The Effect of Tax Incentives on Pension Saving

Laurence O'Brien *

February 15, 2022

Abstract

This paper estimates the responsiveness of private pension saving to tax incentives for UK employees, using employer-provided data on pension contributions between 2005 and 2019. Exploiting a kink in the income tax schedule and using a first-pound instrument for the upfront marginal income tax price of pension saving, we find an intensive-margin elasticity of around -0.1 and an extensive-margin elasticity of -0.05 for the earlier part of our sample period, 2005 to 2012. In 2013 to 2019, after the introduction of automatic enrolment into workplace pension plans, we find a lower average elasticity, consistent with those being brought into pension saving by this policy being passive savers. In general, employees do not respond strongly to this upfront tax incentive to save.

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

*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 pension saving would respond to changes in tax incentives.

In this paper, we estimate how responsive 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 has been affected by automatic enrolment into workplace pensions. Specifically, we test the extent to which employees in the UK change their pension contributions at a kink in the income tax schedule, where the upfront tax price of pension saving changes discontinuously. We provide graphical evidence that both the intensive- and extensive-margin elasticities are small, which is supported by results from panel regressions. Further, we show how the magnitude of the elasticity has become smaller still after the introduction of automatic enrolment, which drastically increased the number of passive savers contributing to workplace pension plans.

Our paper builds on a large literature that has studied the determinants of individual pension saving, surveyed recently by Choi (2015). A common finding is that a more generous employer match rate does not lead to a big increase in pension membership or contributions (Choi et al., 2002; Duflo et al., 2006), suggesting that pension saving is relatively insensitive to changes in the price of pension saving relative to take-home pay. However, it is not clear whether this result also extends to variation in the price of pension saving caused by the tax system. There is extensive evidence that most people's pension saving decisions do not match the predictions of frictionless optimising models (Choi et al., 2011; Card and Ransom, 2011), implying that price variation caused by little understood aspects of the tax treatment of pensions may result in smaller responses than very explicit and well-communicated match rates from employers.

The literature examining the effect of tax incentives on pension saving decisions is therefore more directly relevant to our paper. Early contributions to this literature found large elasticities with respect to the marginal tax rate (O'Neil and Thompson, 1987; Venti and Wise, 1988); however, Eaton (2002) highlighted that estimates from these papers could be significantly biased by failing to control for income effects. More recent studies that control for income effects typically find a lower degree of responsiveness. Feng (2014) and Eaton (2002) find only limited effects of tax incentives in Australia and the US, respectively; however, they do not analyse the intensive-margin response. Selin (2012) analyses the intensive margin in a closely related setting to ours, exploiting non-linearities in the Swedish tax system to examine the tax price elasticity of pension saving among the self-employed. He finds an elasticity of -0.51, which is significantly larger than our preferred estimates. Our results contribute to this literature by estimating both the intensive- and extensive-margin elasticity for employees, who make up a larger share of workers than the self-employed, and by showing that they respond much less to this tax incentive.

Another related paper is Chetty et al. (2014), who study how pension contributions respond

to a reduction in a subsidy for contributing to pension accounts in Denmark for individuals in the top income tax bracket. They find that pension contributions did fall fairly sharply in response to the reform, but that this was driven entirely by just 19% of contributors. They highlight that, when it comes to pension saving, we can think of there being active and passive savers, with the former likely to respond to the price of saving and the latter not. We build on this finding by analysing the responsiveness of pension saving before and after a large reform that drastically increased the number of passive savers, and show that the estimated elasticity is significantly lower post-reform.

To produce our estimates, we use panel data on employees from a survey completed by employers for the period 2005 to 2019. We exploit a large non-linearity in the income tax schedule that creates a sharp discontinuity in the upfront tax price of pension contributions to examine how pension saving responds to this price at both the intensive and extensive margin.

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. We address this by employing the standard first-pound price instrument (Feldstein and Taylor, 1976). The second challenge is the possible simultaneous choice of income and pension saving; we address this by using individual-employer fixed effects to restrict identification to changes in the tax price for individuals working for an unchanged employer.

We estimate the intensive- and extensive-margin price elasticities separately for the period before the introduction of automatic enrolment into workplace pensions (2005 to 2012) and for the period during which this policy was being rolled out or was in place (2013 to 2019). For the prior period our estimates of the intensive- and extensive-margin price elasticities are -0.1 and -0.05 respectively, giving a total elasticity of -0.15. For the latter period our estimates of the elasticities fall to essentially 0.

These results suggest that individuals around the income tax non-linearity that we examine – which is just above the 90th percentile of the income distribution – do not on the whole respond to this incentive to save. This is even less the case since the introduction of automatic enrolment, 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. The additional cost of upfront tax relief accruing to these savers therefore arises mainly from the mechanical effect of their higher tax rate and their propensity to save more because they earn more, rather than because of a behavioural response to the tax incentive. However, additional research is required to know whether this is the same higher up the earnings distribution, where individuals may be more financially astute, but where the greater upfront incentive to save in a pension is less likely to be a long-run incentive (as they are more likely to pay higher rates of income tax in retirement).

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

2 Institutional background

2.1 Private pension saving in the UK

The UK public pension system now only provides 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 taken out by individuals, the vast majority of pension saving for employees is facilitated by employers – either in pensions set up specifically for their employees, or by arranging access to a pension facilitated by a private provider (normally an insurance company). For public sector employees these pensions are 'defined benefit' (DB) in nature: the scheme rules specify the contributions that must be made and the benefits that will be paid in retirement, with those benefits normally being determined by years of service and a measure of salary. For private sector employees, workplace pensions are predominantly 'defined contribution' (DC) schemes: schemes in which contributions are paid into a fund, accrue an investment return, and can then be flexibly accessed in retirement. Employees in DC schemes can typically choose their level of contributions (which may be influenced by matching arrangements provided by some employers). Some DC pension schemes are run by the employer directly, while others are run by external pension providers, with the employer facilitating membership. We refer to these two types of pension as 'Occupational' and 'Other' DC pensions, respectively. In 2005, the start of the period we examine, 44% of private-sector employees were in a private pension: 24% in a DB scheme, 9% in an occupational DC scheme, and 12% in other DC schemes.

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 membership (Cribb and Emmerson, 2020): in 2019, the end of the period we examine, 79% of private-sector employees were in a private pension (12% in a DB scheme, 24% in an occupational DC scheme, and 44% in other DC schemes).

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 another 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.

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 for incomes up to £100,000 - the schedule for earlier years we analyse is very similar. Income tax is based on individual income. Individuals have a personal allowance, on which no income 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 our analysis we focus on individuals earning below this.

Figure 2.1: *UK income tax schedule 2019-20 up to £100,000*

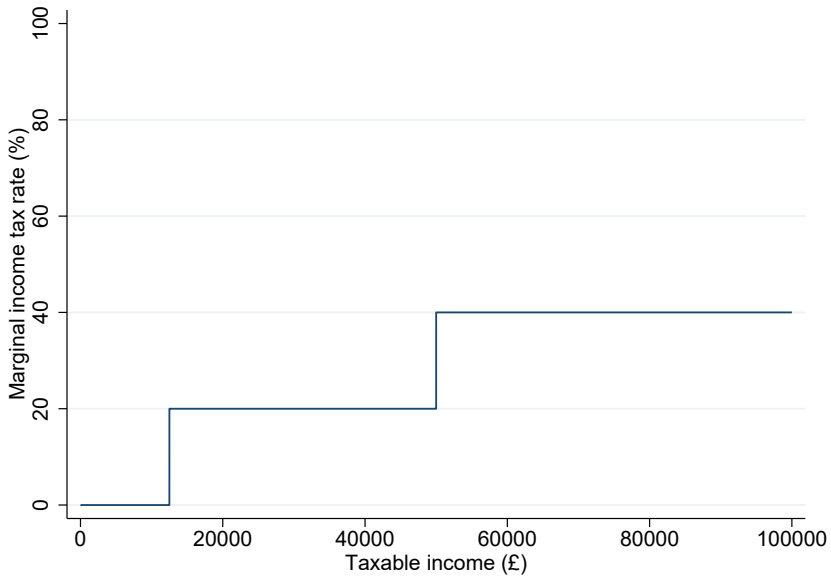

Source: https://www.gov.uk/income-tax-rates

Figure 2.2 shows that the real level of the higher rate threshold changed little over our analysis period, fluctuating between £45,000 and £55,000. In this figure, we also show that the higher rate threshold lies towards the top of the income distribution, with about 6 to 8% of adults earning more than the higher rate threshold in the UK, depending on the year. Given that we estimate a local treatment effect of the responsiveness of pension saving around the higher rate threshold, our estimate is relevant to well-off, but not super-rich, individuals.

Figure 2.2: *Higher rate threshold: its level and the proportion of adults with income above it*

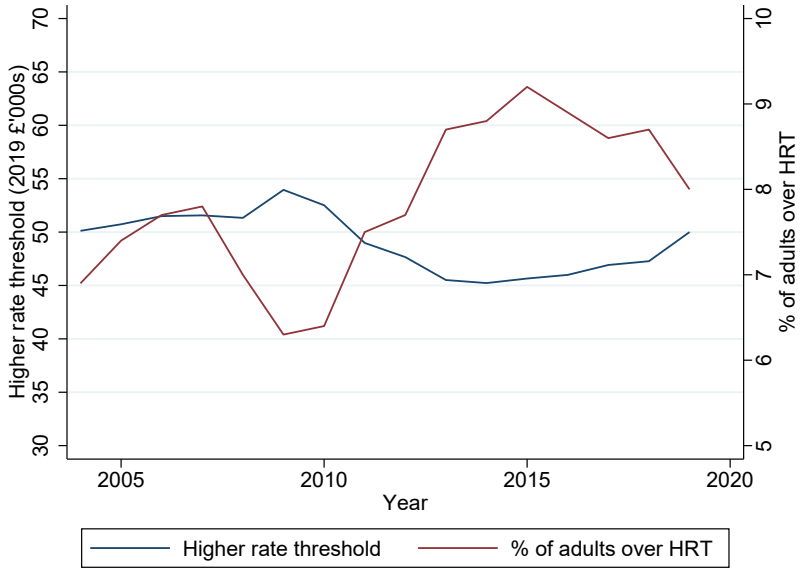

Notes: The higher rate threshold is the level of taxable income at which the UK marginal income tax rate increases to 40%. Data on the percentage of adults with income above this level comes from https://www.gov.uk/government/statistics/number-of-individual-income-taxpayers-by-marginal-rate-gender-and-age.

2.3 Framework

In this paper we 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, we abstract from the intertemporal dimension of how much to save for the future.[1]

Given that most income tax on pension contributions is deferred, rather than relieved outright, this means that our price elasticities should be interpreted as elasticities with respect to the upfront price of saving, identified using a discontinuous change in that upfront price. That is not the same as the long-run price, which will itself depend on the income tax rate paid by an individual in retirement. However, the majority of those saving in a DC pension would be expected (at least given the current income tax system) to only be basic-rate taxpayers in retirement, as there are limits on the amounts that can be saved in private pensions. This means that, while we cannot explicitly measure it, the discontinuity in the upfront price of pension contributions also likely reflects a discontinuity in the long-run price of that pension saving.

3 Data

Our data comes from the Annual Survey of Hours and Earnings (ASHE) for the years 2005 to 2019 (Office for National Statistics, 2021). The ASHE is a compulsory annual survey, filled out by employers, that contains accurate information on employees' individual and job charac-

[1]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).

teristics, 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 160,000 individuals per year, and can be used for longitudinal analysis.[2]

Our outcome of interest is the monetary value of pension contributions made by an employee. In ASHE, employers report the monetary value of employer and employee contributions to the employees' workplace pension made during the reference period.[3] In the UK there are two arrangements through which tax relief on pension contributions can be administered. In 'net pay' schemes, pension contributions are deducted before tax is calculated on the employee's pay. Conversely, in 'relief at source' schemes, the pension contribution is deducted after tax is calculated, and HMRC then sends an additional 25% contribution to the pension scheme to make up for the tax paid.[4] We cannot observe whether a scheme is 'net pay' or 'relief at source'. For 'relief at source schemes', we observe pension contributions before the additional top-up from HMRC, and so underestimate the total amount entering the employees' pension scheme. For these people, we will also mismeasure their taxable income, 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. Furthermore, by using only within-job variation in our preferred specification, together with log contributions as an outcome, the fact we underestimate pension contributions by 25% will only cause a problem for employees whose pension saving arrangement changes within job, which is rare.

We also require a measure of individual income, both for calculating the tax price of pension contributions, and because income itself affects pension saving decisions. 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, we use annual earnings to proxy annual income throughout this paper. Data from the Survey of Personal Incomes suggests that, around the higher rate threshold, earnings makes up over 95% of income for over 80% of employees, implying that our proxy will be accurate for the vast majority of employees.

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. We aggregate pension contributions and earnings to the annual level, assuming they are unchanged throughout the year.[5]

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

[3]ASHE does not collect information on contributions to personal pensions that are made independently by individuals, but this is uncommon among employees.

[4]For basic rate taxpayers, this top-up will exactly compensate for the income tax paid on pension contributions. Higher rate taxpayers can solicit an extra refund to make up for the extra tax they paid, but this is paid into their bank accounts, and so does not affect their employee pension contributions.

[5]We test whether this is a reasonable assumption 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.

Throughout, our 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.[6] We 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. We restrict the earnings range of our sample to around £15,000 to £25,000 above and below the higher rate threshold to allow us to control more accurately for the effect of earnings on pension savings.

Table 3.1 presents the summary statistics for our two samples. Most, but not all, employees in our sample who are members of a workplace pension have strictly positive employee contributions[7]. Conditional on making positive contributions, employees contribute on average just under 5% of gross pay to their pension. 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	61.01	80.61
% with employee contributions > 0	51.96	74.95
% with employer contributions > 0	58.96	79.22
% member of Occ. DB pension	28.16	20.78
% member of Occ. DC pension	12.15	21.23
% member of Other DC pension	19.25	37.25
Average (conditional) employee contribution	4.90	4.41
Average (conditional) employer contribution	10.75	7.89
% women	27.74	30.30
% aged 22–34	29.96	29.40
% aged 35–49	48.51	45.25
% aged 50–59	21.53	25.35
Total observations	245,929	222,475
People-jobs	95,306	90,164
People	77,046	75,131
People with >1 job	15,272	12,952

Notes: Our 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 our empirical analysis our 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 scheme. Since employee contributions to pension schemes are exempt from income taxes, one pound contributed to a pension will save someone $£\tau_I$ of income tax, where τ_I is the marginal income tax rate.

The reason why we do not use the annual earnings data is that we have no data on pension contributions for the same time frame.

[6] All real terms expressed in 2019 £s.

[7] We classify employees with zero employer and employee contributions as not being a member of a workplace pension.

We 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): we add a fixed amount Δs to their employee pension contributions, and compare their resulting tax liability with their originally calculated tax liability. Throughout, we choose $\Delta s = 10$. More specifically, denoting by $T(y)$ the total income tax paid by an individual with annual taxable earnings y, we calculate the upfront tax price of pension saving as:

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

In practice there are three complications with this calculation. The first is that the tax price of pension saving can also be affected by how employee contributions are made. In the UK, it is possible for employees to 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 called a 'salary sacrifice' arrangement. 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 £τ_I of income tax, but also £τ_{NI} of National Insurance, where τ_{NI} is their marginal employee National Insurance rate. Unfortunately, the ASHE only asks whether employee contributions were made through a salary sacrifice arrangement from 2013 onwards. For earlier years, we assume for simplicity that no one has a salary sacrifice arrangement; however, we do perform a back-of-the-envelope calculate to estimate how sensitive our conclusions are to this assumption.[8]

The second complication is that 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 we are 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 most of the individuals in our 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.[9] For someone with two children, 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. We cannot observe in ASHE whether or not someone receives child benefit, meaning that this will be another source of measurement error in our price variable from 2013 onwards. However data from the Family Resources Survey suggests that only a minority of taxpayers in our sample are eligible for child benefit. Furthermore, this mismeasurement would cause us to overestimate the magnitude of the elasticity, since the fall in the tax price around £50,000 (near

[8]For the 2013-19 period, we 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.

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

the HRT) is much greater for those receiving child benefit.

Finally, as mentioned earlier, ASHE provides no information on unearned income, meaning we approximate taxable income with gross earnings. This will be another source of measurement error; however, for the majority of earners, this approximation will be accurate and we will accurately measure their marginal tax rate.

Figure 3.1 plots the average of our calculated tax price of pension saving by bins of what we call 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.[10] 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.

Figure 3.1: *Upfront income tax price of pension saving*

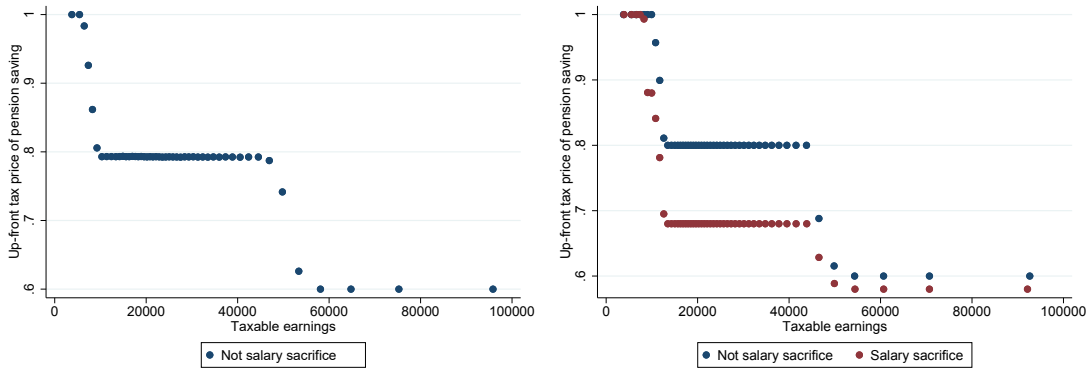

(a) *2005-12* (b) *2012-19*

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

We start by describing graphically how pension membership and contributions vary around the higher rate tax threshold. We analyse the periods 2005 to 2012 and 2012 to 2019 separately for two reasons. First, this is before automatic enrolment into workplace pensions started to be introduced, so the saving environment was more stable and those saving in a pension are more likely to have been 'active' savers. Second, while we cannot observe salary sacrifice arrangements prior to 2013, we do not have the potentially confounding influences of the means-testing of

[10] 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).

child benefit on the tax price of pension saving.

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. 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.

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

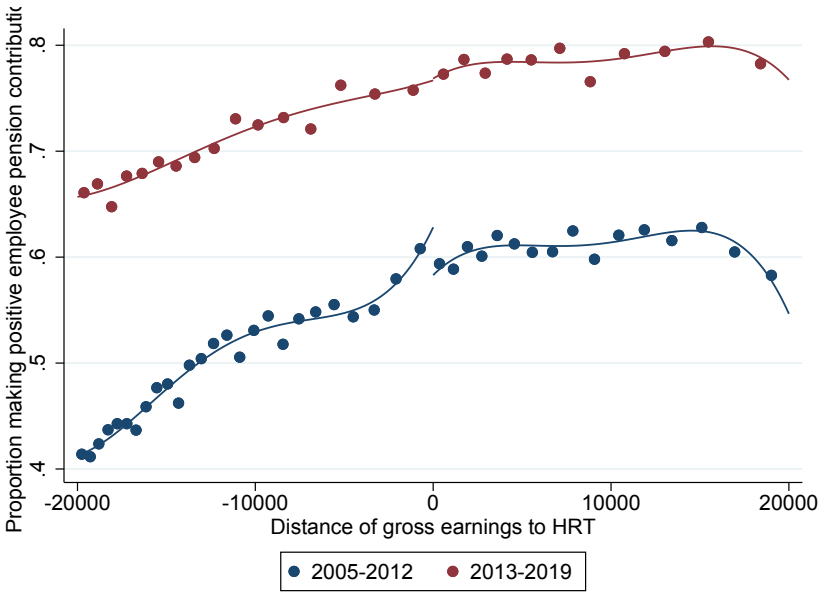

Notes: Shows the proportion of employees that are members of a workplace pension scheme, by bins of gross earnings (normalised relative to the HRT). There is also a fourth-order polynomial fit to the data separately either side of the HRT. Distance of gross earnings to HRT is in real (2019) terms. Data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

To analyse whether pension contributions change at the higher rate threshold, we examine the degree of bunching 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, under certain assumptions, the degree of bunching is positively related to the responsiveness of taxpayers to the tax incentive.

In Figure 4.2, we 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, we 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, we 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 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

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

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

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

where s_{it} is the employee pension contribution of i in year t, α_{ij} and α_t are employee-employer and year fixed effects, respectively, and u_{it} is an idiosyncratic error term. We 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, we 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 ($D_{it} \coloneqq \mathbb{1}(s_{it} > 0)$):

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

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

5.2 Identification challenges

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 in the literature, 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^f_{it}. In our 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:

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

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.

Our identification of the relevant elasticities further relies on the assumption that changes in income are exogenous to employees' desire to save in a pension. In other words, we rule out by assumption employees being motivated to earn more because they want to increase the amount they save into their pension. We also rule out, through the inclusion of employee-employer fixed effects, identification coming from employees who move job to an employment with a different balance between earnings and pension contributions in the compensation package. In other words, we leverage only within-job variation in the tax price, and assume that employees receive an income y_{it}, exogenous of their pension saving decision, 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.

5.3 Regression results, 2005-12

In this section, we estimate formally the elasticity of pension contributions with respect to the upfront tax price of pension saving. We first estimate Equation 5.1 on all employees in our 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[11]. The estimates are shown in Table 5.1, with the coefficient on the log price of pension saving being our estimate of the intensive-margin price elasticity, our main outcome of interest. Column (1) estimates Equation 5.1 using OLS with employee fixed effects, and we 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.

[11] We also control for year fixed effects and the square of age.

In column (2) we 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), we 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) scheme, an occupational defined contribution (DC) scheme, or another type of workplace DC scheme. We can see that the estimated elasticity is not significantly different from zero for those with DB schemes, while we estimate an elasticity of slightly under -0.2 for those in DC schemes. This is consistent with employees having a greater degree of control over their contributions in DC schemes than in DB schemes.

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	0.311***	-0.106***	-0.095***	
	(0.019)	(0.032)	(0.032)	
Log disposable income	0.666***	0.509***	0.468***	0.468***
	(0.014)	(0.017)	(0.017)	(0.016)
Pension price * Occ DB				-0.010
				(0.034)
Pension price * Occ DC				-0.168***
				(0.047)
Pension price * Oth DC				-0.188***
				(0.043)
Observations	116468	116468	111438	111438
R^2	0.821	0.820	0.846	0.846

*Notes: * Signifies significance at the 10% level, ** at the 5% level and *** at the 1% level. 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. Data come from the Annual Survey of Hours and Earnings.*

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 our estimates of β and η in Equation 5.2, and we 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 our preferred specification, where we include 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 scheme. 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	0.267***	-0.032**	-0.024	0.003	-0.021	-0.009
	(0.012)	(0.016)	(0.016)	(0.015)	(0.013)	(0.014)
Log disposable income	0.198***	0.083***	0.045***	0.023***	-0.008	0.026***
	(0.008)	(0.009)	(0.009)	(0.008)	(0.007)	(0.008)
Observations	222726	222726	205654	205654	205654	205654
R^2	0.781	0.780	0.836	0.829	0.699	0.759
Price elasticity	0.514	-0.062	-0.045	0.011	-0.210	-0.059
Income elasticity	0.381	0.159	0.086	0.089	-0.082	0.170

*Notes: * Signifies significance at the 10% level, ** at the 5% level and *** at the 1% level. 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* $\log(p^f_{it})$ *with* $\log(p_{it})$*. Robust standard errors clustered at the employee level. Data come from the Annual Survey of Hours and Earnings.*

To summarise, in our preferred specifications (column (3) in both Tables 5.1 and 5.2), we 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, we 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 membership by 1% and an increase in pension contributions by 2.5%. Taking an average employee earning £60,000 a year and contributing £3,000 into their pension, 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, we 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. If we assume that 50% of employees did in fact save 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% we had assumed[12]. This would increase our estimated average elasticity by one third, implying a total elasticity of around -0.2.

[12]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.

In Tables A.1 and A.2, we test how sensitive our results are to the restriction that employees in the sample have real annual gross earnings between £30,000 and £70,000. Specifically, we 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 Regression results, 2013-19

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. Our 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. This is despite the fact that in the 2005-12 regression we calculate the tax price assuming no employee saves using a salary sacrifice scheme, meaning we underestimate the coefficient slightly, as discussed in Section 5.3. The elasticity in column (3) is, however, not underestimated in this way since we can observe whether the employee has a salary sacrifice scheme. Column (4) shows that our point estimates of the price elasticity are slightly larger for DC schemes than for DB schemes, but they are also insignificant and 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	0.041*	-0.009	-0.022	
	(0.023)	(0.052)	(0.052)	
Log disposable income	0.768***	0.749***	0.649***	0.642***
	(0.019)	(0.025)	(0.025)	(0.024)
Pension price * Occ DB				-0.011
				(0.067)
Pension price * Occ DC				-0.059
				(0.062)
Pension price * Other DC				-0.037
				(0.053)
Observations	145598	145598	136332	136332
R^2	0.845	0.845	0.875	0.876

*Notes: * Signifies significance at the 10% level, ** at the 5% level and *** at the 1% level. 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. Data come from the Annual Survey of Hours and Earnings.*

Table 5.4 presents the estimate of Equation 5.2 for the 2013-19 sample, where we include all private-sector employees within the £30,000-£70,000 annual earnings range. Column (3), our 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 we 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	-0.612***	-0.009	-0.005	0.047**	-0.034	-0.021
	(0.014)	(0.025)	(0.025)	(0.019)	(0.023)	(0.027)
Log disposable income	-0.186***	0.063***	0.043***	0.048***	-0.019	0.012
	(0.011)	(0.014)	(0.014)	(0.010)	(0.013)	(0.015)
Observations	198414	198414	183640	183640	183640	183640
R^2	0.643	0.636	0.689	0.825	0.717	0.718
Price elasticity	-0.817	-0.012	-0.007	0.239	-0.174	-0.062
Income elasticity	-0.249	0.084	0.057	0.242	-0.096	0.034

*Notes: * Signifies significance at the 10% level, ** at the 5% level and *** at the 1% level. 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. Data come from the Annual Survey of Hours and Earnings.*

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. In some ways this is perhaps unsurprising given the large impact that automatic enrolment was having on the UK pension saving landscape during this time period. The proportion of employees saving in a workplace pension increased dramatically in this period, so that the difference in pension membership rates between workers depends much less on earnings than previously (Bourquin et al., 2020). Furthermore, 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.

5.5 The responsiveness of employer contributions

Up to now we have focused solely on how employee pension contributions respond to the change in the tax price at the higher rate threshold. In this section, we 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 we find a low elasticity for employee pension contributions, this matching mechanism is unlikely to lead to a large elasticity for employer contributions. A 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[13]. Note that since our preferred specification includes employee-employer fixed effects, throughout we 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 5.5 shows the estimated elasticity of employer contributions to the upfront tax price of pension saving. Throughout, we instrument the log price of pension saving using the first-pound price, and include year and employee-employer fixed effects. We estimate very small intensive- and extensive-margin elasticities that are not 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 we find for employees, who have a larger incentive to respond.

[13]Note that here we are referring to employer contributions decided by the employer. In our analysis, employer contributions do *not* include employee contributions made through a salary sacrifice agreement (which are treated as employer contributions for tax purposes.)

Table 5.5: *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	-0.042	-0.019	0.023	0.001
	(0.030)	(0.014)	(0.046)	(0.025)
Log disposable income	0.543***	0.057***	0.658***	0.029**
	(0.018)	(0.008)	(0.023)	(0.014)
Observations	126539	205654	144776	183640
R^2	0.879	0.854	0.919	0.660
Price elasticity	-0.042	-0.032	0.023	0.002
Income elasticity	0.543	0.096	0.658	0.036

*Notes: * Signifies significance at the 10% level, ** at the 5% level and *** at the 1% level. 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. Data come from the Annual Survey of Hours and Earnings.*

6 Conclusion

Many governments seek to use a variety of mechanisms to encourage saving for retirement. It is important to understand how effective these are. On the whole the existing empirical evidence suggests individuals are relatively insensitive to price incentives, whether that is in the form of employer matches (Choi, 2015) or tax incentives (Chetty et al., 2014; Selin, 2012). However, given that the exact size of responses are likely to vary according to the population examined and the salience of the incentive in question, further empirical research is beneficial to understand how generalisable these findings are to different settings, and what drives the responsiveness in different settings.

We have examined the response of UK employees to the incentive to save in a private pension generated by upfront relief of pension contributions from income tax. One attractive feature of our setting that we focus on employees, unlike much of the past literature, who make up the majority of workers in all countries. It is also directly relevant to policy debate in the UK around whether the rate of upfront income tax relief given to those with different levels of earnings is appropriate. Furthermore, by repeating the analysis before and after the introduction of automatic enrolment in the UK, we show how this elasticity is affected by other pension saving policies.

Consistent with the existing literature, we find that the responsiveness of pension saving with respect to the tax price is low – with an intensive margin tax price elasticity of -0.1 and an extensive elasticity of -0.05 for the period 2005 to 2012. Over the following period, where the proportion of employees saving in a pension was greatly expanded as a result of automatic enrolment into private pensions, the elasticities are estimated to be near zero, indicating that the elasticity can be affected by features of the policy landscape that determine who is a pension saver.

Therefore, the additional cost of upfront tax relief accruing to savers around the higher rate threshold arises mainly from the mechanical effect of their higher tax rate and their propensity

to save more because they earn more, rather than because of a behavioural response to the tax incentive. One avenue for future research would be to examine how this responsiveness differs at other parts of the earnings distribution, and how it is affected by other changes to the pension saving landscape that affect the composition of pension savers.

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.

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.

Bourquin, P., J. Cribb, and C. Emmerson (2020). Who leaves their pension after being automatically enrolled? *The Institute for Fiscal Studies - Briefing Note 272*.

Card, D. and M. Ransom (2011). Pension plan characteristics and framing effects in employee savings behavior. *The Review of Economics and Statistics 93*(1), 228–243.

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.

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

Choi, J. J., D. Laibson, and B. C. Madrian (2011). $100 Bills on the Sidewalk: Suboptimal Investment in 401(k) Plans. *The Review of Economics and Statistics 93*(3), 748–763.

Choi, J. J., D. Laibson, B. C. Madrian, and A. Metrick (2002). Defined contribution pensions: Plan rules, participant choices, and the path of least resistance. *Tax Policy and the Economy 16*, 67–113.

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.

Duflo, E., W. Gale, J. Liebman, P. Orszag, and E. Saez (2006). Saving Incentives for Low- and Middle-Income Families: Evidence from a Field Experiment with H&R Block. *The Quarterly Journal of Economics 121*(4), 1311–1346.

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.

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

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.

Office for National Statistics (2021). Annual survey of hours and earnings, 1997-2020: Secure access. 18th Edition. UK Data Service. SN: 6689.

O'Neil, C. J. and G. R. Thompson (1987). Participation in individual retirement accounts: An empirical investigation. *National Tax Journal 40*(4), 617–624.

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

Selin, H. (2012). Marginal tax rates and tax-favoured pension savings of the self-employed: Evidence from sweden. *The Scandinavian Journal of Economics 114*(1), 79–100.

Venti, S. F. and D. A. Wise (1988). The Determinants of IRA Contributions and the Effect of Limit Changes. In *Pensions in the U.S. Economy*, NBER Chapters, pp. 9–52. National Bureau of Economic Research, Inc.

A Appendix tables

Table A.1: *Testing the sensitivity of our 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	-0.095***	-0.106**	-0.179***	-0.041
	(0.032)	(0.041)	(0.069)	(0.026)
Log disposable income	0.468***	0.451***	0.373***	0.541***
	(0.017)	(0.027)	(0.058)	(0.011)
Observations	111438	79453	49203	173504
R^2	0.846	0.832	0.830	0.876

*Notes: * Signifies significance at the 10% level, ** at the 5% level and *** at the 1% level. 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.*

Table A.2: *Testing the sensitivity of our 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	-0.024	-0.035*	-0.022	0.003
	(0.016)	(0.019)	(0.027)	(0.013)
Log disposable income	0.045***	0.047***	0.084***	0.056***
	(0.009)	(0.014)	(0.027)	(0.005)
Observations	205654	138811	83347	380679
R^2	0.836	0.833	0.830	0.848
Price elasticity	-0.045	-0.067	-0.043	0.005
Income elasticity	0.086	0.090	0.162	0.109

*Notes: * Signifies significance at the 10% level, ** at the 5% level and *** at the 1% level. 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.*