

Tax reform and retirement saving incentives: take-up of Stakeholder Pensions in the UK

by

R. Disney[†], C. Emmerson* and M. Wakefield*

Abstract

In April 2001, the UK government introduced Stakeholder Pensions – a new private pension arrangement. The reform also changed the structure of tax-relieved pension contribution ceilings, increasing their generosity for lower-earners. We examine the impact of these changes on private pension coverage using individual level data. We use a difference-in-differences strategy with an estimator that is modified to allow for dichotomous outcomes. Contrary to the conventional wisdom that the Stakeholder Pension reforms had little or no impact on saving behaviour, our results indicate that the change to the contribution ceilings affected private pension coverage rates among lower-earners, especially among women.

Key words: Retirement saving pensions tax incentives

JEL classification: E21 G23 H24

* Institute for Fiscal Studies, London

† Institute for Fiscal Studies, London and University of Nottingham

Corresponding author: Disney, School of Economics, Sir Clive Granger Building, University of Nottingham, Nottingham, NG7 2RD, UK. Email: richard.disney@nottingham.ac.uk. Financial support from the ESRC-funded Centre for the Microeconomic Analysis of Public Policy at IFS (grant number M535255111) is gratefully acknowledged. The ESRC Data Archive at Essex made available the Family Resources Survey data used here, and the data are copyright of the Archive. The authors are very grateful to Woojin Chung for assistance with setting up the data used in this paper at an earlier stage of the research project. The authors also thank Erich Battistin, Richard Blundell, Monica Costa Dias, Kevin Milligan and Sarah Smith, and seminar participants at the Centre for Applied Microeconometrics (CAM) at the University of Copenhagen, at FMG at the London School of Economics, at the Institute for Fiscal Studies, the RES Conference 2005, the RTN Economics of Ageing meeting in Paris and the 2007 NBER Summer Institute for useful assistance and comments. Any errors and all opinions expressed are those of the authors.

Tax reform and retirement saving incentives: take-up of Stakeholder Pensions in the UK

1. Introduction

Given demographic pressures on social security finances in most OECD countries, a common policy target is to encourage households to increase their own private retirement saving. In the United Kingdom (UK), many households already rely on private sources rather than social security for much of their retirement income (Banks *et al*, 2005) but, as in other countries, there has been well-publicised concern as to the extent of a ‘savings gap’ between how much working age individuals should save for retirement and what they are actually saving.¹ There is, however, little agreement in the evaluation literature as to *what* saving policies are effective. There has been a substantial debate around this question in the United States (as in, *inter alia*, Bernheim and Scholz, 1993; Poterba, 1994; Journal of Economic Perspectives, 1996; Benjamin, 2003; and, Chernozhukov and Hansen, 2004), but relatively little econometric evidence relating to this question for the UK despite the plethora of recent reforms affecting pensions and other savings instruments there.²

Greater financial incentives to encourage retirement saving are an obvious policy instrument. But it is difficult to target incentives on the marginal saver, so that more generous incentives may actually reduce private retirement saving for the intra-marginal saver through a wealth effect. The cost to the exchequer of providing incentives will also mitigate the impact of incentives on aggregate (i.e. public plus private) saving. Several changes to the UK’s tax regime governing retirement saving have been implemented in the last two decades but, from an evaluation perspective, it is hard to disentangle the effects of these tax regime changes from other reforms taking place at the same time.

In this paper, we consider a reform that embodied a differential change in tax incentives: the introduction of Stakeholder Pensions in the UK in April 2001. This reform, which was intended to encourage overall take-up of, and contributions to, private pensions, contained several provisions that are discussed briefly in the next section. An important feature of the change in the tax regime that accompanied the reform was that it only affected a sub-set of the population. We can therefore use a standard difference-in-differences evaluation technique to examine the impact of this component of the reform by comparing the behaviour over time of those who were potentially affected by the tax reform relative to those who were unaffected. The paper

closest in spirit to the present one is Milligan (2003), which focuses on differential limit changes and Canadian retirement saving.

The format of the paper is as follows. In Section 2, we briefly describe the overall Stakeholder Pension reform but focus in particular on the change in the tax regime for private pensions that occurred simultaneously with the introduction of Stakeholder Pensions. Section 3 describes the data and the estimator used to examine the impact of the change in the tax regime on take-up of private pensions. Section 4 discusses the empirical results of the effect of the reform on pension coverage and retirement saving. Section 5 provides a brief conclusion.

2. Stakeholder Pensions

Stakeholder Pensions were introduced by the UK government from April 2001 as a new tax-subsidised retirement saving instrument. Like existing personal pensions, and some occupational pension schemes³, Stakeholder Pensions are ‘defined contribution’ schemes (i.e. an individual’s pension benefits depend on the accumulated value of her fund). Individual contributors to Stakeholder Pensions could choose to opt-out of part of the social security programme, and also to obtain tax relief on contributions to their Stakeholder Pension, up to an annual ceiling described in more detail shortly. Stakeholder Pensions differed from pre-existing Personal Pensions (introduced for employees in 1988) in having compulsory minimum standards, a different governance structure, guaranteed workplace access for those working for medium-scale or large employers, and a simpler and more uniform administrative cost structure. Further details on these aspects of the Stakeholder Pensions pension reform are given in Disney *et al* (2007).

The government’s Green Paper (Department of Social Security, 1998) which proposed the introduction of Stakeholder Pensions identified middle income earners – defined as those earning between £9,000 and £18,500 per annum in 1998 prices – as a target group for the reform, although Stakeholder Pensions are available to everyone. High income earners, it was assumed, already had access to other retirement saving instruments, and lower-earners were generally seen by the Green Paper to be better off accumulating rights in the public second tier pension (the State Earnings-Related Pension Scheme, SERPS, superseded in April 2002 by the more redistributive State Second Pension, S2P), rather than opting for a private pension arrangement.⁴ Furthermore other pension reforms introduced since 2001, most notably the introduction of the means-

tested Pension Credit in October 2003, also make public provision more generous for people with low lifetime incomes, although there is a group of lower-earners – those with rich spouses – who would be less likely to gain from the introduction of the Pension Credit.⁵

The Green Paper also proposed a number of other changes to the pension regime, including a reform to the structure of tax reliefs that was also implemented in April 2001. This reform forms the main focus of the present paper. The broad features of the tax regime, before and after the advent of Stakeholder Pensions, are as follows. The UK has an individual-based income tax system. Tax privileged pension contributions, up to certain ceilings, can be made by the employer or the individual. Employer contributions receive up-front tax relief at the individual's marginal rate against income tax and National Insurance.⁶ Individual contributions receive 'relief' (deferral) at the higher-rate of income tax for higher-rate taxpayers and at the basic-rate of income tax for *all* other individuals, i.e. including those with individual incomes below the threshold at which income tax becomes payable. Returns are broadly tax-exempt and pensions are then subject to income tax at withdrawal, except for up to 25% of the fund that can be withdrawn tax-free.

Until April 2001, as depicted in Table 1, the ceiling on individual contributions to pension plans was proportional to earnings and more generous for older individuals. Individuals without earnings could not gain tax relief on pension contributions.

Table 1:
Pre-2001 tax reliefs for defined contribution pension plans
Maximum contributions as a % of earnings by age


Age at start of tax year	Maximum contributions as % of earnings
35 or under	17.5%
36 to 45	20.0%
46 to 50	25.0%
51 to 55	30.0%
56 to 60	35.0%
61 to 74	40.0%

Notes: Contributions were subject to an overall earnings cap. In 2005–06, this was set at £105,600. Maximum contributions include contributions by both the employer and employee.

An important difference between the post-Stakeholder Pension tax regime (which applied to all personal pensions including Stakeholder Pensions) and the previous tax regime was that *every* individual, irrespective of any earnings, was able to make gross contributions of up to £3,600 a year to his or her private pension (which, for an individual receiving tax relief at the basic rate, would require a net contribution of

£2,808). In the new regime, individuals were then allowed higher contributions in line with their earnings as in the previous regime in Table 1. Overall, the effect of this change was to raise contributions limits significantly for lower-earners, especially for younger age groups (since maximum contributions as a proportion of earnings are lower). Figure 1 depicts the effect of the change post-2001 on the maximum gross contribution limits by gross relevant earnings for the various age groups in Table 1. Note that individuals with no earnings could also contribute up to the £3,600 maximum. It is worth re-emphasising that the UK's tax system is individual-based so that each individual in a couple could contribute up to this maximum.⁷

Figure 1:
Maximum annual gross contribution limit, by annual gross relevant earnings and age, personal pension tax regime from 2001 to 2006



Although this change in tax regime was highlighted less than the ‘targeting’ of middle earners in the 1998 Green Paper and subsequent discussion of the legislation, an implication of it was described in the Green Paper:

“The changes will also make it easier for partners to contribute to each other’s pensions, again within the overall contribution limits, should they choose to do so.” (p.63)

More precisely, Emmerson and Tanner (1999) noted that:

“The proposals... may be of most benefit to high earners with non-working spouses who have already used up their own tax-free contribution limits or

who want to maximise the value of their joint personal allowances in retirement.” (p.12)

In general, the consensus in the pensions industry and among analysts has been that the Stakeholder Pension reform was a failure, with little impact on either the take-up of private pensions or the real value of retirement saving. This apparent failure has contributed to the decision to enact yet another change in the UK’s retirement saving regime, with legislation in process to enrol employees whose employers do not offer an occupational pension arrangement automatically in a new system of Personal Accounts, unless the employee explicitly chooses otherwise. The rationale for this change towards greater prescription of retirement saving has rested in part on the presumption that the Stakeholder Pension episode illustrated the limited effectiveness of retirement saving policies that relied wholly on individual responses to perceived incentives.

Nevertheless, the introduction of Stakeholder Pensions provides a policy ‘experiment’ with, in effect, both a visible targeting of a new retirement saving instrument on a specified group and a change in the tax regime for pensions for another group; the latter reform was less publicised at the time but for those affected had a direct impact on financial incentives to save in a pension. We briefly show in the next section (and in more detail in Chung *et al*, 2007) that the former aspect (i.e. the headline ‘targeting’ of the reform on middle earners) had no effect on take-up of private pensions. This result accords with the general view on the limitations of the reform.

We then focus the rest of the paper on the impact of the change in contribution limits on pension take-up among the affected group. In this latter respect, our line of argument follows closely the literature in North America that has attempted to identify behaviour in relation to retirement saving off differential changes in contribution limits across sub-groups of the population (as in Venti and Wise, 1987; Gale and Scholz, 1994; Milligan, 2003). In analysing this change in tax reliefs, we show that the reform did indeed induce changes in behaviour, with significant increases in participation rates among sub-groups of the population who benefited from the increase in the ceiling on contributions eligible for tax relief. The consensus that the reform had little or no impact, and the inferences drawn from that conclusion, may thereby have been premature.

3. Empirical strategy

As mentioned in the introduction, our empirical strategy is to use a standard difference-in-differences evaluation technique to examine the impact of the reform to pension contribution ceilings. In this section we describe the data and methods that are the basis of our estimations. The dichotomous nature of our outcome variable (the decision of whether or not to save in a pension) necessitates the somewhat involved estimation technique that we describe. Nonetheless, for interpretation of the results presented in section 4 it is important to keep in mind the underlying principle of difference-in-differences, which is to evaluate the reform by comparing the behaviour over time of those who were potentially affected by the tax reform to the behaviour of those who were unaffected.

3.1. *Data sources and descriptive analysis*

We investigate the determinants of the household decision to take-out a private pension using information from the Family Resources Survey (FRS). The FRS is a large-scale repeated cross section survey used to construct the UK's official statistics on income inequality and income poverty and so it elicits a rich set of information on each household's demographic characteristics, incomes (by detailed component) and other economic circumstances. The FRS asks individual respondents who are in work or who have ever worked (below age 65) whether they or their employer contributes to a pension scheme. The pension arrangements that are explicitly delineated are a 'personal' pension, a company-run pension scheme, or a stakeholder pension. In addition respondents are asked whether the scheme is contributory or non-contributory, when they joined it and if it is 'portable', as well as more detailed questions about own contributions, contracted-out rebates paid into a Personal or Stakeholder Pensions (since individuals can have such schemes without making any additional contributions) and, in the case of a Stakeholder Pension, whether it was organised by the employer or the respondent.⁸

Table 2 Panel A provides data from the Family Resources Survey for the (tax) years 1999–2000 to 2002–03 on pension holdings by type for all employees. According to the table, overall coverage by private pensions has declined slightly over the period. Coverage by employer-provided plans has been constant, and a decline in coverage by Personal Pensions has been not quite offset by the introduction of Stakeholder Pensions and by a slight rise in the number of people with multiple plans.

Panel B reveals the striking finding that coverage has fallen among the high and medium earnings groups over the period (these are the bands delineated by the Green Paper, of £18,500+ and £9,000 to £18,500 respectively) for those aged over 21 and below state pension age (the sampling frame we subsequently use).⁹ Coverage has *risen* among lower-earners and even (marginally) among those reporting no earnings who are below state pension age. Finally, Panel C splits the data by the fraction of the sample that might potentially be affected by the new limit increase, as illustrated in Figure 1. While pension coverage was relatively stable among those did, or would have, received an increase in their pension contribution limit it was falling among those with higher earnings.

Table 2
Pension coverage by type of pension and earnings band: 1999–2000 to 2002–03
Panel A: Employees only

<i>Year</i>	<i>1999–2000</i>	<i>2000–01</i>	<i>2001–02</i>	<i>2002–03</i>	<i>Increase</i>
Type of pension:	%	%	%	%	% point
Personal Pension	11.9	10.8	10.3	8.9	– 3.0
Stakeholder pension	0.0	0.0	0.8	1.3	+ 1.3
Occupational pension	50.3	50.3	50.3	50.2	– 0.1
Combined	2.0	2.2	2.3	2.5	+ 0.4
Aggregate coverage (%)	<i>64.3</i>	<i>63.3</i>	<i>63.6</i>	<i>62.9</i>	<i>– 1.4</i>
Sample size	<i>19,549</i>	<i>18,711</i>	<i>20,418</i>	<i>21,648</i>	<i>80,326</i>

Panel B: All aged 22 to state pension age

<i>Year</i>	<i>1999–2000</i>	<i>2000–01</i>	<i>2001–02</i>	<i>2002–03</i>	<i>Increase</i>
Coverage by earnings band	%	%	%	%	% point
Zero	3.8	4.0	3.9	3.9	+ 0.1
Low	33.7	34.0	35.4	34.2	+ 0.5
Medium	61.9	60.8	61.2	60.6	– 1.3
High	84.5	84.5	83.3	82.2	– 2.3
Aggregate coverage (%)	<i>47.2</i>	<i>46.8</i>	<i>47.4</i>	<i>46.9</i>	<i>– 0.2</i>
Sample size	<i>27,259</i>	<i>25,887</i>	<i>28,026</i>	<i>29,657</i>	<i>110,829</i>

Panel C: All aged 22 to state pension age

<i>Year</i>	<i>1999–2000</i>	<i>2000–01</i>	<i>2001–02</i>	<i>2002–03</i>	<i>Increase</i>
Coverage by limit increase	%	%	%	%	% point
Zero earnings	3.8	4.0	3.9	3.9	+0.1
Limit increase	46.8	46.1	46.9	46.5	– 0.2
No limit increase	81.3	80.9	80.3	79.2	– 2.1
Aggregate coverage (%)	<i>47.2</i>	<i>46.8</i>	<i>47.4</i>	<i>46.9</i>	<i>– 0.2</i>
Sample size	<i>27,259</i>	<i>25,887</i>	<i>28,026</i>	<i>29,657</i>	<i>110,829</i>

Note: The sample includes individuals aged 22 and over up to the state pension age, although a few individuals have to be excluded due to missing data. The sample in Panel B is that used for the

regressions reported in later sections. Rounding explains why figures in the right-hand column may be slightly different from the difference between the 1999–2000 and 2002–03 columns.

Source: own calculations, Family Resources Survey 1999–2000 to 2002–03.

At first sight, some of findings in Table 2 are paradoxical given the stated aims of the policy. The data suggest that the introduction of Stakeholder Pensions has had no effect on overall coverage and indeed coverage by any kind of pension has fallen among the initial ‘target’ group of middle earners. Nor can these declines be explained by a decline in employer-provided occupational pension provision, since this remains constant. Finally, despite the Green Paper suggesting that lower-earners might be better off in the second pillar of the social security programme rather than opting for a personal pension, this is the only group to see a noticeable increase in private pension coverage.

In Chung *et al* (2007) we formally test whether the reform affected coverage among the ‘target group’ of middle earners using a variant of the methodology outlined in the next section. We show that the only earnings group for which there was a significant positive effect on pension coverage were the group delineated as ‘low’ earners by the Green Paper, as suggested in Table 2, Panel B. We calculate that pension coverage amongst this group increased by 3.6 percentage points as a result of the reform. These results are therefore a more precise estimate of the differential trends that can be observed in table 2.

The discussion in Section 2 combined with these results lead us to think that it might be the change in the contribution limits, rather than the targeting of particular earnings bands, that had an impact on pension coverage. This possibility was highlighted in the two quotations in section 2 which pointed to the intra-household incentives to contribute to spouses’ pensions. Accordingly, we examine the effect of the change in the contribution limits in the remainder of this paper. This is done by comparing the pension coverage over time of those who were potentially affected by the tax reform to coverage over time amongst those who were unaffected; the data underlying our analyses are described in Panel C of Table 2.

3.2. *Modelling the pension take-up decision*

Our test is whether the change in contribution limits enacted as part of the Stakeholder Pension reform in April 2001 affected pension saving decisions within households. This reform can be considered as the policy ‘treatment’. Our identification strategy is straightforward, insofar as we can broadly identify the individuals (i.e. those

with earnings below the limits at the time of the reform) who were, and who were not, affected by the tax change. Hence, our basic method of analysis is a difference-in-differences approach. However, our implementation strategy requires further explanation.

To formalise the test, write a general model of retirement saving in which Y_{it}^* is the outcome of the retirement saving decision of individual i at time t , related to a set of individual and household characteristics (X_{it}), an appropriate measure of earnings (Z_{it}), and, to capture trends over time flexibly, to a vector of time dummies (d_t):

$$Y_{it}^* = \beta X_{it} + \gamma Z_{it} + \tau d_t + \varepsilon_{it} \quad (1)$$

Our outcome variable is a dichotomous indicator of whether or not an individual saves in a private pension at a particular point in time. Denoting this dichotomous outcome as Y_{it} , then we think of the (continuous) Y_{it}^* as a latent variable that measures whether or not an individual gains positive utility from saving in a private pension. With a normally distributed error term, this set-up can be analysed using a ‘probit’ model.¹⁰

The hypothesis tested here is that the probability of purchasing a private pension changed differently for those who were potentially affected by the contribution limit increase compared to those who were not. The required counterfactual assumption is that in the absence of the reform the purchase probabilities for those who were and were not affected by the policy change would have followed a common trend. To implement the difference-in-differences exercise, we need to define a ‘treatment’ variable and a ‘post-reform’ variable. The latter is defined as the indicator I_t that measures whether the individual is observed after the beginning of April 2001. The ‘treatment’ variable is another 0–1 indicator, L_{it} , that measures whether, given an individual’s age and annual earnings, the post-April 2001 contribution limit rules would have been more generous to him/her than the pre-April 2001 contribution limit rules.¹¹ If we could estimate the linear relationship (1), then the difference-in-differences model would be estimated by equation (1’) in which the extent of any difference-in-differences would be measured by the coefficient α on the interaction between ‘had a limit increase’ (L_{it}) and the indicator (I_t) for the post April–2001 period:¹²

$$Y_{it}^* = \beta X_{it} + \gamma Z_{it} + \tau d_t + \varphi L_{it} + \alpha L_{it} I_t + \varepsilon_{it} \quad (1')$$

However, in a non-linear model such as the probit (used here), calculated ‘marginal effects’ on interaction terms cannot give a difference-in-differences measure analogous to the coefficients from a linear model. With the discrete outcome set-up, the common trends assumption may not hold for the expectations of Y_{it} (the saving probabilities), since even when affected by the same factors different baseline probabilities are likely to move at different rates simply because they are different distances from the bounds on probabilities of zero and one. Rather, common trends may hold for a transformation of the distribution of the outcome variable and specifically for the inverse probability function, which is assumed to be known and for the probit is $\Phi^{-1}(\cdot)$.¹³ In other words, the assumption of common trends is made for the index rather than for the probability itself. Following Blundell *et al* (2004) this can be written formally as, that in the absence of any ‘treatment’ the following would hold¹⁴:

$$\begin{aligned} &\Phi^{-1} [E(Y_{it} | X_{it}; L_{it}=1, I_t=1)] - \Phi^{-1} [E(Y_{it} | X_{it}; L_{it}=1, I_t=0)] = \quad (2) \\ &\Phi^{-1} [E(Y_{it} | X_{it}; L_{it}=0, I_t=1)] - \Phi^{-1} [E(Y_{it} | X_{it}; L_{it}=0, I_t=0)] \end{aligned}$$

where variables are defined as above.

The right hand side of this equality can be estimated from observations of the ‘control group’ (those not affected by the limit increase) before and after April 2001. Using the common trends assumption as now formulated, this information can in turn be used to construct a counterfactual of how the index would have evolved for each treatment group individual had the change in pension contribution limits not occurred. The impact of the policy can then be evaluated as:¹⁵

$$\begin{aligned} I(X) = &E(Y_{it} | X_{it}, Z_{it}, d_{it}; L_{it}=1, I_t=1) - \Phi \{ \Phi^{-1} [E(Y_{it} | X_{it}, Z_{it}, d_{it}; L_{it}=1, I_t=0)] + \\ &\Phi^{-1} [E(Y_{it} | X_{it}, Z_{it}, d_{it}; L_{it}=0, I_t=1)] - \Phi^{-1} [E(Y_{it} | X_{it}, Z_{it}, d_{it}; L_{it}=0, I_t=0)] \} \quad (3) \end{aligned}$$

Blundell *et al* (2004) propose a method for implementing this ‘difference-of-differences’ estimator of the effect of the policy. A different relationship between the outcome and the observed characteristics is estimated for each group of agents defined according to the various interactions of whether or not the reform would have increased their contribution limit (which depends on their age and earnings) and whether they were observed before or after the reform was implemented. These relationships encapsulate the behavioural patterns of each group and the impact of the reform once it had been enacted. By predicting the outcomes for the ‘treated, after’ group (i.e. individuals characterised by $L_{it}=1$ and $I_t=1$ who were observed after 2001 and whose age and

earnings were such that they were affected by the contribution limit increase) using the behavioural equations for the pre- and post- reform ‘control’ groups, one obtains an estimate of how the underlying index would have changed for individuals in the treated group in the absence of the reform. This can be used in combination with the behavioural equations for the treated group to construct the estimated effect (3). The final estimate of the effect uses predictions made for the ‘treated, after’ group and weighted according to characteristics (X_{it}) in this group. It can therefore be thought of as representing the average impact of treatment on the treated.

3.2.1. *Implementing the estimator: data considerations*

Before examining the results of the difference-in-differences exercise, we return to a *caveat* mentioned in footnote 10. Weekly earnings from the FRS are grossed up to obtain annual earnings, so there is measurement error in the indicator of whether the individual is affected by the change in contribution limits. This error would be particularly pertinent if we had sought to measure the impact of the reform solely by identifying individuals who were on their maximum contribution limits pre-April 2001 and who then benefited from the increase in limits in April 2001, since these individuals could only be identified with error. In any event, such a limited definition of the ‘treated’ group could lead to an underestimate of the overall effect of the policy since it is not true that all of those who might respond to the policy would be observed exactly on their contribution limit in the pre-reform regime. Some lower-earners might, for example, have been induced to take up a private pension rather than doing no pension saving because the large increase in their contribution limit meant that the benefits of making contributions could now outweigh the threshold fixed costs of purchasing a pension.

In contrast to some studies of limit increases in the United States, therefore, we do not see potentially ‘treated’ individuals simply as those who individuals who we calculate to be at the pre-reform contribution. Instead, we assume that *all* individuals or households who received a limit increase as a result of the reform, irrespective of whether their calculated contributions would have been at the pre-reform ceiling, are potentially ‘treated’. From this more general viewpoint, the change in the limits had a large potential coverage. This is illustrated in Figure 2, which shows the fraction of respondents treated by age given our more general measure.

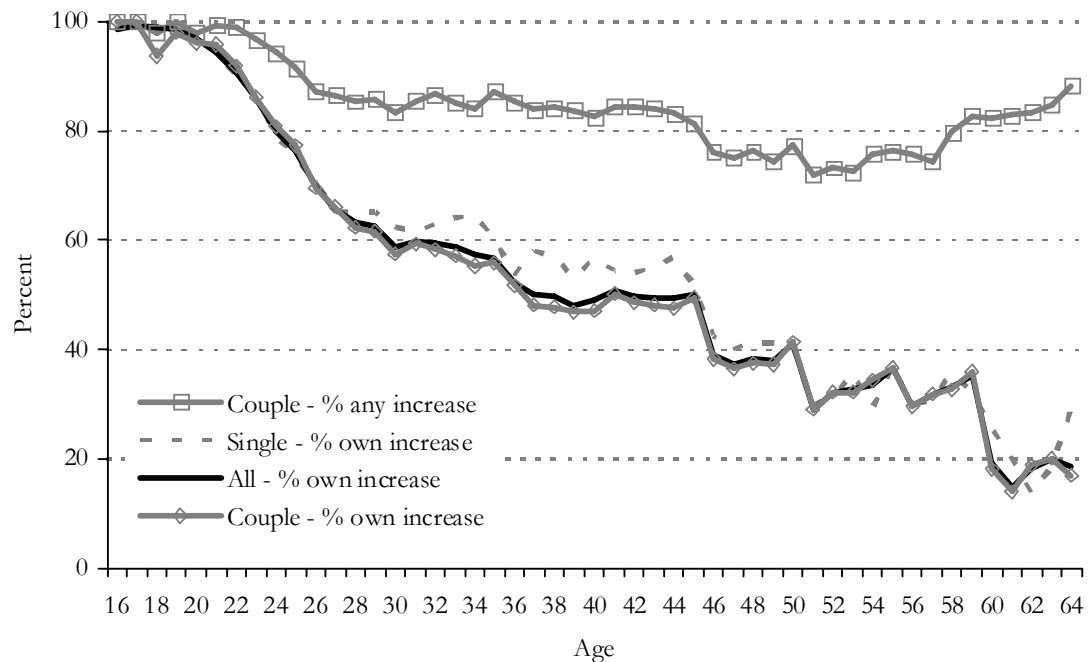
Figure 2 should be read in conjunction with Table 1 and Figure 1, which illustrate that the impact of the limit increase is driven by both age and earnings levels. At the ages

of 36, 46, 51, 56 and 61 there are step changes in the likelihood of being affected by the limit increase because of the higher earnings proportions that can be contributed at each of those ages. These steps are observable in the figures, especially at the three intermediate age ranges.


Eligibility for the limit increase is otherwise driven by two factors: the increase in earnings with age, and the proportion of households with zero earners in each age band. Panel A, which focuses on individuals with positive earnings, shows how rising earnings over the life cycle and the age dependent changes (described in Table 1 and Figure 1) in the proportion of earnings that can be contributed tax-relieved to a pension, combine to reduce the likelihood of being affected by the limit increase when only own earnings are considered. This also implies that, when we consider only positive earners, the ‘control’ group are on average older than the ‘treated’ group; an issue to which we return later. Panel A also shows that almost 100% of individuals aged 21 or under are affected by the limit increase, and so we exclude them entirely from the regression estimates on coverage.

Figure 2: Percentage of individuals receiving limit increase by age and household type

Panel A: Positive earners only



Panel B: All individuals



Note: As Table 2.

Source: As Table 2.

Figure 2 Panel B, which incorporates zero earners within households, gives a somewhat different picture of the impact of the changes in limits by age. Essentially because of cohort differences in participation rates, a greater fraction of older households have at least one non-earner who would be eligible for the limit increase under the reform regime even if the primary earner is above the limit threshold.¹⁶ Thus the (negative) association of the probability of being treated by age is much less marked.

4. Did the reform increase private pension coverage?

4.1. Empirical results

We estimate the difference-in-differences model described by equation (3). This involves estimating separate probits for the treated and controls, pre- and post-treatment date (2001), using the Family Resources Survey for the four years 1999–2000 and 2000–01 (‘pre-treatment’) and 2001–02 and 2002–03 (‘post-treatment’), in order to calculate the estimated treatment effect, as described in the previous section.¹⁷ The regressors comprise whether the individual is single or in a couple, age dummies of both the individual and, where relevant their partner (structured such that the bands coincide with the age ranges for the contribution ceiling bands described in Table 1), sex, age left

school, partner's education, dummies for each period within the pre- and post-treatment regime, earnings and a full set of age-education interactions. We consider those aged over 21 but under the State Pension Age (65 for men and 60 for women). Rather than present a full set of probit estimates (the results are available on request), we provide key calculated treatment effects in Table 3 along with statistical significance (which is estimated using bootstrapping) and sample sizes.

The results in Table 3 confirm that there is a positive effect of the change in the contribution limits on take-up of private pensions among the treated group. The table further analyses the sub-groups which are disproportionately affected.

The first row of Table 3 suggests that the overall result of the reform is an increase in private pension coverage of 2.1 percentage points among those affected by the limit increase. This result averages across positive and zero earners, and indirectly explains the result in Chung *et al* (2007), insofar as this effect is concentrated among affected positive earners (+3.3 percentage points, statistically significant) rather than zero earners (+0.4 percentage points, not statistically significant). We did not find any evidence that the reform had a statistically significant impact on take-up among sub-groups of zero earning respondents (e.g. by sex or whether or not in a couple) and so do not consider single zero earners further.

Among single respondents with positive earners, the total effect on private pension coverage is in an increase of 3.9 percentage points, which is statistically significant at the 5% confidence level. Strikingly, however, the effect among single women is large and statistically significant (+5.4 percentage points) whereas among single men, the effect is again positive but not statistically significant. Among couples, the overall effect is also positive and statistically significant (+3.0 percentage points); again, the effect is stronger (and statistically significant) in couples where the woman is the 'treated' member (i.e. below the contribution limit), whereas the impact on treated men is much smaller.

Finally, we disaggregate couples further by the earnings of the partner. In one specification, we use current earnings of the partner. However since current earnings are subject to transitory shocks we also use partner's education as a proxy for lifetime earnings, where more years of schooling are assumed to be associated with greater earnings and/or participation. A necessary *caveat* to this interpretation is that average years of schooling have increased cohort-by-cohort. The 'treated' group of positive

earners is relatively young on average (see Figure 2.A.) and is therefore more likely to have greater schooling than the ‘control’ group, *ceteris paribus*.

Table 3

Results: Impact of reform of contribution limits on take-up of private pensions for selected groups, using Blundell *et al* (2004) procedure

By zero or positive earnings: respondent

	Zero earners only		Positive earners only		All	
	<i>Predicted level after treatment</i>	<i>Estimated impact of reform</i>	<i>Predicted level after treatment</i>	<i>Estimated impact of reform</i>	<i>Predicted level after treatment</i>	<i>Estimated impact of reform</i>
All	3.9%	+0.4ppt	46.7%	+3.3ppt	28.5%	+2.1ppt

Positive earnings only (respondent)

	Men		Women		All	
	<i>Predicted level after treatment</i>	<i>Estimated impact of reform</i>	<i>Predicted level after treatment</i>	<i>Estimated impact of reform</i>	<i>Predicted level after treatment</i>	<i>Estimated impact of reform</i>
Singles	39.2%	+3.1ppt	40.7%	+5.4ppt	40.1%	+3.9ppt
Couples	47.4%	+0.3ppt	50.3%	+4.3ppt	49.5%	+3.0ppt

Couples, positive earnings only, by partner's earnings/partner's education

	Zero earner		Low-mid earner		High earner	
	<i>Predicted level after treatment</i>	<i>Estimated impact of reform</i>	<i>Predicted level after treatment</i>	<i>Estimated impact of reform</i>	<i>Predicted level after treatment</i>	<i>Estimated impact of reform</i>
Couples	41.3%	+5.7ppt	49.5%	+2.4ppt	55.4%	+2.9ppt
	Low education		Medium education		High education	
	<i>Predicted level after treatment</i>	<i>Estimated impact of reform</i>	<i>Predicted level after treatment</i>	<i>Estimated impact of reform</i>	<i>Predicted level after treatment</i>	<i>Estimated impact of reform</i>
Couples	47.5%	+3.7ppt	54.4%	+7.5ppt	51.6%	-2.2ppt

Notes: Statistical significance estimated using bootstrapping with 1,000 repetitions. Significant treatment effects at 5% level in **bold**. Controls for age, sex, school leaving age, and full age-schooling interactions, whether single or in a couple, and, where relevant, partner's age, partner's school leaving age and partner's earnings. ‘Low-mid’ earner is defined as in 1998 Green Paper ‘low-middle’ earner. Education: ‘low’ = leaving school age 15/16; ‘medium’ = age 18, ‘high’=>18 (tertiary). For sample sizes see Appendix Table A.1.

Using partner's earnings and education, we note that the largest and sole statistically significant effect by partner's earnings group is on positive earners with a zero earning partner. We do not find clear-cut evidence that lower-earners with high earning

partners had a particularly strong response to the reform (inasmuch as the coefficient for this group is positive but not statistically significant at the 5% level)¹⁸; that is, we cannot find conclusive evidence for the suggestion by Department of Social Security (1998) and Emmerson and Tanner (1999) – see section 2 above – that the reform allowed couples jointly to utilise the increased limits by having pension coverage for the lower-earning partner. When we use education as a proxy for long-term earnings of the partner, we get strong but slightly hard to interpret results. Treated respondents whose partners are educated up to 18 exhibit the largest positive responses to the reform (+7.5 percentage points) while a smaller estimated impact is found among those whose partner left school at (or before) the school leaving age. The impact on those partners have the highest education levels is actually negative, albeit not statistically significant.

Overall, the results show statistically significant effects, and for some groups the proportionate effects on pension coverage are substantial. In particular, we observe large effects on pension coverage for lower-earning women (both single and in couples). These results contrast with the conventional wisdom that the overall Stakeholder Pension reform had little or no impact on pension coverage. They also suggest that looking at aggregate trends and at ‘target’ groups such as middle earners has led analysts and policy-makers to overlook significant effects elsewhere in the earnings distribution. The lack of effects on middle earners likely arose because the effective policy ‘treatment’ was the reform of tax relief, which affected lower-earners. The analysis of the next subsection is designed to increase our confidence that the observed effects really were due to this policy reform.

4.2. Alternative explanations and sensitivity analysis

This section considers alternative explanations for our results. In particular, it considers the plausibility in this context of the ‘common trends’ assumption that lies at the heart of the ‘difference-in-differences’ model, and undertakes some sensitivity analysis which focuses on possible implications of pre and post-treatment heterogeneity of the population.

4.2.1. The validity of the common trends assumption

Our results assume that the trend in pension coverage (or more accurately in the index of the underlying propensity to contribute to a private pension) among those who received an increase in the contribution limit as a result of the reform would, in the

absence of the reform, have been the same as among those who were unaffected by the reform to contribution limits. We are overstating (understating) the effect of the tax relief change to the extent that the positive change in coverage among the below-limit group would have been higher (lower) than that of the above-limit group in the absence of the 2001 reform. How plausible is this common trends assumption?

Undoubtedly, the *overall* trend in pension coverage was affected by the change in the financial climate from the beginning of 1998 to the end of 2002, during which the FTSE 100 index of UK equities fell by 31% whereas, on average, UK house prices rose by almost exactly 50% (using the Halifax plc index). This might have induced savers to switch away from saving through private pensions (which were at this stage largely in equity-dominated portfolios) to invest in housing. So if above-limit households either exhibited greater substitutability in their asset portfolios, for example through economies of scale or greater financial acumen, or had different asset portfolios (i.e. more equity-dominated) this might explain the disparate trends.

There is mixed evidence for this alternative ‘story’ in our data. In comparisons of treated earners against controls (for example, of single people) it is certainly true that the ‘typical’ control is older and more likely to have greater housing equity, given the fall in potential treatment rates with age illustrated in Figure 2. This would imply that treated households, with lower housing equity (and perhaps less financial acumen, although younger households tend to have greater formal education to offset this factor) might be less likely to reduce their pension saving in response to house price changes. On the other hand, it is hard to utilise this reasoning to explain away the differential results by the sex of the respondent and for couples by spouses’ income and education.

If, instead, we focus on changes that might affect the ‘treated’ group rather than the control group, there is one reason for thinking that we might *understate* rather than overstate the magnitude of the effect. Here there *is* an important change in the benefit regime that coincides with the introduction of Stakeholder Pensions in 2001. In 2002, the State-Earnings Related Pension (SERPS), the second tier pension, was replaced by the State Second Pension (S2P); a change that had also been announced in the 1998 Green Paper. S2P is more explicitly redistributive towards low lifetime earners in its design. In addition the means-tested benefit for pensioners was formally indexed to earnings rather than prices from April 1999 (unlike the rest of the pension programme),

so increasing its real value for those with low lifetime incomes and reinforcing the disincentive to save for retirement.

Analysis therefore suggests that, whereas replacement rates cohort-by-cohort for the public pension programme have already peaked for average earners, lower-earners are likely to see increasingly generous replacement rates from the public programme for several decades yet (Disney and Emmerson, 2005). To the extent that single women, or couples that include a zero earner or have low education (to take some of the groups where we have found a significant effect) are disproportionately likely to gain from these reforms to the public pension programme relative to the control group, we might expect the treated group's take-up of private pensions to have fallen faster than that of higher income groups, in direct contrast to the results found here.

4.2.2. *Changing heterogeneity among the treated and control groups*

Another possible alternative explanation for the observed 'treatment effect' lies in changes in the composition of households between the treated and untreated group. Household heterogeneity can arise from measured characteristics and from unobserved characteristics, such as individual preferences for saving. If it is assumed that such characteristics are fixed over time then our difference-in-differences estimator should be robust to their presence, although since our data have no panel dimension we cannot directly control for unobserved individual effects and it is for this reason that we are careful not to generalise our empirical measure of the 'treatment effect' to the untreated group. On the other hand, compositional effects of whatever kind could potentially bias our results if the composition of households within the treated group and the controls changed after the reform.

This bias would be most pertinent if observed earnings are endogenous to the reform. As demonstrated in Table 1, prior to 2001, an individual required *some* earnings to be eligible for tax relief on retirement saving; as Figure 1 shows, after 2001 there is an initial lump sum tax relief independent of earnings. It is in theory possible that some people with accumulated cash who wished to engage in retirement saving pre-2001 had to work, at least part-time, in order to obtain eligibility for tax relief; after 2001 this was no longer necessary. In this case observed earnings would be endogenous to the treatment. This possibility, although unlikely, might be pertinent for, say, older women who had low retirement savings and who were not regular workers. If the identification of the treatment is endogenous to the outcome variable, this clearly poses a potential

problem for the analysis. Other possible scenarios, such as greater volatility of intra-household earnings over time, might also have similar effects on the composition of the treated relative to the controls.

We test this possible proposition using the FRS in a manner designed to test for common and differential trends in composition across earnings groups. We pool each ‘treatment’ group (low/zero earners below the contribution limit) with the ‘control’ (those who would not have received an increase in their pension contribution limit) group. Now write the model:

$$J_{it} = \theta L_{t=1} + \gamma_1' X_{it} + \alpha_1' X_{it} L_{t=1} + \varepsilon_{it} \quad (4)$$

Where J_{it} is an indicator variable of whether the observation of the individual occurred during the period in which Stakeholder Pensions were available, L is an indicator of whether or not the reform increased the pension contribution limit of the individual, X_i is a vector of explanatory variables such as age, schooling and partner’s characteristics. This model tests two possibilities. First, the significance (or otherwise) of the vector of coefficients $[\gamma_j]$ tell us whether the characteristics of those observed among the pooled groups before the reform occurred are different from the characteristics of those observed after the reform was implemented. Second, the (lack of) significance of the vector of coefficients $[\alpha_j]$ tells us whether any changes in characteristics over time occurred differently between the control group and the treatment group. It is this second test which is important for our analysis – our results may not imply the policy effects we wish to identify if changes in the characteristics of individuals we observe are occurring differentially across the treatment and control groups. In fact the absence of differential changes in group composition is a more stringent test than is required to support the common trends assumption¹⁹, and so passing the test would be reassuring for our analysis.

Equation (4) can be estimated first using data from the control group and the treated group who did have some earnings, then for the treated group who did not have any earnings and the (whole) control group. In both of these models the results (available on request) show that our sample is, on average, slightly older and slightly better educated in the period after stakeholder pensions were implemented. There is little evidence of a differential change in characteristics between those in the control group and those in the treated group, especially among the treatment group with some earnings. For positive earners, the coefficients on the interaction terms [i.e. vector α_j]

are not jointly different from zero at conventional levels of statistical significance ($\chi^2(18) = 25.08$, $\text{Prob}>\chi^2 = 0.13$). Looking at the individual interaction coefficients, the only significant coefficient is on male, which is negative, suggesting slightly fewer men in the control group relative to the treatment group after the reform. Since men are more likely to contribute to a private pension, this might slightly offset our treatment effect. On the other hand, since men are unlikely to be marginal workers, it provides no support for the argument that our treated group pre-treatment contained a preponderance of earners who were only working to obtain earnings on which to obtain tax relief for their pension contributions.²⁰

4.2.3. *The composition of the control and treatment groups*

Our final sensitivity test attempts to handle the implication of the fact that, for positive earners, our controls are on average older than our treated group (see Figure 2, Panel A). Although we interpret our measured effects as a ‘treatment of the treated’ rather than a consequence of differential limits across age groups, it is interesting to see whether the calculated effects are affected by choosing comparable controls facing a common rate regime. To do this, we focus only on the 36-45 age group. We choose this group because it is a relatively broad (10- as opposed to 5-year) age band that is faced with a single pension contribution ceiling of 20% of earnings in the pre-Stakeholder Pension regime and, conveniently, it contains around 50% of treated individuals and 50% controls (Figure 2, Panel A).

Full results for this age group comparable to Table 3 are not shown in detail.²¹ However they can be summarised as follows. First, as before, the results are only statistically significant for positive earners and, for single people, only among women. The latter coefficient is particularly large (+14.6 percentage points) and statistically significant. Among couples, the largest effects are found, as in Table 3, amongst those whose partner is a zero earner (+5.0 percentage points) and among those whose partner is ‘middle’ educated – that is, educated to age 18 (+8.0 percentage points). These results therefore lead us to believe that the magnitudes in Table 3 for the whole sample are broadly correct.

On balance, therefore, we do not believe that our key results in this section are driven by compositional changes or by the heterogeneity of our control and treatment groups.

4.3. *Reflections on the results*

Our analysis of the composition of our control and treatment groups, coupled with careful consideration of factors that could have disturbed the crucial ‘common trends’ assumption, reinforces our confidence that the results reported in subsection 4.1 can be attributed as causal effects of the stakeholder pension reform package. The effects that we found are for lower-earners and, as argued above, the lack of effects on middle earners likely arose because the effective ‘treatment’ in the policy reform package was the change in incentives arising from the reform of tax relief on lower-earners, rather than the highlighted targeting of middle earners.

Over all individuals affected by the change in contribution limits, the effect on pension coverage has more or less exactly offset what would otherwise have been a falling trend in pension coverage (see Table 2 above). For some groups of the population, the effects on pension coverage are substantial in proportional terms: we found effects of up to 5 – and in one case 7.5 – percentage points, for coverage in groups where around one-half of individuals have private pensions. An important pattern is that we observe particularly strong increases in take-up of private pensions among lower-earning women.

To supplement these results on pension coverage, it would be nice to have results on how the reform has affected the level of pension saving. Unfortunately our data provides only limited information that could be used to explore this issue, as it only records information on individual (and not employer) contributions in to what are reported to be personal or stakeholder pensions. Using these data, in Disney *et al* (2007) we tentatively concluded that an aggregate £160 million²² decline in reported pension contributions of this kind might have been 30% larger in the absence of the policy reform. This tentative figure supports the strong evidence on coverage rates presented in the current paper. The evidence shows that, contrary to the conventional wisdom, the stakeholder pension reform package *did have* important effects on the pension saving of some groups of the population.

5. **Conclusions**

Our starting point was the policy debate concerning the best ways of encouraging people to save for their retirement. Stakeholder Pensions, introduced in 2001, were targeted by the government on middle earners as a means of filling a perceived gap in retirement saving products. The introduction of Stakeholder Pensions was also

associated with a change in pension contribution limits which allowed lower-earners to make larger tax-relieved contributions to private pension schemes. Our analysis represents the first systematic attempt, to our knowledge, to examine the impact of these policy reforms on the probability of households engaging in retirement saving.

Aggregate data suggest that the introduction of Stakeholder Pensions had little impact on the overall propensity to save for retirement. The number covered by private pensions was static and there was a downward trend in pension saving over the period 1999 to 2002, either side of the reform in 2001. This apparent failure of the Stakeholder Pension reform was one factor that encouraged the government to introduce further pension reforms on 6th April 2006²³, and to announce yet more reforms to the retirement saving regime on 22nd May 2006 (see Department for Work and Pensions, 2006a and 2006b) with greater emphasis on a ‘default option’ that encourages individuals to invest a minimum amount in a private pension plan.

Our results suggest that analysts have been too quick in assuming that the 2001 reform had no effect, and also in assuming that individuals failed to respond to the changes in tax incentives. Exploiting a difference-in-differences estimator that allows for the dichotomous nature of the saving decision, we show that these aggregate trends conceal a more complex picture. In particular, our results show that a trend fall in coverage was partially counteracted by the introduction of Stakeholder Pensions, primarily through the associated change in contribution ceilings that disproportionately benefited low and zero earners. We provide some evidence that women, both single and in couples, have benefited from the increase in the joint contribution limits within households, which was an additional intention of the policy. To put this in context, there was an underlying decline in private retirement saving in the early part of the decade (for reasons that we briefly discuss in the text) that would have been greater had it not been for the tax changes associated with the Stakeholder Pension reform in 2001.

In general, our results also suggest that individuals respond to tax incentives in making retirement saving decisions – a result incidentally confirming much of the US literature on the impact of contribution limits on saving in Individual Retirement Accounts (see again Journal of Economic Perspectives, 1996, and the literature cited therein). The results also highlight that it is sometimes important to know the details of a given policy reform, rather than just the ‘headline’ target, in order to understand how the policy might work in practice. Since it is common for large impacts of policy reforms

to be highlighted that turn out to be illusory on subsequent closer analysis, it is ironic that the introduction of Stakeholder Pensions in 2001 was quickly written off as having had little impact on retirement saving: our evidence suggests that some of the changes that were part of the Stakeholder Pension reform package had non-trivial effects on particular groups of the population and thus on aggregate retirement saving.

References

- Ai, C. and Norton, E. C. (2003), ‘Interaction terms in logit and probit models’, *Economic Letters*, vol. pp. 123–129.
- Attanasio, O. P., Banks, J., and Wakefield, M., (2005), ‘The Effectiveness of Tax Incentives to Boost (Retirement) Saving: Theoretical Motivation and Empirical Evidence’, *OECD Economic Studies*, No. 39: Special Issue on Tax-favoured Retirement Saving, pp. 145–172.
- Attanasio, O. P. and Rohwedder, S. (2003), ‘Pension Wealth and Household Saving: Evidence from Pension Reforms in the U.K.’, *American Economic Review*, 93, 5, pp.1499–1521.
- Banks, J., Blundell, R., Disney, R. and Emmerson, C. (2002), *Retirement, Pensions and the Adequacy of Saving*, IFS Briefing Note No. 29. (www.ifs.org.uk/pensions/bn29.pdf).
- Banks, J., Emmerson, C., Oldfield, Z., and Tetlow, G. (2005) *Prepared for retirement? The adequacy and distribution of retirement resources in Britain*, IFS Report no. 67.
- 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. (1992) *Is the baby boom generation preparing adequately for retirement? Technical report*, Unpublished paper, New York: Merrill Lynch.
- 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., Costa Dias, M., Meghir, C. and Van Reenen, J. (2004) ‘Evaluating the employment impact of a mandatory job search program’, *Journal of the European Economic Association*, 2, June, pp.569–606.
- 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.
- Chung, W., Disney, R, Emmerson, C. and Wakefield, M. (2008) ‘Public policy and retirement saving incentives in the UK’ in De Menil, G., Pestieau, P. and Fenge, R. (eds), *Strategies for Pension Reform*, Cambridge: MIT Press: *forthcoming*.
- Clark, T. and Emmerson, C. (2003) ‘The tax and benefit system and the decision to invest in a stakeholder pension’, Institute for Fiscal Studies *Briefing Note* No. 28, London.
- Department of Social Security (1998), *A new contract for welfare: Partnership in pensions*, (<http://www.dwp.gov.uk/publications/dss/1998/pengp/pdfs/pensions.pdf>), Cm 4179, London: DSS.
- Department for Work and Pensions (2002), *Simplicity, Security and Choice: Working and Saving for Retirement*, Cm. 5677, London: The Stationery Office (<http://www.dwp.gov.uk/consultations/consult/2002/pensions/gp.pdf>).

- Department for Work and Pensions (2006a), *Security in retirement: towards a new pensions system*, Cm 6841, London: The Stationery Office, (<http://www.dwp.gov.uk/pensionsreform/whitepaper.asp>).
- Department for Work and Pensions (2006b), *Personal accounts: a new way to save*, Cm. 6975, London: The Stationery Office, (http://www.dwp.gov.uk/pensionsreform/new_way.asp).
- Disney, R. and Emmerson, C. (2005) 'Public pension reform in the United Kingdom: what effect on the financial well being of current and future pensioners?', *Fiscal Studies*, 26, March, 55–82.
- Disney, R., Emmerson, C. and Wakefield, M. (2001) 'Pension reform and saving in Britain' *Oxford Review of Economic Policy*, 17, Spring, pp.70–94.
- Disney, R., Emmerson, C. and Wakefield, M. (2007) 'Tax reform and retirement saving incentives: Evidence from the introduction of Stakeholder Pensions in the United Kingdom', Institute for Fiscal Studies *Working Paper* W07/19, London. (http://www.ifs.org.uk/publications.php?publication_id=4093).
- Emmerson, C. and Tanner, S. (1999) 'Stakeholder pensions', Institute for Fiscal Studies, *Briefing Notes* BN01, (http://www.ifs.org.uk/publications.php?publication_id=1761), London.
- 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.
- Journal of Economic Perspectives* (1996) Symposium on 'Government incentives for saving', vol. 10, Fall, pp.73–138.
- Milligan, K. (2003) 'How do contribution limits affect contributions to tax-preferred saving accounts?' *Journal of Public Economics*, 67, 253–281.
- Oliver, Wyman & Company (2001), *The Future Regulation of UK Savings and Investment*, London. Executive summary is available at (<http://www.abi.org.uk/oliverwymreport.pdf>).
- 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>).
- Pensions Commission (2005), *A New Pension Settlement for the Twenty-First Century: The Second Report of the Pensions Commission*, London: The Stationery Office (<http://www.pensionscommission.org.uk/publications/2005/annrep/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.
- 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.

Appendix Table A.1. Sample sizes for analysis in Table 3.

	Treatment		Control	
	Before	After	Before	After
All	33,987	36,619	19,159	21,062
Positive earners	19,101	21,002	19,159	21,062
Zero earners	14,886	15,617	19,159	21,062
Positive earners, singles				
All	5,353	6,154	3,956	4,303
Single women	3,184	3,748	1,638	1,980
Single men	2,169	2,406	2,318	2,323
Positive earners, couples				
All	13,748	14,848	15,203	16,759
Couples women	9,666	10,327	3,969	4,562
Couples men	4,082	4,521	11,234	12,197
Couples – zero earning partner	3,163	3,405	4,105	4,453
Couples – low or mid earning partner	6,085	6,712	7,250	7,916
Couples – high earning partner	4,500	4,731	3,848	4,390
Couples – low education partner	9,007	9,386	7,869	8,368
Couples – mid education partner	2,553	2,950	3,502	3,786
Couples – high education partner	2,156	2,473	3,777	4,540

Endnotes

¹ For the UK, Oliver, Wyman & Company (2001), the Pensions Commission (2004; 2005) and Department for Work and Pensions (2002; 2006a) expressed concerns over the ‘adequacy’ of retirement saving. Banks *et al* (2005) are more sanguine. Much the same debate has occurred over a long period in the United States: see, on the one hand, Bernheim (1992) and, on the other, Engen *et al* (1999) and Scholz *et al* (2006). The same issue has also arisen in countries such as Australia and New Zealand.

² See Disney, Emmerson and Wakefield, 2001, Attanasio and Rohwedder, 2003, and Attanasio, Banks and Wakefield, 2005 for some findings on the impact of earlier pension reforms on saving in the UK.

³ ‘Occupational pensions’ are, in UK parlance, employer-provided pension plans.

⁴ Under the United Kingdom’s social security programme, individuals can choose to opt-out of the second tier of the public programme. In the case of a personal pension, the DWP pays part of the social security contribution (which is a proportion of earnings within a given range) into the opted-out pension account in return for the individual forgoing that part of second-tier pension benefit that would have accrued had they remained ‘contracted-in’. The employee can then contribute further amounts to their account, accruing the tax reliefs described in the text. The individual is not required to opt out of the social security programme to open a defined contribution pension plan. For further details on opting-out incentives, see Chung *et al* (2007).

⁵ The need to target middle-earners had been queried at the time, since this group already had high rates of pension coverage (see Disney, Emmerson and Tanner, 1999 and Table 2 in the text). For the 30% who were not covered, unstable incomes and less accessible savings made pension saving less attractive (see Banks, Blundell, Disney and Emmerson, 2002).

⁶ ‘National Insurance Contributions’ (NICs) is the name given to social security contributions in the UK.

⁷ Clark and Emmerson (2003) discuss other features of the tax treatment of Stakeholder Pensions. An even more sweeping reform to the ceilings on pension contributions was introduced in April 2006. Under these provisions designed to unify the tax regime for all types of private pensions – whether of the defined benefit or defined contribution form – there is an annual limit on contributions of 100% of earnings up to a ceiling of £215,000 (with the floor of £3,600 remaining) and a new lifetime limit on the value of the pension fund of £1.5m, rising over time. A further change which affected *all* potential contributors, was the abolition of the ‘carry-over’ provision for unused limits. Milligan (2003) examines the implications for intertemporal behaviour of a similar reform in Canada.

⁸ As a cross check, we examined responses from the General Household Survey (GHS), which asks somewhat different questions, primarily about coverage and membership, and also looks at aggregate data on pension scheme membership and contributions from Inland Revenue sources. The FRS provides detailed information on contributions, unlike the GHS. Both household surveys give significantly lower numbers for pension coverage and (more significantly, in the case of the FRS) for contributions than aggregate data from the Inland Revenue, perhaps reflecting under sampling in household surveys of contributors who make large contributions (i.e. the rich) and of other groups who may be contributing but are not asked about their contributions in the survey. However, it can be noted that aggregate data on total pension saving has been heavily revised downwards in recent years (although this applies more to data reported by the Office of National Statistics).

⁹ We gross up weekly earnings data to provide these annual earnings bands. This inevitably produces measurement error – for example some people will wrongly be attributed ‘zero’ earnings for the year based on current zero earnings. In addition, the Green Paper sometimes refers to ‘£20,000’ and sometimes to ‘£18,500’ as the highest income of ‘middle earners’. In general we work with the latter definition in the FRS data, revalued over time in line with average earnings growth in the sample.

¹⁰ It is important that our modelling strategy allows for the discrete nature of our outcome variable: a linear probability specification may very likely lead to the prediction that those with zero or very low earnings have a negative ‘probability’ of saving in a pension.

¹¹ For example, for an individual aged 35 or younger, this variable takes the value 1 if gross earnings are less than £20,571. This number is derived from the fact that before April 2001 individuals in this age range could contribute no more than 17.5% of their earnings to a pension, but after April 2001 this

limit became the maximum of 17.5% of earnings or £3,600. £3,600 is (to the nearest pound) 17.5% of £20,571. Similar values are constructed for individuals in other age bands.

¹² The variable I_t is not entered independently in this regression since the time dummies capture this variation.

¹³ In addition to the points raised in the text about the common trends assumption, it is also the case that the ‘marginal effects’ on interaction terms in non-linear models that are automatically generated by software packages (in our case by STATA version 9.2) often do not give a true measure of ‘interaction effects’. For more details see Ai and Norton (2003).

¹⁴ Time dummies are no longer included since separate probits are run for those observed pre-reform and those observed post-reform. Similarly analysis is done separately by whether or not the individuals would have received an increase in their private pension contribution limit and since this depends on earnings these are also excluded from this specification. Partners earnings, where relevant, is included.

¹⁵ Despite the similarity to the linear case, the nonlinear assumption exploited here entails two additional restrictions on the nature of the error terms: only group effects are allowed for and the groups being compared are assumed to have the same residual variance. See Blundell *et al, ibid*, p.580.

¹⁶ This raises another small measurement issue that applies to those with no earnings who, pre-2001, should not have been contributing to a private pension and receiving tax relief – see the regime described in Table 1. Table 2 nevertheless suggests that we observe a few individuals who are contributing pre-2001, which arises (we surmise) because they had some earnings during the year even though we observe no current earnings at the time the individual was surveyed for the FRS. Under certain assumptions (notably concerning the volatility of earnings), this measured proportion of take-up among those with no earnings pre-2001 may be taken as an upper bound on the measurement error involved in grossing up weekly earnings to obtain annual earnings, both before and after the reform.

¹⁷ Data limitations rule out using earlier years before 1999–2000 in calculating the ‘pre-treatment’ probits.

¹⁸ With the bootstrapped 5% confidence interval, we can only confirm that the effect lies between -1 and $+7$ percentage points.

¹⁹ The common trends assumption could still be valid even if characteristics changed differentially over time across the control and treatment groups, but these specific characteristics did not affect the take-up of private pensions, or if differential changes in characteristics just happened to cancel out in their *net* impact on average take-up among the control and treatment groups.

²⁰ For zero earners, we have evidence that fewer low educated people are likely to be observed after the reform. The coefficients on the interaction terms are jointly different from zero at conventional levels of statistical significance ($\chi^2(18) = 34.70$, $\text{Prob} > \chi^2 = 0.01$). Inspection suggests that our control group has a higher proportion of less educated people after the reform – that is, the overall fall in the fraction with low education was disproportionately concentrated among the treated group. Since people with more schooling are more likely to contribute to a pension, this facet tends to raise the measured treatment effect but in fact we observe no significant treatment effect when looking at the impact of the policy change on those without any earnings, as reported in Table 3.

²¹ Available on request from the authors.

²² This amount is for contributions reported to the FRS, then grossed up using survey weights to give a national figure.

²³ Notably this reform again changed the limits on the amount of tax-relieved contributions that individuals can make in such a way that the vast majority of individuals will now be able to make tax-relieved contributions equivalent to their full year’s earnings each year.