

TAX REFORM AND RETIREMENT SAVING INCENTIVES:
EVIDENCE FROM THE INTRODUCTION OF
STAKEHOLDER PENSIONS IN THE UK

Richard Disney
Carl Emmerson
Matthew Wakefield

Tax reform and retirement saving incentives: Evidence from the introduction of Stakeholder Pensions in the UK

by

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

Abstract

Faced with ageing populations, OECD governments are seeking policies to increase individual retirement saving. In April 2001, the UK government introduced Stakeholder Pensions – a low cost retirement saving vehicle. The reform also changed the structure of tax-relieved contribution ceilings, increasing their generosity for lower earning individuals. We examine the impact of these changes on private pension coverage and on contributions to personal pension accounts using individual level micro data. We use a difference-in-differences strategy, and where necessary our estimator is modified to allow for dichotomous outcomes. The results suggest that the change to the contribution ceilings affected both coverage rates and contributions to private pensions among lower earnings individuals, especially among women, and those in couples.

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: Evidence from the introduction of Stakeholder Pensions in the UK

1. Introduction

Demographic pressures are putting pressure on social security programme finances in many countries. In this context, a common policy target is to encourage households to increase their own private retirement saving which, it is hoped, will ameliorate both political expectations and cost pressures on the public budget. 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 work and do not work. There has been a substantial debate around this question in the United States (see, *inter alia*, Bernheim and Scholz, 1993; Poterba, 1994; Journal of Economic Perspectives, 1996), 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 (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).

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 total (i.e. public plus private) impact of incentives on saving. Several changes to the UK’s tax regime governing retirement saving have been implemented in the last two decades but, from an evaluative point of view, 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 standard evaluation techniques 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 personal 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 coverage while Section 5 examines the impact of the tax reform on the level of contributions paid into personal pensions. Section 6 provides a brief conclusion.

2. Stakeholder Pensions

Stakeholder Pensions were introduced in April 2001 after being proposed in a government Green Paper (Department of Social Security, 1998). Like all personal pensions, and some occupational pension schemes², Stakeholder Pensions are ‘defined contribution’ schemes, in that pension benefits depend on the accumulated value of the fund. They 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.

Since April 2001, companies employing at least five people that do not offer occupational pensions have been required to: nominate a Stakeholder Pension provider after consultation with employees; provide employees with information on Stakeholder Pensions; and to arrange for any contributions that employees wish to make to be channelled to the nominated pension provider. Neither employees nor employers are compelled to contribute to a Stakeholder Pension and indeed firms employing less than five people were completely exempted from the requirement to nominate a provider. Stakeholder Pensions have a simple charging structure: an initial annual cap on charges was set at 1% of the fund, with no charges either upfront or on withdrawals from the fund.³ Moreover, contributors can start and stop contributing at any time and schemes

have to accept all contributions of £20 or more. Compulsory minimum standards are intended to provide greater uniformity between Stakeholder Pensions offered by different pension providers than previous pension arrangements. By making pension providers offer a relatively uniform product it was hoped that there would be less need for individuals to seek independent financial advice before taking out a Stakeholder Pension, and that a more transparent system would encourage greater competition to reduce charges. Further details of the reform are described in Chung *et al* (2007).

The Green Paper 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 pension schemes are open to everyone. High income earners, it was assumed already had access to other retirement saving instruments, whereas lower earners were assumed 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 Indeed 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 low-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, and which 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's direct tax system is individual-based and contributions whether by employer or employee to defined contribution pension plans obtain tax relief (deferral) against income tax and, in the case of employer contributions, National Insurance⁶ relief, up to a ceiling of earnings. The contributions are made net of tax, with the government then contributing the equivalent basic rate income tax to the individual's account. Higher rate income taxpayers can go on to claim more relief in line with their higher marginal income tax rate. Returns are broadly tax-exempt and pensions are then taxed 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 low earning individuals, 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

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

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



Emmerson and Tanner (1999) also 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 a subsequent defence of the Stakeholder Pension reform in response to the apparently disappointing take-up rate, especially among middle earners (see the next section), the Treasury continued to emphasize the targeting of groups in the reform defined by earnings levels without explicitly referring to the change in tax reliefs, but subtly modified the way that the earnings groups were defined. Rather than referring to the impact on those earning between £9,000 and £18,500, the Treasury referred to all those earning *less than* £20,000 a year, thus including lower as well as middle earners:

“Stakeholder pensions have been a success for their target audience and their introduction has made a good value personal pension vehicle widely accessible for the first time....Stakeholders are being bought by their target market of moderate earners; around two-thirds of those taking up the pension are workers earning less than £20,000 a year....” (Treasury evidence to House of Commons Treasury Select Committee, 2006)

Despite this defence, the apparent failure of the reform to increase overall private retirement saving has contributed to the decision to enact yet another change in the UK's retirement saving regime, with the 2006 proposal for a new National Pension Savings Scheme (NPSS). Under this proposal employees offered the opportunity to join an occupational pension scheme will be automatically enrolled into that scheme – that is they will have to choose not to join a scheme rather than having to choose to join a scheme. Those employees whose employer does not offer an occupational pension arrangement will be automatically enrolled into the NPSS and, unless they choose otherwise, will contribute 4% of their earnings (between a floor and a ceiling), to this defined contribution scheme (their employer will contribute a further 3% and another 1% in tax relief will come from the Government). The rationale for this change towards greater prescription of retirement saving in 2006 also 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.

The introduction of Stakeholder Pensions therefore 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 less publicised at the time but also having a potential impact on retirement saving incentives. We briefly show in the next section (and in more detail in Chung *et al*, 2007) that the former aspect (i.e. the 'targeting' of the reform on middle earners) had no effect on take-up of private pensions and therefore focus the rest of the paper on the impact of the change in contribution limits on pension take-up and contributions among the affected group. In this latter respect, our line of argument follows closely the literature in the United States 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, and Gale and Scholz, 1994).

3. Empirical analysis

3.1. Data sources and descriptive analysis

We investigate the determinants of the household decision to take-out a pension and to save for retirement 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 low 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:	%	%	%	%	<i>% point</i>
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	%	%	%	%	<i>% point</i>
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	%	%	%	%	<i>% point</i>
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, these combined findings from 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 low 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 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 the table.

The discussion in Section 2 combined with these results lead us to think that it might be the change in the contribution limits that had an impact on coverage rather than the targeting of particular earnings bands, as highlighted in the two quotations in that section 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 relative differences in the pension behaviour of those who did, or who would have, received an increase in their pension contribution limit to those who did not (as illustrated in Panel C of Table 2).

3.2. *Modelling the pension take-up decision*

We wish to test formally whether the change in contribution limits incorporated into the Stakeholder Pension reform in April 2001 affected pension saving decisions within households. This reform can be considered as a policy ‘treatment’ in that it affected only some individuals (i.e. those with earnings below the limits at the time of the reform), and so our basic method of analysis is the difference-in-differences exercise that we describe below.

To formalise this test we write a general model of retirement saving in which Y_{it}^* is the outcome of the retirement saving decision of individual i and time t , which we relate 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):

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

In this section our outcome variable is a dichotomous indicator of whether or not an individual saves in a private pension at a particular point in time. (In the next

section we consider the question of how much is saved in the private pension). If this dichotomous outcome is Y_{it} , then we may 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 that we wish to test here is whether or not 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 counterfactual assumption required to allow differential differences to be attributed to the effect of the policy 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 be thought of as giving 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) but for a transformation of the distribution of the outcome variable; 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 saying that in the absence of any ‘treatment’ the following would hold¹⁴:

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

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 it is 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

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

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 observables 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). Since 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 8. 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 the value of contributions post-April 2001 that were in excess of what the contribution limits would have been in the absence of the reform. In any event, such a measure of the impact could lead to an underestimate of the overall effect of the policy, since the large rise in the contribution limits may have triggered the take-up of a private pension because the higher value of tax-relieved contributions now outweighed the threshold costs of purchasing a pension when the individual was previously only able to make small values of tax-relieved contributions to their pension account. Thus the potential 'treated' in the analysis are not simply those individuals who we might calculate to be at the pre-reform contribution ceiling and who thereby benefited from the increase in contribution ceilings as a result of the reform, but 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. 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 by 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.

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



Note: As Table 2.
Source: As Table 2.

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 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 with age is much less marked.

4. Did the reform increase private pension coverage?

4.1. Empirical results

We estimate the model described in equations (1'), (2) and (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 (where appropriate) 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 with bootstrapped standard errors and sample sizes.

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: Standard errors calculated 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.

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, statistically insignificant). 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 a 5% confidence interval. 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’ are more likely to be younger among those with positive earners (see Figure 2) and therefore more likely to have greater schooling than the ‘control’ group, *ceteris paribus*.

Using partner’s earnings and education, we note that the largest effect by earnings group (and indeed the only statistically significant impact) is on positive earners with a zero earning partner. We do not find clear-cut evidence that low 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 to increase pension coverage among the low 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 between the school leaving age and 18 years exhibit the largest estimated positive responses to the reform (+7.3 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 that the change to the contribution limits had a statistically significant impact; some of the proportionate effects on coverage among sub-groups are quite substantial. In particular, the change in limits seems to have particularly boosted pension coverage among 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 and among disaggregated groups. As we have argued above, the lack of effects on middle earners may well be due to the details of the policy reform package which focussed its ‘treatment’, measured as the change in financial incentives due to the reform of tax relief, on lower earners.

4.2. Alternative explanations and sensitivity analysis

This section considers alternative explanations for our results. In particular, it considers the plausibility of the ‘common trends’ assumption that lies at the heart of the ‘difference-in-differences’ model in its present context, and undertakes some sensitivity analyses of the results.

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 that among those who did not receive an increase in their contribution limit as a result of the reform. 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 had largely 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 in a less favourable position to reduce their pension saving as a response to the increase in housing equity. 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.

In addition, if instead we focus on changes that might affect the ‘treatment’ 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. As described briefly in Section 2, the period saw the replacement of SERPS, the second tier pension, by the State Second Pension (S2P – introduced in April 2002 but 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, known as the Minimum Income Guarantee (MIG), was formally indexed to earnings rather than prices from April 1999 (unlike the rest of the pension programme),

so increasing its real value for low earners and reinforcing the disincentive to save for retirement. Analysis of these, and subsequent, trends suggests that, at least under the current regime, whereas replacement rates cohort-by-cohort for the public pension programme have already peaked for average earners, low 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 to understate take-up as a result of the change in contribution limits.

Another possible alternative explanation for the ‘treatment effect’ lies in changes in the composition of households within earnings bands. Suppose, for example, that earnings volatility grew over the period: at the beginning of the period, in this scenario, the rich were persistently rich and the poor persistently poor but that by the end of this period, there had been greater mobility so that the lower earners contained a larger fraction of transient rich, and *vice versa*. Suppose, too, that this was reflected in a growing disparity in *intra* household incomes over time, with a greater preponderance of low earners with high income spouses (and *vice versa*). If the rich, irrespective of current situation, tended to take-up private pensions while the poor did not, a greater number of transitions would lead to some slight convergence in take-up rates, which is in fact what we observe in Table 2 and the subsequent analysis.

Given that the FRS is not a panel, we cannot test for greater earnings volatility over the period (although we find little evidence of it from some experiments with the British Household Panel Survey and the rotating panel element of the Labour Force Survey). So we test the 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:

$$I_{it} = \theta L_{i=1} + \gamma_1' X_{it} + \alpha_1' X_{it} L_{i=1} + \varepsilon_{it} \quad (4)$$

Where I_{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 $[\gamma_i]$ tell us whether or not 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 $[\alpha_i]$ tell us whether or not any changes in characteristics over time occurs differently between those in the control group and those in the treatment group. It is this second test which is important for our analysis – our results will be less likely to imply the policy effects we wish to identify if changes in the characteristics of those observed in the control and treatment group are occurring differentially by the relevant 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 is 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 show that our sample is, on average, slightly older and slightly better educated in the period after stakeholder pensions were implemented. This is not surprising given that over time the UK population is ageing and that successive cohorts are achieving more education qualifications.

There is however little evidence of a differential change in characteristics between those in the control group and those in the treated group, especially for the treatment group with some earnings. For positive earners, the coefficients on the interaction terms [i.e. vector α_i] 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 some of our results where we pool men and women.

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 3.²⁰

Our final sensitivity test attempts to handle the possibility 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’, 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 age group 36-45, which faced a single contribution ceiling of 20% of earnings in the pre-Stakeholder Pension regime, contains a 10 year as opposed to a 5 year bandwidth and which, conveniently, contains around 50% individually treated 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 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, 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.

5. Did the reform increase private pension saving?

5.1. Identifying the ‘treatment’ group

In this section, we provide some evidence on the impact of the increase in the ceiling on tax-relieved contributions on the *level* of contributions among the treated group. However, the evidence is somewhat tentative insofar as there are further possibilities for mis-measurement and misclassification of the ‘treatment’ in relation to

contribution levels in addition to those described in determining the pension take-up effect. Since misclassification of treatments and controls generally leads to a downward bias in estimated outcomes, it is tempting to argue that any significant result that we find is a lower bound to the ‘true’ effect; however we prefer simply to be cautious.

A priori, the relevant group can be defined as individuals who are contributing to, or could contribute to, a personal pension, with the ‘treatment’ defined as whether such individuals received an increase in their contribution limit. However the Family Resources Survey only measures the value of individuals’ private pension contributions to personal and stakeholder pensions. Therefore we have to exclude those who report that they contribute to an occupational pension from the analysis of pension contributions. One potentially important issue is that by excluding this group we may exclude some individuals from the treatment group who misreport that they have an ‘occupational pension’ when in fact they have an employer-sponsored stakeholder pension from an independent insurer (see Section 2). We do however test for evidence of whether or not this exclusion of occupational pension members affects the pre- and post-reform composition of the treatment group relative to the control group, the results of which are presented below.

There are other reasons why we may impart misclassifications and measurement errors to our derivation of the treatment group using the FRS data. We also do not know about the value of any employer’s contribution to an individual’s private pension. This may lead us to mis-measure contributions to personal and stakeholder pensions.²² Second, contributions may be more ‘lumpy’ than earnings and the point that only weekly amounts are measured in the FRS while the tax relief limits are based on annual earnings becomes even more pertinent. Note that we do not use the contributions data to determine whether the individual is treated or not – this is solely a function of his or her age and earnings – nevertheless the change in measured contributions may be biased by this lumpiness. Fortunately we can exclude most individuals in this position since ‘one-off’ or lumpy contributions can be identified from the data and are excluded from the sample.

We apply a linear treatment model, which captures an overall treatment impact that averages across two distinct responses – a take-up effect and a contribution effect conditional on take-up. Unfortunately, without panel data it is difficult to separate out these two effects.

5.2. Contribution levels: empirical results

We estimate a linear ‘difference-in-differences’ models of weekly contributions to pensions before and after the tax relief reform using the FRS as before. As mentioned previously, we exclude individuals who reported (or whose partner reported) that they contributed to an occupational pension²³ and we may thereby exclude people who are affected by the reform but who were misclassified from their response to the question of ‘type of pension’.²⁴

In Table 4, we provide the results of this exercise in several ways. We estimate the model for singles (Column 1) and then estimate the model for couples in two ways. First (Column 2), we estimate a model of the joint contributions of the couple, with the ‘treatment’ being whether either (or indeed both) members of the couple received the increase in contribution limits. Second (Column 3), we estimate a model of the contributions of each respondent in the couple, whereby we allow the impact on contributions of a limit increase for the respondent to vary by the earnings of his or her partner.

Before analysing the treatment effects, there are several common features to the various specifications in Table 4. Relative to the default age group (21-35), average contributions to pensions peak in the age range 51-55, both for respondents and the age of their partners (where appropriate). The sex of the respondent makes no difference, and lower schooling levels reduce contribution rates. Contributions are strongly related to the level of earnings, although in some cases the relationship is non-linear in polynomials. Most strikingly, there is a *downward* trend in contribution levels among the control group over time: the reduction from 1999 to 2002 is over £1 per week for singles, rather more for individuals in couples and indeed is estimated at £9 per week for couples jointly. The positive calculated ‘treatment’ effects on the change in contribution limits should be set against this fact.

Table 4:
Impact of reform on weekly pension contributions (£): Difference-in-differences

Method: OLS	(1) Single persons		(2) Couples (total family contribs)		(3) Couples (individual contribs)	
	Coeff.	Std.Err.	Coeff.	Std.Err.	Coeff.	Std.Err.
<i>Respondent</i>						
Age 36 to 45	1.220*	0.163	1.565*	0.260	0.436	0.267
Age 46 to 50	2.551*	0.229	2.424*	0.415	1.458*	0.383
Age 51 to 55	2.300*	0.226	3.469*	0.388	2.312*	0.418
Age 56 to 60	1.645*	0.238	2.681*	0.382	1.774*	0.469
Age 61 to 64	1.321*	0.376	2.148*	0.542	1.346*	0.619
Male	-0.139	0.129	0.010	0.011	0.294	0.212
Schooling years = 15/16	-0.843*	0.171	-1.524*	0.401	-1.642*	0.286
Schooling years = 17/18	-0.238	0.205	-0.464	0.542	-1.012*	0.317
Earnings	0.0004*	0.00001	0.0006*	0.00001	0.0006*	0.00002
Earnings ²	-3.86e ⁻¹⁰ *	-5.34e ⁻¹¹	-2.80e ⁻¹⁰	-9.83e ⁻¹⁰	-5.84e ⁻¹⁰ *	-1.19e ⁻¹⁰
Earnings ³	-3.91e ⁻¹⁷	-4.05e ⁻¹⁷	-2.78e ⁻¹⁶	-1.07e ⁻¹⁵	-6.64e ⁻¹⁷	-1.25e ⁻¹⁶
Zero earner	2.393*	0.184	3.835*	0.793	3.641*	0.248
<i>Partner</i>						
Age 22 to 35	n/a	n/a	-0.369	0.444	0.601	0.856
Age 36 to 45	n/a	n/a	1.191*	0.433	1.794*	0.883
Age 46 to 50	n/a	n/a	2.052*	0.539	1.820*	0.927
Age 51 to 55	n/a	n/a	3.098*	0.500	2.083*	0.942
Age 56 to 60	n/a	n/a	2.493*	0.498	2.011*	0.968
Age 61 to 64	n/a	n/a	1.786*	0.645	1.623	1.053
Schooling years = 15/16	n/a	n/a	-1.511	0.405	0.085	0.286
Schooling years = 17/18	n/a	n/a	-0.477	0.546	0.523	0.317
Earnings	n/a	n/a	0.565*	0.073	0.159*	0.020
Earnings ²	n/a	n/a	-0.0003	0.001	-4.18e ⁻⁰⁶	0.0001
Earnings ³	n/a	n/a	-2.77e ⁻⁰⁷	-1.06e ⁻⁰⁶	-5.10e ⁻⁰⁸	-1.10e ⁻⁰⁷
Zero earner	n/a	n/a	3.842*	0.792	0.355	0.244
<i>Year effects</i>						
Year: 2000–01	-0.378*	0.177	-1.703*	0.508	-0.843*	0.234
Year: 2001–02	-1.015*	0.351	-8.626*	3.728	-1.482*	0.451
Year: 2002–03	-1.150*	0.350	-9.174*	3.716	-1.728*	0.449
Had limit increase? (Y/N)	-1.612*	0.332	n/a	n/a	1.253*	0.426
Limit increase? * post	0.793*	0.355	n/a	n/a	1.011*	0.473
Family limit increase?	n/a	n/a	-6.608*	3.222	n/a	n/a
Family limit increase? * post	n/a	n/a	7.335*	3.695	n/a	n/a
Partner's earnings * limit increase	n/a	n/a	n/a	n/a	-0.125*	0.020
Partner's earnings * post	n/a	n/a	n/a	n/a	-0.101*	0.024
Partner's earnings * limit increase * post	n/a	n/a	n/a	n/a	0.110*	0.026
R ²	0.1510		0.2208		0.2154	
No. of observations	21,894		26,132		26,573	
F(df, n)	228.91		n/a		227.69	

Notes: Data source: *Family Resources Survey*. Default category is woman aged 21 to 35, with tertiary education, in 1999–2000 whose income lies above the contribution limit. Constant estimated. Sample excludes people with only occupational pension, and those who report making a one-off lump sum pension contribution.

* = coefficient significant at 5% level. Standard errors in (2) are clustered at family level.

Column (1) of Table 4 suggests that, on average, treated single individuals contributed £1.6 less a week than controls and that, as a result of the reform, they increased their contributions by £0.8 per week. This would not be a large increase, but it should be noted that the mean *level* of contributions among single people per week to a private pension averaged only £0.7 per week during the period 1999 to 2002. This stems from the fact that only 10% of single people (once we have excluded those belonging to occupational pensions) contributed to a private pension in this period; among those that did contribute, however, the average contribution was £13 per week. Since Table 3 suggests that the reform increased coverage by almost 4 percentage points, it seems very likely that this average effect is an increased coverage effect rather than an increase in contributions among existing holders of private pensions. Note, too, that the average change in contributions attributable to the reform just about counteracts the fall in average contributions between the pre-reform period (1999–2000) and the post-reform period (2001–02); the net effect of the reform of contribution limits for single householders was to maintain contributions at their pre-reform level, whereas contributions among the control group fell over the period.

Column (2) of Table 4 considers the joint contributions of couples, thus the structure of coefficients is much the same as that for column 1. Here there is a much larger underlying fall in contributions among the control group. On average, households with at least one member affected by the limit change contributed £6.6 per week less than the control group and the effect of the reform was to raise contributions by treated households by £7.3 per week. Again natural benchmarks are the average contribution over the period among all households whether contributing or not (which is £5.7 per week) and among contributing households (which is £26 per week). This estimated treatment effect is just short of the increase in contributions needed to offset the fall in contribution rates among the controls before and after the reform (which is about £8 per week from Table 4).²⁵

The final results in Table 4, in Column (3) again examine couples, but look at individual contributions. *Ceteris paribus* now becomes important in interpretation of the treatment effects. An individual eligible for the treatment actually contributed more than a control, and on average increased his or her contributions as a result of the 2001 reform by almost £1 per week plus an additional £0.1 for every £1,000 of partner's earnings. So if the partner obtained earnings of £20k, the net effect of the reform was to increase the respondent's contributions by £3 per week. The total effects (analogous to

column 2) then depend on the composition of each family's income – if both partners were affected by the treatment and had earnings of this magnitude (as is certainly possible for households aged under 36), the combined increase in contributions from this calculation would be close to the average effect calculated in column (2). However, it should be borne in mind that the estimates of treatment effects, while clearly significantly positive, are not very precisely determined and may vary from specification to specification.

5.3. *Sensitivity analysis and aggregate implications*

This sub-section considers several issues concerning the interpretation of these results on contribution limits. First, we again test whether the results arise from possible composition changes in the treatment and control groups post-treatment which might affect our conclusions. We repeat the method described in equation (4) in Section 4.2 to test first whether the composition of the sample has changed pre- and post-treatment and second, whether this change in composition differs between the treated and the untreated. Note that we are now excluding the sub-set of households where one or more member contributes to an occupational pension plan from the analysis of contribution levels here. Formally, we easily can reject the possibility of compositional changes that differ across the treated and untreated for both earners ($\chi^2(18) = 21.51$, $\text{Prob}>\chi^2 = 0.25$) and zero earners ($\chi^2(18) = 23.91$, $\text{Prob}>\chi^2 = 0.16$). Inspection of coefficients again suggests weak evidence that the treated group of positive earners have a greater proportion of men in the post-treatment period, and that the treated zero earners tend to have more education in the post-treatment period.

One specific concern arises from comparing grossed-up FRS data to statistics from HM Revenue and Customs (HMRC) on employee contributions to Personal and Stakeholder Pensions. Inspection of data from HMRC²⁶ confirms that individual employee contributions to personal and Stakeholder Pensions fell slightly between 2000–01 and 2001–02 but that contributions rose sharply in 2002–03. We do not observe this last rise in the FRS data and the fall at the time of the reform seems to be larger than in HMRC data. This may be because pension contributions are disproportionately paid by sub-set of individuals who are underrepresented in the data, or because some individuals in the FRS misreport becoming a member of a Stakeholder Pension after 2001 as becoming a member of an occupational pension plan. Since we exclude members of occupational pensions from the analysis of this section, we may thereby underreport

post-reform contributions to such pensions. This would be a concern in measuring the treatment effect if we believed that these misreporting trends differed across the treated and the controls but of course we have no direct evidence on measurement error.

A second general concern would arise if we were to interpret the average increase in contributions as a ‘pure’ contribution effect rather than a coverage effect. There is a natural downward limit to the fall in contributions that could have occurred in the absence of the treatment since pension contributions cannot be negative. Many individuals in our sample of people who are not in occupational pensions do not contribute to a private pension at all – indeed the median weekly contribution is zero and mean contributions are around £1.50 per week. This is particularly pertinent for our treated group, as it is apparent from Table 4 that the treated group had a lower contribution level than the average respondent; individuals in the control group had a greater probability of contributing, and larger contributions which they tended to reduce over the period. Thus even in the absence of the policy reform we might have expected a smaller fall in the value of contributions among the treatment group relative to the control group simply because for some in the treatment group the fall in contributions would be effectively truncated at the point at which they stop saving in a pension. This means we have to be careful to consider both contribution and coverage effects when interpreting the results reported in Table 4.

The focus on individual coefficients other than the ‘treatment effect’ is slightly misleading because we should take account of other characteristics that are significant predictors of contribution rates. In particular, to test the validity of our estimate of the differential contribution changes of the treated and controls, we can gross up the Family Resources Survey and compare directly the reported contributions before and after the reform by the ‘treated’ and ‘control’ groups, and also the change in contributions for the ‘treated’ predicted by applying the average treatment effect reported in Table 4. In addition, of course, these grossed up figures give us some idea of the overall effect of the policy reform. The results of grossing up reported contributions for the two groups can be seen in Table 5.

These results, based on the aggregated contributions data in the FRS, confirm a difference in trend in total contributions (which may arise from a coverage effect, a contribution effect or a combination thereof) between the group affected by the change in the contribution limits and the group that were not. It can be seen that the change in

contribution limits did not offset the measured decline in contributions reported in the FRS, but that the decline in contributions might have been some 30% larger had it not been for the reform to contribution ceilings [(i.e. $0.05 / (0.21 - 0.05)$ in Table 5).. Applying instead the calculated average treatment effect for individuals who are single and in couples appropriately to our treated group, we get an increase in contributions of +£0.08 billion. This is somewhat higher than the estimate derived from simple grossing-up of reported contributions in the FRS but the discrepancy is likely related to the fact that the calculation from the average treatment effect does not correct for the conflation of contribution and coverage effects which we discussed previously.

Table 5
Estimates of Policy Impact from Grossed-up Contributions
(Two years 1999–2000/2000–01 v 2001–02/2002–03)

Group	Average annual contributions (£)	Grossed up by FRS weights (£ billion)
Treated: 2 year pre-SP	46.8	0.93
2 year post-SP	48.1	0.98
<i>Net effect</i>	<i>+1.3</i>	<i>+0.05</i>
Controls: 2 year pre-SP	600.5	2.83
2 year post-SP	522.0	2.62
<i>Net effect</i>	<i>-78.5</i>	<i>-0.21</i>

Note: Own calculations from FRS data. Note that the figures aggregate average annual and total contributions added over the two years preceding the reform and the two years after the reform.

6. 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 the contribution limits which, essentially, 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 recent policy developments on the probability of households engaging in retirement saving and on the amount of that saving.

Aggregate data suggest that the introduction of Stakeholder Pensions had little impact on the overall propensity to save for retirement. The numbers 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 2006²⁷, 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. In similar vein, we show that the level of contributions among those benefiting from the higher contribution limits did not fall in contrast to the rest of the sample who did not belong to occupational pension schemes. 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 perhaps ironic that the introduction of Stakeholder Pensions in 2001 was quickly written off as having had little impact on retirement saving when our evidence suggests that at least some of the associated changes had non-trivial effects on sub-sets of the relevant population and 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.
- 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.
- Chung, W., Disney, R., Emmerson, C. and Wakefield, M. (2007) '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 (2003) *Pensions 2002: Public attitudes to pensions and saving for retirement*, DWP Research Report 193, London: July.
- 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. (2006) 'Pension provision and retirement saving: Lessons from the United Kingdom', presented at SEDAP conference at McMaster University, Hamilton, Canada, November.
- 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.
- HM Treasury (2004), *Charge Caps For New Stakeholder Products*, Press Release, 17th June 2004, London: HM Treasury (http://www.hm-treasury.gov.uk/newsroom_and_speeches/press/2004/press_57_04.cfm).
- 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.
- Office of National Statistics (2005) 'Private pension contributions: Updated estimates, 1996-2004', *Economic Trends*, 622, 29-34.
- Oliver, Wyman & Company (2001), *The Future Regulation of UK Savings and Investment*, London. Executive summary is available at <http://www.abi.org.uk/oliverwymanreport.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; 2006) 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 ‘Occupational pensions’ are, in UK parlance, employer-provided pension plans.

³ In 2004, after lobbying from the finance industry, the Treasury increased this charge cap from 1% to 1½% for the first 10 years that a product is held. For more details see HM Treasury (2004).

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 social security 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 which applies to those with no earnings who, pre-2001, should not have been contributing to a Personal 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 standard error, we can only confirm at the 5% confidence interval that the coefficient 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 A necessary reminder is that zero (annual) earners could not take advantage of any tax relief on contributions prior to the reform, although we measure weekly not annual earnings.

²¹ They are available on request from the authors.

²² However, recent data confirm that, unlike occupational pensions, the bulk of contributions to personal and stakeholder pensions are from employees rather than employers: see Disney, Emmerson and Wakefield (2006).

²³ Excluding those who have an occupational pension reduces the sample of earners by around 40%.

²⁴ We also exclude the small number of individuals who explicitly reported that they made a ‘one off’ lump sum contribution to their pension since clearly the weekly data give no guide to their annual contribution.

²⁵ That is $(0.5 \times 8.626 + 0.5 \times 9.174) - (0.5 \times 1.703)$.

²⁶ <http://www.hmrc.gov.uk/stats/pensions>, especially Table 7.16 as revised in May 2006. It should be noted that these data have been revised on at least one occasion.

²⁷ 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.