Institute for Fiscal Studies

Laurence O'Brien

Working paper 23/10

The effect of tax incentives on private pension saving

Economic and Social Research Council

The Effect of Tax Incentives on Private Pension Saving*

Laurence O'Brien†

February 20, 2023

Abstract

In many countries, saving in a private pension is tax-advantaged to encourage saving for retirement; however, there is mixed evidence on the extent to which this increases saving. This paper estimates the responsiveness of private pension saving to tax incentives for employees in England and Wales using employer-reported panel data on pension contributions between 2005 and 2019. For identification, I exploit a kink in the income tax schedule where there is a large, discontinuous change in the marginal price of pension saving. In the earlier part of the sample period, 2005 to 2012, I estimate an intensive-margin elasticity of around -0.1 and an extensive-margin elasticity of -0.05, suggesting that employees do not respond strongly to this tax incentive to save. In 2013 to 2019, after the introduction of automatic enrolment into workplace pension plans, I find a lower average elasticity, consistent with those being brought into pension saving by this policy being passive savers.

Keywords: retirement saving; incentive effects of taxation
JEL classification: H2, H3

*Funding from the Nuffield Foundation under grant (WEL /FR-000000374), and co-funding from the Centre for the Microeconomic Analysis of Public Policy (ES/T014334/1) at the Institute for Fiscal Studies, is gratefully acknowledged. The Nuffield Foundation is an independent charitable trust with a mission to advance social well-being. The author would like to thank Jonathan Cribb, David Sturrock, Carl Emmerson, Cormac O'Dea and especially Rowena Crawford for helpful feedback and comments. The Annual Survey of Hours and Earnings is provided by the Office for National Statistics (ONS) and accessed through the Secure Research Service. The use of the ONS statistical data in this work does not imply the endorsement of the ONS in relation to the interpretation or analysis of the statistical data. Correspondence to laurence.obrien@ifs.org.uk
†Institute for Fiscal Studies and University College London

1 Introduction

In common with many countries around the world, saving in a private pension in the UK is relatively tax favoured to encourage saving for retirement. The government revenue foregone from this tax advantage is sizeable, and the benefits often accrue disproportionately to relatively higher earners. Given this, it is unsurprising that the tax treatment of pension saving is a common topic of public and policy debate. Key to understanding the merits of the current system, and the possible effects of any reforms, is knowledge of how saving would respond to changes in tax incentives. However, despite extensive research on the determinants of private pension saving in recent years (Choi, 2015), there is still uncertainty about the extent to which these tax subsidies raise total wealth accumulation (Bernheim, 2002; Gelber, 2011) and how they interact with other retirement saving policies.

In this paper, I estimate how responsive private-sector employees are to upfront tax incentives to save in a pension, focusing on a policy-relevant part of the tax system in the UK, and show how this responsiveness has changed since the introduction of automatic enrolment into workplace pensions. To produce these estimates, I use employer-reported panel data on private-sector employees for the period 2005 to 2019. I exploit a jump in the marginal income tax rate from 20% to 40% at around the 90th percentile of the income distribution that creates a sharp discontinuity in the upfront tax price of pension contributions to examine how pension saving responds to this price on both the intensive and extensive margin. To assess the extent to which automatic enrolment affected the responsiveness of pension saving to tax incentives, I repeat all the analysis for two time periods: 2005 to 2012, before the implementation of automatic enrolment, and 2013 to 2019, during which time many private-sector employees had to be automatically enrolled into workplace pensions.

I start by providing graphical evidence that both the intensive- and extensive-margin elasticities are small in both the time periods I analyse. On the extensive margin, I show that there is no discontinuous increase in the share of private-sector employees with strictly positive employee pension contributions at the point where the upfront tax price of pension saving increases discontinuously. On the intensive margin, I show that those employees who are saving in a pension are not manipulating their pension contributions in such a way as to bunch their taxable income at the kink in the tax schedule. These results implies little to no extensive- or intensive-margin responsiveness of pension saving to tax incentives.

I confirm the lack of responsiveness suggested by the graphical analysis using a panel regression approach. Estimation poses two main identification challenges. First, the tax price of pension savings is endogenous, since employees can reduce their taxable income, and marginal tax rate, by increasing their pension saving. I address this by employing the

standard first-pound price instrument (Feldstein and Taylor, 1976). The second challenge is the possible simultaneous choice of earnings and pension saving; I address this by using individual-employer fixed effects to restrict identification to changes in the tax price for individuals working for an unchanged employer. In other words, identification comes from employees who move from one tax bracket to another while working for the same employer.

For the period 2005 to 2012, before automatic enrolment into workplace pensions, the estimate of the intensive-margin price elasticity is -0.1, while the estimated extensive-margin elasticity is -0.05, and statistically insignificantly different from zero. The total elasticity estimated for this time period is therefore -0.15, which is economically small. For the period 2013 to 2019, during which time automatic enrolment was being rolled out or was in place, estimates of both elasticities fall to essentially 0.

These results suggest that individuals around the kink in the income tax schedule that I examine – which is just above the 90th percentile of the income distribution – do not on the whole respond strongly to this tax incentive to save. The reduction in the intensive-margin elasticity since the introduction of automatic enrolment is consistent with those being brought into pension saving by the introduction of automatic enrolment being (even) more passive savers than those saving in a pension prior to the introduction of that policy. These new savers are potentially precisely the group that policymakers might worry are undersaving for retirement, but the results in this paper suggest that using tax incentives to increase their saving may be particularly ineffective.

This paper builds on a growing literature that has studied how private pension saving responds to tax incentives. Early contributions in the U.S. focused on whether tax-advantaged savings plans, such as IRAs and 401(K) accounts, encouraged greater saving, or whether the balances accumulated in these accounts were simply shifted from other assets. Early results were mixed (Poterba et al., 1995; Engen et al., 1996; Hubbard and Skinner, 1996), but Gelber (2011) finds becoming eligible for a 401(K) account does cause a significant increase in 401(K) saving, but finds little evidence of an increase in total saving, perhaps due to a lack of precision.

In general, more recent research suggests that tax incentives for retirement saving do lead to an increase in saving in subsidised accounts, but in many cases this is partly offset by a reduction in saving in other accounts or an increase in debt. For example, Chetty et al. (2014) show that a decrease in a subsidy for contributing to capital pension accounts for individuals in the top income tax bracket in Denmark did lead to a sharp decrease in capital pension contributions; however this decrease was driven entirely by just 19% of contributors, and nearly all the reduction in capital pension contributions can be accounted for by higher saving in other accounts. Andersen (2018) shows that an introduction of a cap on tax-favoured retirement contributions in Denmark also led to a significant reduction in retirement contributions, with this being fully offset by both

an increase in other types of savings and a decrease in gross debt. Similarly, Lavecchia (2019) finds that the introduction of a tax-advantaged saving account in Canada leads to an increase in saving in these accounts, but little to no new total saving due to substantial crowding-out.

However, other studies find rather larger effects of tax incentives on total saving. For example, Goodman (2020) exploits a discontinuity in retirement contribution limits based on date of birth in the U.S. He finds that constrained individuals' retirement saving does increase in response to a higher contribution limit, and can rule out crowd-out greater than approximately 0.38 to 0.57 at the 95% confidence level. Chan et al. (2022), Christensen and Ellegaard (2022) and Messacar (2018) similarly find substantial increases in retirement saving with only partial crowd-out among older workers making large contributions in Australia, middle-income workers in Denmark and unionised workers in Canada, respectively.

This paper contributes to this literature in two ways. First, in contrast to much of the recent literature, the tax incentive to pension saving that I analyse only has a very small to null effect on pension contributions. In other words, the tax incentive does not lead to a large increase in saving even in the "subsidised" account. This suggests that the institutional setting and the salience of the tax incentive can have a large effect on estimated elasticities, as pointed out in Chetty et al. (2013). Second, I highlight how tax incentives for retirement saving interact with another major retirement saving policy: automatic enrolment. Automatic enrolment has been shown to have a large impact on pension participation in a variety of settings (Madrian and Shea, 2001; Cribb and Emmerson, 2020), leading some to conclude that savings policies focusing on behavioural "nudges" are perhaps more effective than tax subsidies (Friedman, 2017). However, to date no paper has analysed how these two policies interact. I provide evidence that tax price elasticity of saving is lower after the introduction of automatic enrolment into workplace pensions in the UK, and in fact insignificantly different from zero, suggesting it is important to consider the two policies in tandem to assess their effectiveness.

The remainder of the paper is organised as follows. Section 2 describes the institutional context, while Section 3 describes the data sources. Section 4 contains graphical evidence of how pension saving responds to tax incentives, with the empirical methodology and main results outlined in Section 5. Section 6 concludes with a discussion of the implications of the empirical results.

2 Institutional background

2.1 Private pension saving in the UK

The current UK public pension system provides almost all older individuals with a flat-rate benefit, irrespective of earnings, that amounts to just under 30% of median earnings. Most individuals must therefore save additionally in private pensions if they want to smooth their living standards in retirement.

While pensions can be arranged by individuals, the vast majority of pension saving for employees is facilitated by employers. For public-sector employees these pensions are typically 'defined benefit' (DB) in nature, while for private-sector employees, workplace pensions are now predominantly 'defined contribution' (DC) plans. Employees in DC plans can typically choose their level of contributions (which may be influenced by matching arrangements provided by some employers). Some DC pension plans are run by the employer directly, while others are run by external pension providers, with the employer facilitating membership; I refer to these two types of pension as 'Occupational' and 'Other' DC pensions, respectively. In 2005, the start of the period I examine, 44% of private-sector employees were in a private pension: 24% in a DB plan, 9% in an occupational DC plan, and 12% in other DC plans.

Between 2012 and 2018 the UK government rolled out a policy of automatic enrolment into workplace pensions. This means employees are automatically enrolled into a pension and have to choose to opt out, and if they remain in the pension making minimum contributions then they are also eligible for some minimum contributions from their employer. This reform substantially increased pension participation (Cribb and Emmerson, 2020, 2021): in 2019, the end of the period I examine, 79% of private-sector employees were in a private pension (12% in a DB plan, 24% in an occupational DC plan, and 44% in other DC plans).

2.2 Tax treatment of pensions

The UK income tax system treats pensions as deferred earnings. This means that contributions are not subject to income tax and returns are not taxed, but income tax must be paid on pension income when it is drawn in retirement. Pensions are also relatively tax favoured compared to other forms of income because employer contributions to pensions are not subject to the payroll tax (national insurance), 25% of all pension saving can be withdrawn free of income tax in retirement, and pensions are taxed relatively lightly on death.[1]

[1]There are some restrictions on the amount that can be saved into a pension free of income tax, namely the annual and lifetime allowances. However, these are set at high enough levels to be essentially irrelevant for the employees in the analysis sample. Since 2014, for employees in the income range I analyse, the annual allowance limits tax-free contributions to £40,000 per year. Prior to 2014, the

The non-linearity of the income tax schedule mean that there are sharp discontinuities in the upfront income tax relief on pension contributions. Figure 2.1 shows the income tax schedule in 2019-20 in England and Wales for incomes up to £100,000 - the schedule for earlier years I analyse is very similar (Scotland is not included in the analysis as it has a different income tax schedule). Income tax is based on individual income. Individuals have a personal allowance, on which no income tax is due. On earnings above that but less than the 'higher rate threshold', individuals pay the basic rate of income tax, which was 22% until 2007-08 and 20% since then. For earnings above the higher rate threshold a higher rate of income tax of 40% is due. Since 2010-11 there have also been even higher effective marginal income tax rates on those earning £100,000 or more; however, throughout the analysis I focus on individuals earning below this level.

Figure 2.1: Income tax schedule 2019-20 in England and Wales (up to £100,000)

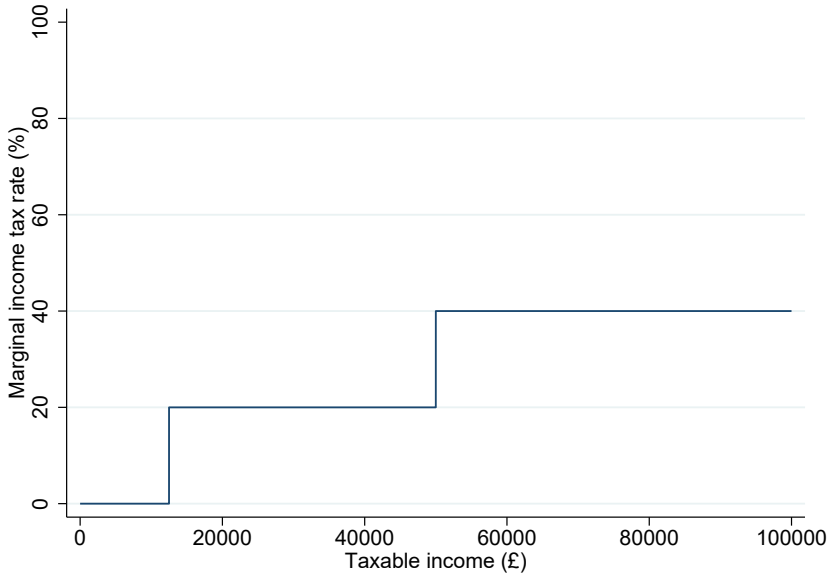

Source: https://www.gov.uk/income-tax-rates

Figure 2.2 shows that the real level of the higher rate threshold varied during the analysis period, fluctuating between £45,000 and £55,000. this figure also shows that the higher rate threshold lies towards the top of the income distribution, with about 6 to 9% of adults earning more than the higher rate threshold in the UK, depending on the year. Given that I estimate a local treatment effect of the responsiveness of pension saving around the higher rate threshold, the estimate is relevant to well-off, but not super-rich, individuals by UK standards.

annual allowance was even higher. The lifetime allowance, set at £1,073,100 in 2021-22, determines the maximum total value of an individual's private pensions before attracting high tax rates.

Figure 2.2: The level of the higher rate threshold, and the proportion of adults in the UK with income greater than it, over time

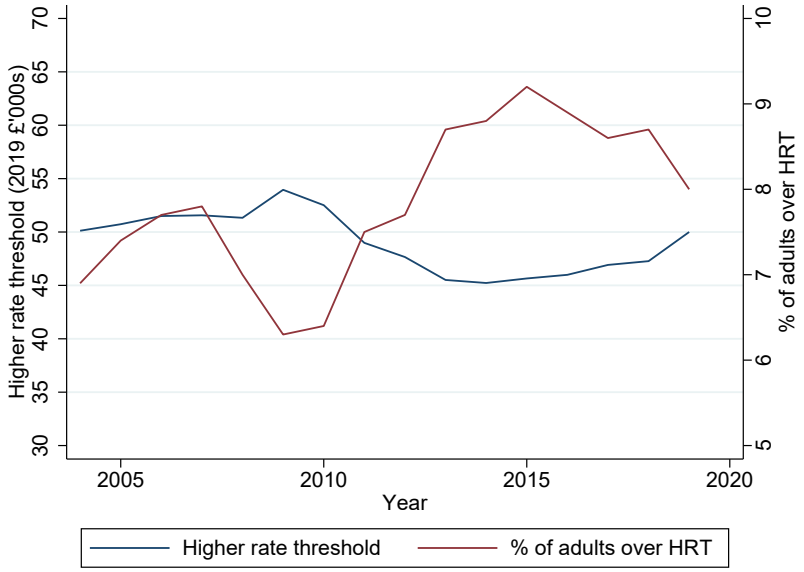

Notes: The higher rate threshold is the level of taxable income at which the UK marginal income tax rate increases to 40%. Data on the percentage of adults with income above this level comes from https://www.gov.uk/government/statistics/number-of-individual-income-taxpayers-by-marginal-rate-gender-and-age.

2.3 Framework

In this paper I examine the decision about how much to save in a pension as a static decision about how much of earnings to contribute to pension saving and how much to take for current consumption. In other words, in keeping with much of the literature in this area, I abstract from the intertemporal dimension of how much to save for the future.[2]

As mentioned, income tax on pension contributions is deferred, rather than relieved outright. This paper relies on a large discontinuity in the upfront income tax relief on pension contributions for identification. This will also reflect a large discontinuity in the overall long-run price of pension saving as long as the discontinuity in the upfront tax relief on pension contributions does not correspond exactly to a discontinuity in the marginal tax rate paid on pension income in retirement. This will almost certainly be the case; in fact, given the current income tax system, the majority of individuals saving in a DC pension in the UK will be at most basic-rate taxpayers in retirement (with a 20% marginal tax rate) (Adam et al., 2023).

[2]Little of the existing literature attempts to model responses to the price of saving in a lifecycle framework. One exception is Engelhardt and Kumar (2007).

3 Data

The data comes from the Annual Survey of Hours and Earnings (ASHE) for the years 2005 to 2019 (Office for National Statistics, 2022). The ASHE is an annual survey, filled out by employers, that contains accurate information on employees' individual and job characteristics, including earnings and workplace pension saving information. The sample frame for the ASHE is always the same 1% random sample of employees in the UK, meaning that it has a large sample size of around 180,000 individuals per year, and can be used for longitudinal analysis.[3]

The outcome of interest is whether an employee is enrolled in a workplace pension and, if they are, the monetary value of their pension contributions. In the ASHE, employers report the monetary value of employer and employee contributions to the employees' workplace pension made during the pay period. The ASHE does not collect information on contributions to personal pensions that are made independently by individuals, but this is uncommon among employees.

I also require a measure of individual income, both for calculating the tax price of pension contributions, and because income itself affects pension saving decisions. The ASHE contains good information on individual total earnings (including basic earnings, overtime earnings and bonus earnings - from multiple employers where relevant), but no information on unearned income from other sources such as rental income from property, or income from saving and investments. As a result, I use annual earnings to proxy annual income throughout this paper. I analysed data from the Survey of Personal Incomes and found that around the higher rate threshold earnings makes up over 95% of income for over 80% of employees. This implies that this proxy will be accurate for the vast majority of employees.

Throughout, the main analysis sample consists of 22- to 59-year-old private-sector employees with annual gross earnings between £30,000 and £70,000 in real terms (throughout I express real terms in 2019 £s). I focus on private-sector employees because most public-sector employees in the UK save in defined benefit pensions, where the employee has little autonomy over how much to contribute to the pension each year. I restrict the earnings range of the sample to around £15,000 to £25,000 above and below the higher rate threshold (depending on the year) to control more accurately for the effect of earnings on pension savings; robustness tables in the appendix show this has limited effect on the empirical results.

Table 3.1 presents the summary statistics for the two samples. Most, but not all, employees in the sample who are members of a workplace pension have strictly positive employee contributions.[4] Conditional on making positive contributions, employees

[3]There is, however, attrition caused by employer non-response, with 20-30% of employees in one year of the data not in the following year.

[4]I classify employees with zero employer and employee contributions as not being a member of a

contribute on average just under 5% of gross pay to their pension. Average conditional employer contributions reduced from just under 11% of gross pay in the 2005-12 period to less than 8% of pay by the later period.

Table 3.1: Sample Summary Statistics

Characteristic	2005-12	2013-19
% member of workplace pension	60.2	80.2
% with employee contributions > 0	51.2	74.6
% with employer contributions > 0	58.2	78.8
% member of Occ. DB pension	27.6	20.4
% member of Occ. DC pension	12.1	21.1
% member of Other DC pension	19.1	37.3
Average employee contribution among members (% pay)	4.9	4.4
Average employer contribution among members (% pay)	10.7	7.8
% women	28.0	30.6
% aged 22–34	30.4	29.9
% aged 35–49	48.1	45.0
% aged 50–59	21.4	25.1
Observations	230,585	208,538
People-jobs	89,651	85,021
People	72,386	70,810
People with >1 job	14,444	12,220

Notes: The samples contain 22-59 year-old private-sector employees with annual gross earnings between £30,000 and £70,000 in real (2019) terms. Average (conditional) employee and employer contributions means average contributions conditional on strictly positive contributions. Data come from the Annual Survey of Hours and Earnings.

3.1 Calculating the tax price of pension saving

In the empirical analysis the main independent variable of interest will be the upfront tax price of pension saving. This is how much contemporaneous disposable income the employee forgoes by contributing one more pound to their pension plan. Since employee contributions to pension plans are exempt from income taxes, one pound contributed to a pension will save someone £τ_I of income tax, where τ_I is the marginal income tax rate.

I calculate employee i's upfront tax price of pension saving in year t using the standard method of the literature (Kleven and Schultz, 2014; Almunia et al., 2020): I add a fixed amount Δs to their employee pension contributions, and compare their resulting tax liability with their originally calculated tax liability. Throughout, I choose $\Delta s =$ £10. More specifically, denoting by $T(y)$ the total income tax paid by an individual with annual taxable earnings y, I calculate the upfront tax price of pension saving as:

$$p_{it} = 1 - \frac{T(z_{it} - s_{it}) - T(z_{it} - s_{it} - \Delta s))}{\Delta s} \tag{3.1}$$

workplace pension.

In the UK, the tax price of pension saving can also be affected by how employee contributions are made, that is, whether they are made through 'salary sacrifice' or not.[5] The ASHE asks whether employee contributions were made through a salary sacrifice arrangement from 2013 onwards, so for these years I adjust the tax price measure accordingly.[6] For earlier years, I assume for simplicity that no one has a salary sacrifice arrangement; however, I calculate how sensitive the conclusions are to this assumption.

Figure 3.1 plots the average of the calculated tax price of pension saving by bins of real 'taxable earnings' (gross earnings after subtracting employee pension contributions). Panel (a) shows the price for the period 2005-12, while panel (b) shows the price over the period 2013-2019, separately for employees with and without salary sacrifice agreements. The estimated prices for those without a salary sacrifice arrangement are consistent with the tax rate schedule outlined in Section 2.2 and clearly show the discontinuity in the tax price around the higher rate threshold.[7] For those with a salary sacrifice arrangement, their tax price is affected by the employee NICs rate of 12% between around £10,000 and £50,000, which then falls to 2% for those with higher incomes. Clearly, there is still a large discontinuity in the tax price for those with a salary sacrifice arrangement at the higher rate threshold.

[5]Under a salary sacrifice arrangement, employees agree with their employer to reduce their earnings by an amount equal to their desired employee pension contributions, and for these contributions to be made as employer contributions instead. This is advantageous as, unlike employee pension contributions, employer contributions are not subject to another payroll tax (National Insurance contributions (NICs)). Therefore, for employees with salary sacrifice arrangements, one pound contributed to their pension saves them not only $£\tau_I$ of income tax, but also $£\tau_{NI}$ of National Insurance, where τ_{NI} is their marginal employee National Insurance rate.

[6]For the 2013-19 period, I can calculate the tax price accounting for salary sacrifice using Equation 3.1, where $T(y)$ denotes the total income tax and national insurance paid by someone with earnings y.

[7]Given that the higher rate threshold moves slightly between years, there are multiple bins in panel (a) with a tax price of pension saving between 0.8 or 0.78 and 0.6. Within year, the tax price changes discontinuously at the higher rate threshold. The same is true for panel (b).

Figure 3.1: Upfront income tax price of pension saving

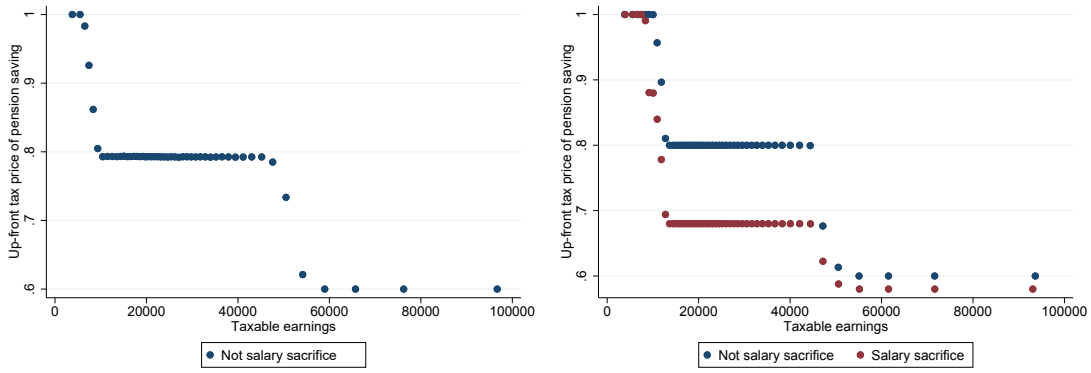

(a) 2005-12 (b) 2012-19

Notes: Shows the average calculated tax price of pension saving for 50 bins of taxable earnings (= gross earnings - employee contributions). Note that the tax thresholds change over time in real terms, and some bins will contain employees on either side of a tax kink, which is why the average for some bins will not equal one of the possible tax prices for any single individual: 1, 0.8, 0.78, 0.68, 0.6, and 0.58.

4 Graphical evidence

I start by describing graphically how pension participation and contributions vary around the higher rate tax threshold. I analyse the periods 2005 to 2012 and 2012 to 2019 separately to explore whether the introduction of automatic enrolment (rolled out from late 2012 onwards) had any discernable effect on pension saving behaviour.

Figure 4.1 plots the proportion of employees with strictly positive employee pension contributions by bins of gross annual earnings around the higher rate threshold, separately for the two time periods.[8] The incentive to contribute one pound to a pension increases discontinuously above the higher rate threshold: for example, it costs the employee £0.78 or £0.80 of disposable income if their gross income is less than the higher rate threshold, but only £0.60 if their gross income is above the threshold, assuming they do not have a salary sacrifice arrangement. Despite this discontinuity in incentives, Figure 4.1 shows no evidence of an increase in the proportion of employees making a positive contribution to their pension above the higher rate threshold in either period. This suggests a low extensive-margin responsiveness to changes in the tax price of pension saving.[9]

[8] I use quantile-spaced bins selected to minimise integrated mean-squared error, as outlined in Cattaneo et al. (2020).

[9] See Crawford et al. (2012) for similar results using survey data showing that the pension participation rate does not increase above the higher rate threshold in the UK.

Figure 4.1: Pension participation around the tax threshold, 2005-12 vs. 2013-19

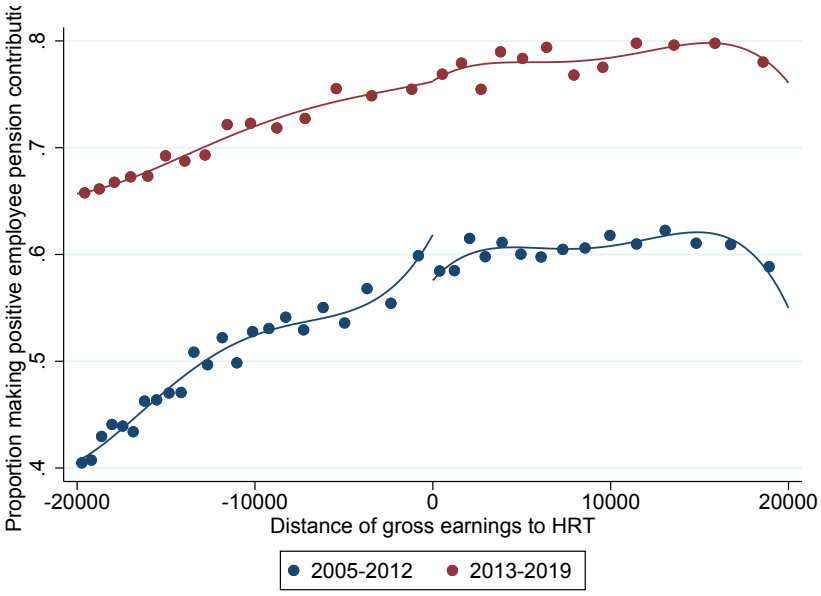

Notes: Shows the proportion of employees that are members of a workplace pension plan, by bins of gross earnings (normalised relative to the HRT). There is also a fourth-order polynomial fit to the data separately either side of the HRT. Distance of gross earnings to HRT is in real (2019) terms. Data come from the Annual Survey of Hours and Earnings.

To analyse whether pension contributions change at the higher rate threshold, I examine the degree of bunching of taxable earnings at this point. To see why, consider an employee with gross earnings above the higher rate threshold, who doesn't save via salary sacrifice. To start with, each pound contributed to their pension costs them £0.60 of contemporaneous disposable income. This is true until the point where their contributions are high enough that their taxable income equals the higher rate threshold: from this point on, each pound contributed costs them £0.78 or £0.80 of contemporaneous disposable income. Therefore, there is a convex kink in their budget set at this point, and we would expect bunching in response to this. Saez (2010) demonstrates that, with well-behaved preferences, the degree of bunching is positively related to the responsiveness of taxpayers to the tax incentive.

In Figure 4.2, I therefore investigate the degree to which employees are making pension contributions in such a way as to bunch their taxable income around the kink created by the higher rate threshold, starting with the 2005 to 2012 period. Specifically, I plot the distribution of gross earnings minus employee pension contributions around the higher rate threshold. Of course, there could also be bunching in this variable at the threshold for reasons unrelated to pension saving; for example, employees could adjust their hours to bunch their earnings at the kink. To account for this, I also plot the distribution of gross earnings around the tax threshold. If employees' pension saving in particular were responding to the change in the tax price at the higher rate threshold, we would

expect a larger degree of bunching in gross earnings minus employee contributions than in gross earnings at the higher rate threshold. However, Figure 4.2 shows no evidence of bunching in either variable, suggesting a low intensive-margin responsiveness of pension saving (and gross earnings) to the tax price. The equivalent figure for the 2013-19 period, Figure 4.3, also shows no evidence of bunching, again indicating a low intensive-margin elasticity. This is consistent with the evidence in Adam et al. (2021), who find little evidence of bunching in taxable income by employees in the UK around the higher rate threshold.

Figure 4.2: Bunching of income around the tax threshold, 2005-12

Notes: Shows the frequency density of gross income and taxable earnings around the higher rate threshold. Data come from the Annual Survey of Hours and Earnings.

Figure 4.3: Bunching of income around the tax threshold, 2013-19

Notes: Shows the frequency density of gross income and taxable earnings around the higher rate threshold. Data come from the Annual Survey of Hours and Earnings.

5 Empirical results

5.1 Empirical specification

The empirical methodology is similar to that of Almunia et al. (2020), who estimate the effect of tax incentives on charitable donations in the UK. I assume the pension saving of individual i working for employer j in year t depends both on the upfront tax price of pension saving, p_{it}, and post-tax earnings, y_{it}.

To estimate the intensive-margin responsiveness of employee pension contributions to the (upfront) tax price of pension saving, I estimate the following equation when contributions are strictly positive:

$$\ln s_{ijt} = \varepsilon_{INT} \ln p_{it} + \eta_{INT} \ln y_{it} + \delta X_{it} + \alpha_{ij} + \alpha_t + u_{ijt} \tag{5.1}$$

where s_{ijt} is the employee pension contribution of i working for employer j in year t, α_{ij} and α_t are employee-employer and year fixed effects, respectively, and u_{ijt} is an idiosyncratic error term. I control for the square of age in X_{it}. Then, ε_{INT} is the intensive-margin elasticity of pension saving to the upfront tax price, and η_{INT} is the intensive-margin income elasticity of pension saving. Note that y_{it} is post-tax earnings for employee i if they made zero pension contributions.

To estimate extensive-margin elasticities, I estimate a similar regression including all employees in the sample, where now the outcome variable is an indicator for whether

their employee contribution is strictly positive or not ($D_{ijt} \coloneqq \mathbb{1}(s_{ijt} > 0)$):

$$D_{ijt} = \beta \ln p_{it} + \gamma \ln y_{it} + \delta X_{it} + \alpha_{ij} + \alpha_t + u_{ijt}. \tag{5.2}$$

Then, to calculate the extensive-margin price and income elasticities, I divide the estimates of β and γ by the proportion of employees in the sample whose employee contributions are strictly positive.

The results in this section focus on how employee pension contributions respond to the change in the tax price at the higher rate threshold. Appendix Section B repeats this analysis for the responsiveness of employer pension contributions.

5.2 Identification

Estimating equations 5.1 and 5.2 by OLS is likely to yield upwardly-biased estimates of the relevant elasticities. This is because $\ln p_{it}$ is endogenous: an increase in pension contributions s_{it} reduces the employee's taxable income, and may therefore increase their tax price p_{it}. This problem has been widely discussed, and the standard solution is to instrument the "last-pound" price of pension saving, p_{it} with the "first-pound" price of pension-saving, p_{it}^f (Feldstein and Taylor, 1976). In this setting, this means the instrument is the tax price of pension saving that the employee would have faced had he or she made no employee pension contribution:

$$p_{it}^f = 1 - \frac{T(y_{it}) - T(y_{it} - \Delta s)}{\Delta s} \tag{5.3}$$

This instrument is highly correlated with the "last-pound" price of pension saving, since most employees do not contribute so much to their pension as to change their marginal tax band. Furthermore, Equation 5.3 shows that the instrument is not mechanically affected by s_{it} in the same way as p_{it} in Equation 3.1.

The identification of the relevant elasticities further relies on the assumption that changes in income are exogenous to changes in an employee's desire to save in a pension (note that including employee-employer fixed effects controls for time-invariant individual characteristics that may affect saving desires, such as their discount factor). That is, I rule out by assumption employees suddenly being motivated to earn more because they decide they want to increase the amount they save into their pension. In addition, I include employee-employer fixed effects in my preferred specification. I therefore exclude variation from employees moving to higher paying employers, who might offer a higher share of remuneration in employer pension contributions, thereby affecting employee pension contributions. In other words, I leverage only within-job variation in the tax price, and assume that employees receive an income y_{it} and then decide how much of this to save into their workplace pension, given their disposable income and the relative tax price

of pension saving, controlling for fixed individual and employer characteristics.

5.3 Responses to tax incentives before Automatic Enrolment (2005-12)

In this subsection, I estimate formally the elasticity of employee pension contributions with respect to the upfront tax price of pension saving during the 2005-12 period. I first estimate Equation 5.1 (the intensive margin equation) on all employees in the sample with strictly positive employee pension contributions for the years 2005 to 2012, either by OLS or IV, and including either employee or employee-employer fixed effects. I also control for year fixed effects and the square of age. The estimates are shown in Table 5.1, with the coefficient on the log price of pension saving being the estimate of the intensive-margin price elasticity, the main outcome of interest. Column (1) estimates Equation 5.1 using OLS with employee fixed effects, and I find a positive estimated elasticity of around 0.3. As explained in Section 5.2, we would expect the estimated elasticity to be upward biased when using OLS because of the reverse causation between pension contributions and the price of pension saving.

In column (2) I instrument the log of the "last-pound" price of pension saving using the log of the "first-pound" price of pension saving, as defined in Equation 5.3. The estimated elasticity becomes negative at around -0.11. Including employee-employer fixed effects, rather than employee fixed effects, as in column (3), reduces the magnitude of the estimate slightly to around -0.10. In column (4), I interact the log price of pension saving with an indicator variable for the type of pension the employee has—either an occupational defined benefit (DB) plan, an occupational defined contribution (DC) plan, or another type of workplace DC plan. We can see that the estimated elasticity is not significantly different from zero for those with DB plans, while I estimate an elasticity of slightly under -0.2 for those in DC plans. This is consistent with employees having a greater degree of control over their contributions in DC plans than in DB plans.

Table 5.1: The intensive-margin responsiveness of pension saving to tax incentives, 2005-12

	(1) OLS	(2) IV	(3) IV	(4) IV
Log price of pension saving	0.312***	-0.114***	-0.101***	
	(0.020)	(0.033)	(0.033)	
Log post-tax earnings	0.669***	0.511***	0.471***	0.471***
	(0.015)	(0.017)	(0.017)	(0.017)
Log price of pension saving * Occ DB				-0.000
				(0.036)
Log price of pension saving * Occ DC				-0.207***
				(0.050)
Log price of pension saving * Oth DC				-0.199***
				(0.045)
Observations	107,507	107,507	102,872	102,872
R^2	0.821	0.820	0.846	0.847
Employee FE	Yes	Yes	No	No
Employee-employer FE	No	No	Yes	Yes

Notes: All columns contain year fixed effects and control for the square of age. Columns (1) and (2) include employee fixed effects, while columns (3) and (4) include employee-employer fixed effects. Column (1) is estimated using OLS, while columns (2) to (4) instrument $\log(p_{it}^f)$ with $\log(p_{it})$ (or the interaction). Robust standard errors clustered at the employee level. Occ DB, Occ DC and Oth DC are dummy variables indicating wheether the employee is paricipating in an occupational DB, an occupational DC, or other DC pension plan. Data come from the Annual Survey of Hours and Earnings. *, **, and *** denote significance at 0.10, 0.05 and 0.01 levels, respectively.

Table 5.2 reports estimates of Equation 5.2 to evaluate the extensive-margin elasticity. The reported coefficients on the log price of pension saving and the log of disposable income correspond to the estimates of β and η in Equation 5.2, and I also report the implied price and income elasticities obtained by dividing these estimates by the share of individuals with strictly positive contributions.

The estimate obtained by OLS estimation in column (1) is again positive; however, columns (2) and (3) demonstrate that using the "first-pound" instrument gives us a negative elasticity. In the preferred specification, which includes employee-employer fixed effects, the estimated coefficient is only -0.024 and statistically insignificant at the 10% level. This implies an extensive-margin price elasticity of only -0.045.

Columns (4) to (6) show how the extensive-margin responsiveness differs by the type of pension. These results are obtained by changing the dependent variable in Equation 5.2 to be an indicator variable for positive employee contributions in the given type of pension plan. All the estimated coefficients on the log pension saving price are small and not statistically significant at the 10% level.

Table 5.2: The extensive-margin responsiveness of pension saving to tax incentives, 2005-12

	(1) OLS	(2) IV	(3) IV	(4) IV: DB	(5) IV: Occ DC	(6) IV: Oth DC
Log price of pension saving	0.269***	-0.025	-0.017	0.010	-0.018	-0.010
	(0.013)	(0.017)	(0.016)	(0.015)	(0.014)	(0.014)
Log post-tax earnings	0.204***	0.094***	0.052***	0.030***	-0.006	0.025***
	(0.009)	(0.010)	(0.009)	(0.008)	(0.007)	(0.008)
Observations	208,694	208,694	192,541	192,541	192,541	192,541
R^2	0.781	0.780	0.836	0.828	0.701	0.758
Employee FE	Yes	Yes	No	No	No	No
Employee-employer FE	No	No	Yes	Yes	Yes	Yes
Price elasticity	0.526	-0.048	-0.032	0.042	-0.185	-0.063
Income elasticity	0.399	0.183	0.102	0.121	-0.065	0.164

Notes: All columns contain year fixed effects and control for the square of age. Columns (1) and (2) include employee fixed effects, while columns (3) to (6) include employee-employer fixed effects. Column (1) is estimated using OLS, while columns (2) to (6) instrument $\log(p_{it}^f)$ with $\log(p_{it})$. Robust standard errors clustered at the employee level. In columns (1) to (3), the outcome variable is a dummy for whether the employee is participating in a workplace pension. In columns (4), (5) and (6), the outcome is a dummy variable for whether the employee is participating in an occupational DB, an occupational DC or other DC workplace pension, respectively. Data come from the Annual Survey of Hours and Earnings. *, **, and *** denote significance at 0.10, 0.05 and 0.01 levels, respectively.

To summarise, in the preferred specifications (column (3) in both Tables 5.1 and 5.2), I have estimated an intensive-margin elasticity of pension contributions to the upfront tax price of pension saving of approximately -0.10, and a corresponding extensive-margin elasticity of -0.05 for the time period 2005 to 2012. Aggregating, I estimate a total elasticity of around -0.15. These elasticities are small. At the higher rate threshold, where the tax price of pension saving falls by 25% after 2008 (from 0.8 to 0.6), these estimated elasticities imply an increase in pension participation by 1% and an increase in pension contributions by 2.5%. Taking an employee earning £60,000 a year and contributing £3,000 into their pension (i.e. assuming a 5% contribution rate, the average rate for employees at this point of the income distribution), this implies they contribute only about £75 more into their pension each year due to the change in the tax price at the higher rate threshold.

Robustness

As described in Section 3, I cannot observe whether employees were saving for their pension using a salary sacrifice arrangement for data years 2005 to 2012. The estimates in section 5.3 were obtained assuming that no employee had a salary sacrifice agreement. In 2013, approximately 50% of private-sector employees were saving in a pension via

salary sacrifice. These employees would have approximately faced only a 15% drop in the tax price of pension saving at the higher rate threshold, as opposed to the 25% I had assumed.[10] If I assume 50% of employees used salary sacrifice in the 2005-12 period, this would therefore increase the estimated average elasticity by one third, implying a total elasticity of around -0.2.

In Tables C.1 and C.2, I test how sensitive the results are to the restriction that employees in the sample have real annual gross earnings between £30,000 and £70,000. Specifically, I rerun the specifications from columns (3) of Tables 5.1 and 5.2 for different samples of earnings. Restricting the sample to employees with annual earnings between £35,000 and £65,000 has little effect on either the intensive or extensive margin price elasticity. Constricting the range to employees with earnings between £40,000 and £60,000 does lead to a slightly higher intensive margin price elasticity of -0.19, but the extensive margin price elasticity is little changed. Finally, including all employees with gross annual earnings between £20,000 and £90,000 reduces the magnitude of both elasticities substantially and also makes them insignificantly different from zero.

5.4 Responses to tax incentives under Automatic Enrolment (2013-19)

I now move on to estimating the responsiveness of employees' pension saving to tax incentives between 2013 and 2019. This coincides with the roll-out of automatic enrolment into workplace pensions between October 2012 and February 2018. This caused a large increase in pension participation (Cribb and Emmerson, 2020) among private-sector employees, with many contributing the default rates under automatic enrolment. In this subsection, I explore whether these changes, in particular the large increase in the number of passive savers enrolled in workplace pensions, have an effect on the responsiveness to tax incentives.

Table 5.3 presents the results from estimating Equation 5.1 for the 2013-19 sample period, with the same table structure as Table 5.1. The sample again includes private-sector employees with real annual earnings between £30,000 and £70,000 and strictly positive employee pension contributions. The preferred estimate of the intensive-margin price elasticity for 2013-19 is only -0.045, less than half the estimated elasticity for 2005-12, and it is also not significantly different from zero. Indeed, if I compare the 2013-19 estimate with the 2005-12 estimate that is corrected for salary sacrifice then the difference becomes even starker. Column (4) shows that the point estimates of the price elasticity are slightly larger for DC plans than for DB plans, but they are also insignificant and

[10] From 2008/9 on, employees with a salary sacrifice agreement have a tax price of 0.68 before they reach the higher rate tax threshold, and a tax price of 0.58 after the upper earnings limit, similar to panel (b) in Figure 3.1 These two thresholds are similar in most years, meaning that the pension price drops by $100 * \frac{0.1}{0.68} \approx 15\%$ at this point.

much smaller than in Table 5.1.

Table 5.3: The intensive-margin responsiveness of pension saving to tax incentives, 2013-19

	(1) OLS	(2) IV	(3) IV	(4) IV
Log price of pension saving	0.052**	0.030	0.004	
	(0.024)	(0.054)	(0.053)	
Log post-tax earnings	0.783***	0.775***	0.672***	0.662***
	(0.019)	(0.026)	(0.026)	(0.025)
Log price of pension saving * Occ DB				0.023
				(0.069)
Log price of pension saving * Occ DC				-0.036
				(0.064)
Log price of pension saving * Oth DB				-0.030
				(0.055)
Observations	135,333	135,333	126,626	126,626
R^2	0.846	0.846	0.876	0.876
Employee FE	Yes	Yes	No	No
Employee-employer FE	No	No	Yes	Yes

Notes: All columns contain year fixed effects and control for the square of age. Columns (1) and (2) include employee fixed effects, while columns (3) and (4) include employee-employer fixed effects. Column (1) is estimated using OLS, while columns (2) to (4) instrument $\log(p_{it}^f)$ with $\log(p_{it})$ (or the interaction). Robust standard errors clustered at the employee level. Occ DB, Occ DC and Oth DC are dummy variables indicating wheether the employee is paricipating in an occupational DB, an occupational DC, or other DC pension plan. Data come from the Annual Survey of Hours and Earnings. *, **, and *** denote significance at 0.10, 0.05 and 0.01 levels, respectively.

Table 5.4 presents the estimate of Equation 5.2 for the 2013-19 sample, where I include all private-sector employees within the £30,000-£70,000 annual earnings range. Column (3), the preferred specification, shows an estimated price elasticity that is approximately zero, even smaller than in the period 2005 to 2012. Columns (4) to (6) show that I again do not estimate a significant negative elasticity for any pension type.

Table 5.4: The extensive-margin responsiveness of pension saving to tax incentives, 2013-19

	(1) OLS	(2) IV	(3) IV	(4) IV: DB	(5) IV: Occ DC	(6) IV: Oth DC
Log price of pension saving	-0.626***	0.000	-0.001	0.043**	-0.035	-0.010
	(0.014)	(0.025)	(0.026)	(0.020)	(0.024)	(0.028)
Log post-tax earnings	-0.188***	0.066***	0.046***	0.044***	-0.016	0.018
	(0.011)	(0.014)	(0.015)	(0.010)	(0.013)	(0.015)
Observations	185,623	185,623	171,637	171,637	171,637	171,637
Employee FE	Yes	Yes	No	No	No	No
Employee-employer FE	No	No	Yes	Yes	Yes	Yes
R^2	0.643	0.636	0.689	0.823	0.716	0.717
Price elasticity	-0.840	0.000	-0.001	0.223	-0.179	-0.028
Income elasticity	-0.253	0.089	0.062	0.229	-0.081	0.053

Notes: All columns contain year fixed effects and control for the square of age. Columns (1) and (2) include employee fixed effects, while columns (3) to (6) include employee-employer fixed effects. Column (1) is estimated using OLS, while columns (2) to (6) instrument $\log(p_{it}^f)$ with $\log(p_{it})$. Robust standard errors clustered at the employee level. In columns (1) to (3), the outcome variable is a dummy for whether the employee is participating in a workplace pension. In columns (4), (5) and (6), the outcome is a dummy variable for whether the employee is participating in an occupational DB, an occupational DC or other DC workplace pension, respectively. Data come from the Annual Survey of Hours and Earnings. *, **, and *** denote significance at 0.10, 0.05 and 0.01 levels, respectively.

Overall, the results in this section point to an even smaller elasticity of pension contributions to the tax price of pension saving in the 2013-19 time period than in the 2005-12 time period, despite correcting for the attenuation in the earlier period caused by mismeasurement of the tax price for those saving using a salary sacrifice agreement. This suggests that the existence of nudges towards default levels of pension saving (through automatic enrolment) can affect the response to tax incentives. Indeed, previous research has demonstrated that the difference in pension participation between workers in the UK now depends much less on earnings than was previously the case (Cribb and Emmerson, 2020). Further, many of the people brought into saving in a pension by automatic enrolment are presumably more passive savers who are less likely to respond on the intensive margin to crossing the higher rate tax threshold.

6 Conclusion

Many governments provide tax incentives for pension saving to encourage saving for retirement. Existing evidence suggests that these incentives do often lead to an increase in saving in tax-advantaged accounts; however, in many cases this extra saving is partially or fully offset by a reduction in other forms of saving (Chetty et al., 2014; Goodman, 2020).

In contrast, this paper provides evidence of a setting where an increase in the tax incentive to save in a pension leads to only a very small increase in pension saving itself. In particular, I find only a very small responsiveness of UK employees' pension saving to crossing the higher rate threshold, which is a kink in the income tax schedule where there is a large, discontinuous fall in the marginal price of pension saving. I precisely estimate an intensive margin tax price elasticity of -0.1 and an extensive margin elasticity of -0.05 for the period 2005 to 2012. To put these elasticities into context, if employees earning £60,000 per year received 20% upfront income tax relief on their pension saving (rather than 40%), then pension participation would only fall by 1 percentage point, from 60% to 59%, and contributions (conditional on participation) would only fall by £75, from £3,000 to £2,925.

Repeating this analysis for the period 2013 to 2019, when automatic enrolment into workplace pensions was being rolled out across the UK, leads both estimated elasticities to fall to essentially 0. This paper is therefore the first to provide evidence on how these two important retirement savings policies interact, which is important for determining the relative merits of each. Studying the mechanisms behind this interaction is an interesting avenue for future research. On the one hand, it is almost certainly the case that automatic enrolment brought more passive savers into pension saving. But, in addition, it is possible that automatic enrolment led active savers to become passive savers, potentially because they view the default contribution rate as a "recommended" saving rate, or because the default contribution rate is close enough to their "optimal" saving rate that it is not worth the cost to change it (whether that be a cognitive or a time cost).

References

Adam, S., J. Browne, D. Phillips, and B. Roantree (2021). Frictions and taxpayer responses: evidence from bunching at personal tax thresholds. *International Tax and Public Finance 28*(3), 612–653.

Adam, S., I. Delestre, C. Emmerson, and D. Sturrock (2023). A blueprint for a better tax treatment of pensions. Institute for Fiscal Studies (IFS) Report 240.

Almunia, M., I. Guceri, B. Lockwood, and K. Scharf (2020). More giving or more givers? the effects of tax incentives on charitable donations in the uk. *Journal of Public Economics 183*, 104114.

Andersen, H. Y. (2018). Do tax incentives for saving in pension accounts cause debt accumulation? evidence from danish register data. *European Economic Review 106*, 35–53.

Bernheim, B. D. (2002). Chapter 18 - taxation and saving. Volume 3 of *Handbook of Public Economics*, pp. 1173–1249. Elsevier.

Cattaneo, M. D., N. Idrobo, and R. Titiunik (2020). *A Practical Introduction to Regression Discontinuity Designs: Foundations*. Elements in Quantitative and Computational Methods for the Social Sciences. Cambridge University Press.

Chan, M. K., T. Morris, C. Polidano, and H. Vu (2022). Income and saving responses to tax incentives for private retirement savings. *Journal of Public Economics 206*, 104598.

Chetty, R., J. N. Friedman, S. Leth-Petersen, T. H. Nielsen, and T. Olsen (2014). Active vs. Passive Decisions and Crowd-Out in Retirement Savings Accounts: Evidence from Denmark. *The Quarterly Journal of Economics 129*(3), 1141–1219.

Chetty, R., J. N. Friedman, and E. Saez (2013). Using differences in knowledge across neighborhoods to uncover the impacts of the eitc on earnings. *The American Economic Review 103*(7), 2683–2721.

Choi, J. J. (2015). Contributions to defined contribution pension plans. *Annual Review of Financial Economics 7*(1), 161–178.

Christensen, C. S. and B. E. Ellegaard (2022). Do tax subsidies for retirement saving impact total private saving? new evidence on middle-income workers. CEBI Working Paper 17/22.

Crawford, R., R. Disney, and C. Emmerson (2012). Do up-front tax incentives affect private pension saving in the united kingdom? IFS Working Papers W12/05.

Cribb, J. and C. Emmerson (2020). What happens to workplace pension saving when employers are obliged to enrol employees automatically? *International Tax and Public Finance 27*(3), 664–693.

Cribb, J. and C. Emmerson (2021). What can we learn about automatic enrollment into pensions from small employers? *National Tax Journal 74*(2), 377–404.

Engelhardt, G. V. and A. Kumar (2007). Employer matching and 401(k) saving: Evidence from the health and retirement study. *Journal of Public Economics 91*(10), 1920–1943.

Engen, E. M., W. G. Gale, and J. K. Scholz (1996). The illusory effects of saving incentives on saving. *The Journal of Economic Perspectives 10*(4), 113–138.

Feldstein, M. and A. Taylor (1976). The income tax and charitable contributions. *Econometrica 44*(6), 1201–1222.

Friedman, J. N. (2017). Tax Policy and Retirement Savings. In *The Economics of Tax Policy*. Oxford University Press.

Gelber, A. M. (2011). How do 401(k)s affect saving? evidence from changes in 401(k) eligibility. *American Economic Journal: Economic Policy 3*(4), 103–22.

Goodman, L. (2020). Catching up or crowding out? the crowd-out effects of catch-up retirement contributions on non-retirement saving. *Journal of Public Economics 188*, 104221.

Hubbard, R. G. and J. S. Skinner (1996). Assessing the effectiveness of saving incentives. *Journal of Economic Perspectives 10*(4), 73–90.

Kleven, H. J. and E. A. Schultz (2014). Estimating taxable income responses using danish tax reforms. *American Economic Journal: Economic Policy 6*(4), 271–301.

Lavecchia, A. (2019). 'back-loaded' tax subsidies for saving, asset location and crowd-out: Evidence from tax-free savings accounts. McMaster University Department of Economics Working Paper Series 2019-04.

Madrian, B. C. and D. F. Shea (2001). The power of suggestion: Inertia in 401(k) participation and savings behavior. *The Quarterly Journal of Economics 116*(4), 1149–1187.

Messacar, D. (2018). Crowd-Out, Education, and Employer Contributions to Workplace Pensions: Evidence from Canadian Tax Records. *The Review of Economics and Statistics 100*(4), 648–663.

Office for National Statistics (2022). Annual survey of hours and earnings, 1997-2020: Secure access [data collection]. 21st edition. UK Data Service. SN: 6689, DOI: 10.5255/UKDA-SN-6689-20.

Poterba, J. M., S. F. Venti, and D. A. Wise (1995). Do 401(k) contributions crowd out other personal saving? *Journal of Public Economics 58*(1), 1–32.

Saez, E. (2010). Do taxpayers bunch at kink points? *American Economic Journal: Economic Policy 2*(3), 180–212.

A Data Appendix

The data used for the empirical analysis in this paper comes from the Annual Survey of Hours and Earnings. I create annual-level panel datasets containing 22 to 59-year-old private-sector employees, one for the years 2005 to 2012, and another for the years 2013

to 2019. The datasets contain information employees' post-tax earnings, their employee and employer pension contributions, the type of pension they are saving into, the (first- and last-) pound price of pension saving, their employer and the square of the employee's age. The main sample is restricted to those employees with annual earnings between £30,000 and £70,000 in real (2019) terms. This data appendix contains some extra detail on the construction of this dataset.

A.1 Aggregating monthly variables to the annual level

The ASHE is filled out by employers with information relating to the employee's pay period that encompasses a particular reference date in April. As the UK tax year starts on 6th April and runs to 5th April the following year, the information in the ASHE relates to effectively the first month of the tax year. I aggregate pension contributions and earnings to the annual level, assuming they are unchanged throughout the year.

I test whether this is a reasonable assumption for earnings using data in the ASHE about the employee's annual gross pay in their current job for the tax year ending on 5 April for the survey year. This may be a better measure of annual earnings for employees whose earnings are volatile. On the other hand, it will underestimate annual earnings for employees who worked in their current job for less than a year. Despite this, the correlation between the two variables is over 0.9, suggesting that aggregated monthly earnings approximates annual earnings well. The reason why I do not use the annual earnings data is that there is no data on pension contributions for the same time frame.

A.2 Net pay vs. relief-at-source pension arrangements

In the UK there are two arrangements through which tax relief on pension contributions can be administered, which will have an effect on the measurement of pension contributions. In 'net pay' plans, pension contributions are deducted before tax is calculated on the employee's pay. Conversely, in 'relief at source' plans, the pension contribution is deducted after tax is calculated, and HMRC then sends an additional 25% contribution to the pension plan to make up for the tax paid - higher rate taxpayers can, however, solicit an extra refund to make up for the extra tax they paid. This extra refund for higher-rate taxpayers is paid into their bank accounts, and so does not affect their employee pension contributions. I cannot observe whether a plan is 'net pay' or 'relief at source'. For 'relief at source' plans, I observe pension contributions before the additional top-up from HMRC, and so underestimate the total amount entering the employee's pension plan by 25%. However, this underestimation will be captured in the employer-employee fixed effects I include in the main specification (since the intensive margin outcome is *log* employee pension contributions). I will also mismeasure the taxable income of relief-at-source employees, and therefore potentially their pension saving price, p_{it}. However,

this mismeasurement of p_{it} will be alleviated by the instrument introduced in Section 5.2, which is calculated based on the "first pound" price of pension saving i.e. before deducting any pension contributions.

A.3 The effect of the benefit system on the tax price

Parts of the benefit system can also affect the upfront price of pension saving. Eligibility for most means tested benefits is assessed against a measure of income that excludes pension contributions. This means that the full tax price of pension contributions is also affected by whether or not individuals would gain extra entitlement to benefits as a result of their contribution. Since I am exploiting on a non-linearity in the income tax schedule that occurs around the 90th percentile of the earnings distribution, most means-tested benefits will not be relevant for any of the individuals in the sample. The exception is child benefit which, since January 2013, is gradually withdrawn from those earning £50,000 or more, such that those earning £60,000 or more receive no benefit.[11] For someone with two children under 16, who would be entitled to £1,752 per year of child benefit in 2013, this amounts to an effective tax rate of 17.5% on earnings between £50,000 and £60,000. Since the measure of income used for this assessment excludes pension contributions, the tax price of pension saving is substantially reduced for those receiving child benefit whose income would otherwise be in this range. I cannot observe in the ASHE whether or not someone receives child benefit and so assume for simplicity that no one in the sample does. This would cause the magnitude of the (very small) estimated elasticities in the 2013-19 period to be an *overestimate*, since the fall in the tax price around £50,000 (near the HRT) is much greater for those receiving child benefit.

B The responsiveness of employer contributions

The analysis in the main part of this paper focuses solely on how employee pension contributions respond to the change in the tax price at the higher rate threshold. In this section, I show how employer pension contributions change in response to this tax price incentive.

There are two main reasons why employer contributions might respond to the change in the tax price. First, if many employees are in pension plans where the employer 'matches' the employee's contribution, then any increase in employee contributions in response to tax incentives might mechanically lead employers to raise their contribution too. However, given I find a low elasticity for employee pension contributions, this matching mechanism is unlikely to lead to a large elasticity for employer contributions. A

[11] In other words, there is an effective tax rate of 1% of child benefit entitlement on each £100 earned over £50,000.

second reason is that employers might change their compensation package in response to a change in the employee's marginal tax rate. Specifically, it is conceivable that employers agree to reduce the employee's pay and increase their employer pension contributions after they cross the higher rate threshold.[12] Note that since the preferred specification includes employee-employer fixed effects, throughout I rule out variation arising from employees who cross the higher rate threshold moving to employers whose compensation package offers higher employer pension contributions in return for lower pay.

Table B.1 shows the estimated elasticity of employer contributions to the upfront tax price of pension saving. Throughout, I instrument the log price of pension saving using the first-pound price, and include year and employee-employer fixed effects, as well as a control for the square of age. I estimate very small intensive- and extensive-margin elasticities that are not statistically significantly different from zero. This suggests that employers do not change the amount they are contributing to an employee's pension in response to the employee crossing the higher rate threshold. This is perhaps not surprising given the small elasticity I find for employees, who have a larger incentive to respond.

Table B.1: The responsiveness of employer contributions to tax incentives

	2005-2012		2013-2019	
	(1)	(2)	(3)	(4)
	Intensive	Extensive	Intensive	Extensive
Log price of pension saving	-0.046	-0.013	0.034	0.005
	(0.031)	(0.015)	(0.048)	(0.025)
Log post-tax earnings	0.553***	0.059***	0.675***	0.035**
	(0.018)	(0.008)	(0.023)	(0.014)
Observations	117,017	192,541	134,506	171,637
R^2	0.878	0.848	0.920	0.661
Employee-employer FE	Yes	Yes	Yes	Yes
Price elasticity	-0.046	-0.023	0.034	0.006
Income elasticity	0.553	0.102	0.675	0.044

Notes: All columns contain year fixed effects, employee-employer fixed effects, and control for the square of age. All columns instrument $\log(p^f_{it})$ with $\log(p_{it})$ (or the interaction). Robust standard errors clustered at the employee level. Data come from the Annual Survey of Hours and Earnings. *, **, and *** denote significance at 0.10, 0.05 and 0.01 levels, respectively.

[12]Note that here I am referring to employer contributions decided by the employer. In the analysis, employer contributions do *not* include employee contributions made through a salary sacrifice agreement (which are treated as employer contributions for tax purposes, but are recorded as employee contributions in the data.)

C Appendix tables

Table C.1: Testing the sensitivity of 2005-12 intensive margin estimates to the sample income range

	(1) 30-70K	(2) 35-65K	(3) 40-60K	(4) 20K-80K
Log price of pension saving	-0.101***	-0.120***	-0.178***	-0.038
	(0.033)	(0.042)	(0.068)	(0.028)
Log post-tax earnings	0.471***	0.451***	0.382***	0.547***
	(0.017)	(0.027)	(0.055)	(0.011)
Observations	102,872	73,706	45,919	158,272
R^2	0.846	0.833	0.833	0.876
Employee-employer FE	Yes	Yes	Yes	Yes

Notes: All columns contain year fixed effects, employee-employer fixed effects, and control for the square of age. All columns instrument $\log(p_{it}^f)$ with $\log(p_{it})$ (or the interaction). Robust standard errors clustered at the employee level. The sample is restricted to observations with real earnings in the range specified in the column title. Data come from the Annual Survey of Hours and Earnings. *, **, and *** denote significance at 0.10, 0.05 and 0.01 levels, respectively.

Table C.2: Testing the sensitivity of 2005-12 extensive margin estimates to the sample income range

	(1) 30-70K	(2) 35-65K	(3) 40-60K	(4) 20K-80K
Log price of pension saving	-0.017	-0.034*	-0.033	0.001
	(0.016)	(0.020)	(0.027)	(0.014)
Log post-tax earnings	0.052***	0.048***	0.070***	0.056***
	(0.009)	(0.014)	(0.027)	(0.005)
Observations	192,541	130,206	78,566	352,019
R^2	0.836	0.833	0.831	0.849
Employee-employer FE	Yes	Yes	Yes	Yes
Price elasticity	-0.032	-0.065	-0.064	0.002
Income elasticity	0.102	0.093	0.136	0.110

Notes: All columns contain year fixed effects, employee-employer fixed effects, and control for the square of age. All columns instrument $\log(p_{it}^f)$ with $\log(p_{it})$ (or the interaction). Robust standard errors clustered at the employee level. The sample is restricted to observations with real earnings in the range specified in the column title. Data come from the Annual Survey of Hours and Earnings. *, **, and *** denote significance at 0.10, 0.05 and 0.01 levels, respectively.