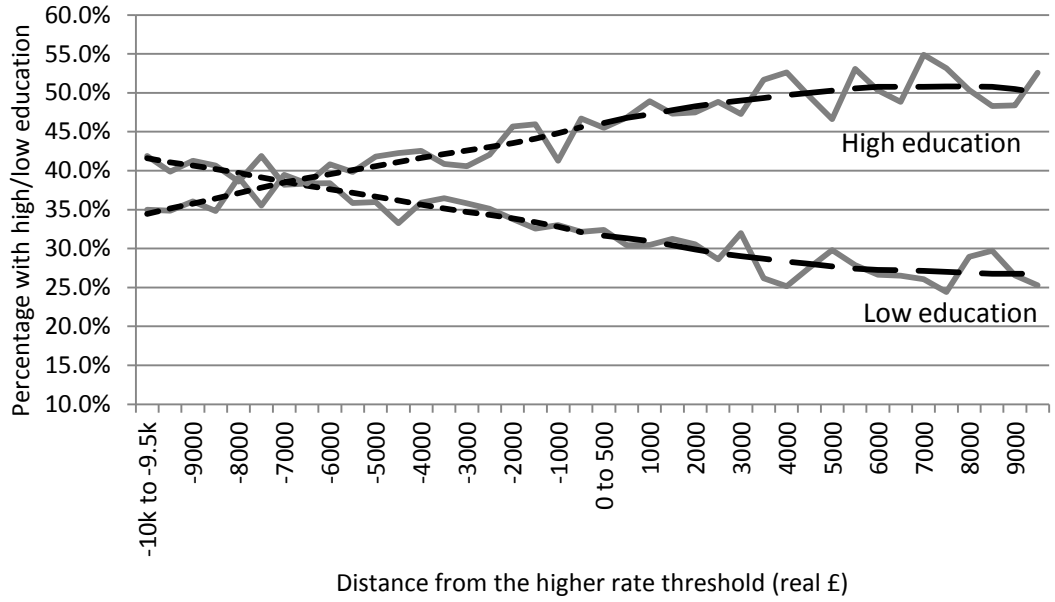

Note: High education is defined as those who left full-time education at age 19 or above, while low education is defined as those who left full-time education at age 16 or below.

Figure A6: Average hours worked, by income

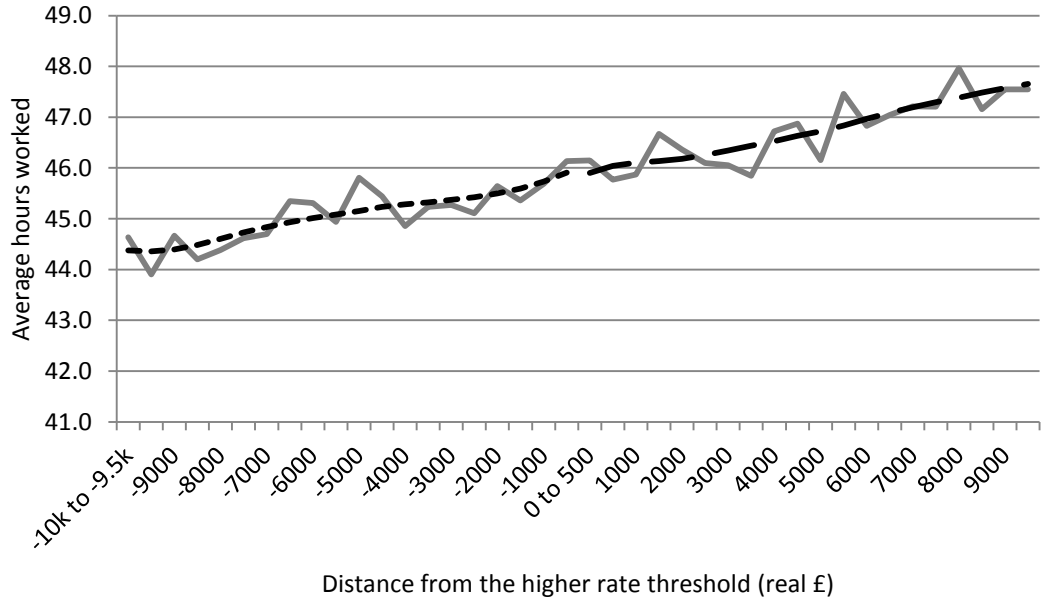

Table A1: Effect on pension coverage of being just above placebo thresholds as opposed to just below

	h = £20,000		h = £10,000		h = £4,000	
Model:	C	D	C	D	C	D
Taxable income						
Threshold = HRT + £10,000						
τ	0.003	-0.010	-0.004	-0.004	0.018	0.038
Std. error	(0.014)	(0.020)	(0.019)	(0.029)	(0.031)	(0.046)
N	15,578	15,578	7,189	7,189	2,740	2,740
Threshold = HRT - £10,000						
τ	-0.007	-0.005	-0.008	-0.012	-0.005	-0.009
Std. error	(0.007)	(0.010)	(0.009)	(0.014)	(0.015)	(0.022)
N	65,900	65,900	32,070	32,070	12,955	12,955
Threshold = HRT + £5,000						
τ	-0.001	0.016	0.013	0.036	0.047**	0.014
Std. error	(0.011)	(0.016)	(0.015)	(0.023)	(0.024)	(0.035)
N	23,190	23,190	10,452	10,452	4,103	4,103
Threshold = HRT - £5,000						
τ	-0.003	-0.004	-0.001	-0.003	0.021	-0.002
Std. error	(0.007)	(0.011)	(0.010)	(0.015)	(0.016)	(0.024)
N	48,071	48,071	23,656	23,656	9,256	9,256
Earnings						
Threshold = HRT + £10,000						
τ	-0.001	-0.028	-0.017	-0.029	-0.002	0.022
Std. error	(0.014)	(0.021)	(0.019)	(0.030)	(0.032)	(0.048)
N	15,018	15,018	6,924	6,924	2,681	2,681
Threshold = HRT - £10,000						
τ	-0.002	0.000	-0.002	-0.001	0.006	-0.006
Std. error	(0.007)	(0.010)	(0.009)	(0.014)	(0.015)	(0.023)
N	65,477	65,477	31,906	31,906	12,959	12,959
Threshold = HRT + £5,000						
τ	0.000	0.005	0.008	0.025	0.036	-0.031
Std. error	(0.011)	(0.017)	(0.016)	(0.024)	(0.025)	(0.037)
N	22,541	22,541	10,043	10,043	3,909	3,909
Threshold = HRT - £5,000						
τ	-0.005	-0.006	0.002	-0.006	0.016	0.003
Std. error	(0.007)	(0.011)	(0.010)	(0.015)	(0.016)	(0.024)
N	47,523	47,523	23,389	23,389	9,143	9,143