Institute for Fiscal Studies

Laurence O'Brien

Working paper 24/45

The effect of tax incentives on retirement saving

Economic and Social Research Council

The Effect of Tax Incentives on Retirement Saving*

Laurence O'Brien†

October 15, 2024

Abstract

This paper estimates the responsiveness of retirement saving to tax incentives for employees in Great Britain and shows how this responsiveness interacts with automatic enrolment into workplace retirement plans. For identification, I exploit a kink in the income tax schedule where there is a large, discontinuous change in the marginal price of retirement saving. Prior to the introduction of automatic enrolment, I estimate an intensive-margin elasticity of -0.21 and an extensive-margin elasticity of -0.12, suggesting that employees respond only weakly to this tax incentive to save even in a world without defaults. In 2013 to 2019, after the introduction of automatic enrolment, I find an even lower elasticity. I show this is because the large number of savers automatically enrolled into retirement accounts are entirely unresponsive to the tax incentive, even though they are potentially precisely the group policymakers might worry are undersaving for retirement. This finding has important implications for the targeting of retirement saving incentives.

Keywords: retirement saving; incentive effects of taxation; automatic enrolment
JEL classification: D14, H24, J32

*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. I would like to thank Rowena Crawford, Jonathan Cribb, Carl Emmerson, Eric French, Cormac O'Dea, Fabien Postel-Vinay, David Sturrock and Eric Zwick 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 many countries around the world, governments subsidise savings into retirement accounts to encourage individuals to shift sufficient resources into old age. The government revenue foregone from this tax advantage is often sizeable, and the benefits tend to accrue disproportionately to relatively higher earners. Given this, it is unsurprising that the tax treatment of retirement saving is a common topic of public and policy debate. Key to understanding the merits of current systems, 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 retirement saving in recent years (Choi, 2015), there is still uncertainty about the extent to which these tax subsidies raise total wealth accumulation (Bernheim, 2002; Friedman, 2017) and, in particular, how they interact with other retirement saving policies such as automatic enrolment.

In this paper, I estimate how responsive private-sector employees are to tax incentives to save in defined contribution (DC) retirement accounts, 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 retirement accounts. 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 tax price of saving to examine how retirement saving responds to this price on both the intensive and extensive margin.

The UK represents a particularly suitable setting for this analysis for two reasons. First, the UK is especially reliant on individuals saving privately for retirement, as the UK public pension provides a fairly low replacement rate for the majority of individuals compared to many other developed economies (OECD, 2021). Second, between 2012 and 2018 the UK rolled out a nationwide policy whereby employers are now obliged to automatically enrol employees into workplace retirement accounts, allowing me to analyse the interaction between tax incentives and nudges in encouraging retirement saving.

To assess the extent to which the responsiveness of retirement saving tax incentive changes when employees are automatically enrolled into retirement accounts, 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 retirement accounts. There are two reasons why we might expect a different intensive-margin elasticity of retirement saving contributions to the tax price after automatic enrolment. The first is that the new employees brought into retirement saving by automatic enrolment might well respond differently to tax incentives than existing savers. The second is that automatic enrolment might directly change the responsiveness of existing savers to tax incentives, for example because the new default contribution rate acts as an implicit recommendation as to how much to save. To disentangle these, I use samples of existing and new savers to understand the mechanisms through which automatic enrolment affects the measured elasticity.

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 contributions to DC retirement accounts at the point where the tax price of saving, that is, the amount of disposable income foregone by contributing one extra pound to a retirement account, decreases discontinuously. On the intensive margin, I show that those employees who are saving in a DC retirement account are not choosing the amount they save in such a way as to bunch their taxable income at the kink in the tax schedule. Importantly, this kink constitutes a discontinuity in both the upfront tax price of retirement saving (accounting only for taxes saved today by contributing into a retirement account) and the overall, long-run, price of retirement saving (which also accounts for taxes paid in retirement when funds are withdrawn). This is because taxes on funds saved in UK retirement accounts are deferred until retirement in such a way that the marginal tax rate on funds when they are withdrawn is not directly linked to the tax avoided by saving in these accounts during working life.[1] These results therefore imply little to no extensive- or intensive-margin responsiveness of retirement saving to tax incentives.

I confirm the lack of responsiveness suggested by the graphical analysis using a panel regression approach. I instrument the tax price of retirement saving using the standard first-pound price suggested by Feldstein and Taylor (1976) (i.e. the price of retirement saving if the individual saved nothing in their account). This addresses the endogeneity

[1]Indeed, Adam et al. (2023) show that the tax rate faced by UK retirees varies little, and that 95% of working-age individuals today can expect to have an average tax rate of at most 20% when withdrawing funds from retirement saving accounts, given the current tax system. The marginal tax rate faced by working-age individuals varies much more.

of the tax price arising from the fact that employees can reduce their taxable income, and therefore potentially their marginal tax rate, by increasing their retirement saving. In addition, the regression controls for (logged) earnings and includes employee-employer fixed effects, meaning that 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 retirement accounts, the estimate of the intensive-margin price elasticity is -0.21, while the estimated extensive-margin elasticity is -0.12. The total elasticity estimated for this time period is therefore -0.33, 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 respond strongly to this tax incentive to save, and indeed the responsiveness I estimate is lower than in other papers in the literature.

To understand why the elasticity is lower in the latter time period, I separately analyse the intensive-margin responsiveness of two groups of savers: 'existing' savers, who I observe saving both before 2012 as well as after they become eligible for automatic enrolment, and 'new' savers, who I observe not saving before 2012, but who do start saving after they are automatically enrolled. I show that the intensive-margin elasticity of existing savers is very similar both before and after automatic enrolment and similar to the elasticity estimated on the full sample of savers in the 2005-12 period. In contrast, new savers' elasticity is essentially zero. Therefore, the reduction in the intensive-margin elasticity since the introduction of automatic enrolment is driven entirely by the fact that those brought into retirement saving by the introduction of automatic enrolment are even more passive savers than those saving in a retirement account prior to automatic enrolment.

These findings have important implications for the welfare effects of different methods of encouraging retirement saving. Choukhmane and Palmer (2023) show that these welfare effects are higher for policies that increase retirement saving without crowding out other forms of saving, particularly among those individuals who are undersaving for retirement. The tax incentive I study encourages very little extra retirement saving overall, making the extent of crowdout, which I cannot observe in my data, essentially irrelevant.

The small amount of extra retirement saving as a result of the tax incentive comes from those who were already saving before the introduction of automatic enrolment, while I find an elasticity of essentially zero for new savers. These new savers are, however, potentially precisely the group that policymakers might worry are undersaving for retirement, suggesting that this tax incentive is particularly poorly targeted.

This paper builds on a growing literature that has studied how savings respond to tax incentives, or financial incentives more broadly. Many papers studying this question find that tax incentives for retirement saving often 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 (Attanasio et al., 2004; Gelber, 2011; Chetty et al., 2014; Andersen, 2018; Lavecchia, 2019). However, other papers find rather less evidence of crowding out (Goodman, 2020; Chan et al., 2022; Christensen and Ellegaard, 2023), indicating that tax incentives can be an effective tool for stimulating total savings in certain settings. On the other hand, in other cases it seems that tax incentives do not lead to a substantial increase in saving even in the subsidised account, never mind overall. Duflo et al. (2006) randomise match rates to an IRA account as part of a randomised field experiment in the St. Louis metro area, while Engelhardt and Kumar (2007) use HRS data and a life-cycle-consistent discrete choice regression to estimate the effect of match rates on employee 401(k) contributions. Despite their different methodologies, both papers find that an increase in the match rate by 20 or 25 percentage points increased IRA/401(k) participation by only around 5 percentage points.

The existing literature therefore exhibits a wide variety of estimates on the effectiveness of tax incentives in increasing retirement savings. This paper studies a financial incentive for retirement saving that is an implicit feature of the tax system in the UK, and finds a much lower responsiveness of retirement saving to the incentive than even Duflo et al. (2006) and Engelhardt and Kumar (2007): participation in DC retirement accounts increases by only around 1 percentage point at the higher rate threshold in the earlier period I analyse, where the "match rate" increased by 20 percentage points. This contrasts with previous evidence from how financial incentives affect saving in the UK. Attanasio et al. (2004) show that new tax-advantaged savings accounts were taken up in large numbers in the UK, even if relatively small fractions of the funds in these accounts represented new saving. These results suggest 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).

The second contribution of this paper is to analyse how tax incentives for retirement saving interact with another major retirement savings policy: automatic enrolment, which has been shown to have a large impact on both participation and contributions in a variety of settings, including the UK (Cribb and Emmerson, 2020).

Some models of retirement saving decisions posit two types of individuals: 'active' savers, whose saving responds to financial incentives, and 'passive' savers, whose saving does not deviate from the default, and so would respond to automatic enrolment but not financial incentives (Chetty et al., 2014). In reality, automatic enrolment could affect the passivity of (potential) savers. On the one hand, automatic enrolment introduces a notable focal point for the contribution rate into the system, identical across the UK, which could be interpreted by some as a recommendation for how much they should be saving for retirement. Previous research has highlighted the potential importance of anchoring and cues for individuals' saving decisions, including arising from automatic enrolment (Madrian and Shea, 2001; Bernheim et al., 2015; Choi et al., 2017), and it is possible that savers might be less likely to respond to changes in the price of saving if the anchor is strong enough. This paper therefore contributes to the literature by measuring whether existing savers' responsiveness to tax incentives differs in environments with and without automatic enrolment.

On the other hand, it is also possible that new savers due to automatic enrolment then choose their contribution amount taking financial incentives into account. Indeed, Cribb and Emmerson (2020) show that many new savers do not save at the default rates after being automatically enrolled, suggesting they might not be wholly passive. However, I find little evidence that the price of retirement saving affects the contribution decision of new savers after they have been enrolled, suggesting limited complementarity between automatic enrolment and tax incentives in encouraging saving rates among more passive savers.

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 retirement 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 Retirement saving in the UK and automatic enrolment

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 if they want to smooth their living standards in retirement. Private retirement accounts are a particularly common and, as explained later in this section, tax-advantaged way to do this additional saving.

While these retirement accounts can be arranged by individuals, the vast majority of retirement saving for employees is done in accounts facilitated by employers. For public-sector employees these retirement accounts are typically 'defined benefit' (DB) pensions, while for private-sector employees, they are now predominantly 'defined contribution' (DC) plans. Throughout my analysis, I focus on savings into DC accounts, because employees with these accounts can typically choose their level of contributions (which may be influenced by matching arrangements provided by some employers). In 2005, the start of the period I examine, 44% of private-sector employees were saving in a workplace retirement account: 24% in a DB pension and 20% in a DC plan.

Due to concerns about low private saving for retirement, between 2012 and 2018 the UK government rolled out a policy of automatic enrolment into workplace retirement accounts. The policy was rolled out by employer size from 2012, starting with the largest employers. Once an employer reached their deadline for implementing automatic enrolment, they had to automatically enrol all 'targeted' employees into a workplace retirement account with (at least) minimum levels of contributions, some coming from the employer. Targeted employees are those aged between 22 and state pension age, earning more than a given earnings level (set at £10,000 annually in nominal terms since April 2014), and, in practice, with a job tenure of at least three months. I refer to targeted employees who work for an employer that has reached their deadline for implementing automatic enrolment as employees who are 'eligible' for automatic enrolment. Employers can choose the default contribution rates at which they enrol their employees, subject to minimum rates. These minimum default contribution rates were at least 2% of qualifying pay initially (with at least 1% of pay coming from the employer), but in April 2018 increased to 5% of qualifying pay (with at least 2% coming from the employer) and in April 2019 to 8% of qualifying pay (with at least 3% coming from the employer). See Cribb and Emmerson (2020) and Cribb and Emmerson (2021) for more information on the UK's automatic enrolment policy.

Automatic enrolment substantially increased participation in retirement saving accounts (Cribb and Emmerson, 2020, 2021). In 2019, the end of the period I examine, 79% of private-sector employees were saving in a retirement account (12% in a DB pension and 68% in a DC plan).

2.2 Tax treatment of retirement saving

The UK income tax system is based on individual income and treats retirement savings 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 income when it is drawn from the account in retirement. Retirement savings are also relatively tax-favoured compared to other forms of income because employer contributions to accounts are not subject to another payroll tax (National Insurance), 25% of funds in the account can be withdrawn free of income tax in retirement, and these accounts are taxed relatively lightly on death. Adam et al. (2023) estimate that these tax advantages are worth around £4.4 billion per year compared to a pure EET system (i.e. one where income contributed to retirement accounts is not taxed upfront, the returns are not taxed, but the income is taxed upon withdrawal from the account). Relative to a pure TEE system (i.e. one where income contributed to retirement accounts is taxed upfront, but the returns are not taxed and the income is not taxed upon withdrawal) the current pension tax system is more generous to the tune of around £46 billion a year.

There are some restrictions on the amount that can be saved into a retirement saving account 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 samples, who earn between £30,000 and £70,000 per year (in 2019 prices). From 2014 to the end of my sample period, for employees in the income range I analyse, the annual allowance limits tax-free contributions to £40,000 per year. Prior to 2014, the annual allowance was even higher. Employees in my sample would therefore have to save nearly all of their earnings to get close to hitting the annual allowance. The lifetime allowance, set at £1,073,100 in 2021-22, determines the maximum total value of funds in an individual's retirement account before attracting high tax rates and is again sufficiently high to not affect my sample's retirement saving decisions.

The non-linearity of the income tax schedule means that there are sharp discontinuities in the upfront income tax relief on contributions. In particular, in the sample of incomes that I analyse in this paper, there is one nonlinearity in the income tax schedule: the "higher rate threshold", where the marginal income tax rate changes from the "basic" rate (22% in 2007-08 and earlier, 20% since) to the "higher" rate (40%).[2]

The real level of the higher rate threshold varies during my analysis period, fluctuating between £45,000 and £55,000. This implies 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 retirement saving around the higher rate threshold, the estimate is relevant to well-off, but not super-rich, individuals by UK standards.

The higher rate threshold represents a large discontinuity in the upfront price of saving in DC retirement accounts. This will also reflect a large discontinuity in the overall long-run price of saving as long as the discontinuity in the upfront tax relief on contributions does not correspond exactly to a discontinuity in the marginal tax rate paid on funds when they are withdrawn in retirement. This will almost certainly be the case; in fact, given the current income tax system, Adam et al. (2023) estimate that 86% of current higher-rate taxpayers will be at most basic-rate taxpayers in retirement.

Mirrlees et al. (2011) provide an example calculation showing how the long-run tax incentive to save in a pension increases for people who go from basic-rate to higher-rate taxpayers in working life, while remaining basic-rate taxpayers in retirement. They estimate that employee contributions into a retirement account held for 25 years have an effective tax rate of -8% for those who were basic-rate taxpayers when contributing, compared to -48% for those who are higher-rate taxpayers while contributing. The change in the long-run tax price of saving at the higher rate threshold is therefore substantial.

[2]Since 2017-18, Scotland has had a slightly different tax schedule to the rest of the UK. Since 2018-19, in Scotland there is an "intermediate" marginal tax rate of 21% levied on incomes between approximately £25,000 and the higher rate threshold, and a higher rate of 41% rather than 40%. The level of the higher rate threshold is also slightly lower than in the rest of the UK. I can observe whether an employee is resident in Scotland in my data, so I account for the different tax schedules in my analysis.

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 retirement saving information. The sample frame for the ASHE is always the same 1% random sample of employees in Great Britain (based on the last two digits of the employee's National Insurance number), meaning that it has a large sample size of around 180,000 individuals per year, and can be used for longitudinal analysis. Responding to the survey is a statutory requirement for employers, therefore levels of non-response are low.

The outcome of interest is whether an employee is saving in a workplace DC retirement saving account and, if they are, the monetary value of their contributions. In the ASHE, employers report the monetary value of employer and employee contributions to the employee's retirement saving account made during the pay period. I can observe if the individual has several employee jobs and calculate total contributions made to all workplace retirement accounts. The ASHE does not collect information on contributions to personal retirement accounts that are made independently by individuals, but this is uncommon among employees.[3]

I also require a measure of individual income, both for calculating the tax price of retirement saving, and because income itself affects retirement 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 income from other sources such as income from self-employment, rental income from property, or income from saving and investments. As a result, I use annual earnings from all employment to proxy annual income throughout this paper. I have analysed data from the Survey of Personal Incomes and found that around the higher rate threshold employee earnings make up over 95% of income for over 80% of employees. This implies that this proxy will be accurate for the vast majority of employees, although it could cause some attenuation bias in my results.

One limitation of this data source is that I cannot observe saving in other types of

[3]Data from the Family Resources Survey suggests that only 5% of employees were saving in a personal retirement account in 2021-22. See https://www.gov.uk/government/statistics/family-resources-survey-financial-year-2021-to-2022.

assets. The effect of retirement saving policies on overall saving is crucial for evaluating the desirability of these policies (Friedman, 2017; Choukhmane and Palmer, 2023), and indeed many papers focus on how increases in savings in incentivised accounts crowd out other forms of saving (e.g. Chetty et al. (2014); Christensen and Ellegaard (2023); Choukhmane and Palmer (2023)). However, importantly, in my setting I find very little responsiveness of retirement saving to tax incentives for such saving. This implies there is very little extra saving to be 'crowded out' in the first place.

Throughout, the main analysis samples consist 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 retirement savings; robustness tables in the appendix show this restriction has limited effect on the empirical results.

Table 3.1 presents the summary statistics for the two samples: the 2005-12 and the 2013-19 sample. The proportion of the sample saving in a DC retirement account is significantly higher in the latter period, consistent with the impact of automatic enrolment. Most, but not all, employees in the sample who are participating in a workplace DC retirement account have strictly positive employee contributions (I classify employees with zero employer and employee contributions as not saving in a workplace retirement account). Conditional on making positive contributions, employees contribute on average around 4% of gross pay to their DC retirement account. Average conditional employer contributions reduced from around 7% of gross pay in the 2005-12 period to slightly over 5% of pay by the 2013-19 period.

Table 3.1: Sample Summary Statistics

Characteristic	2005-12 Sample	2013-19 Sample
% saving in workplace retirement account	60.4	80.3
% saving in DC retirement account	31.1	58.4
% with DC employee contributions > 0	25.2	54.0
% with DC employer contributions > 0	30.0	57.5
Average DC employee contribution	4.1	3.7
Average DC employer contribution	6.9	5.4
% women	28.1	30.6
% aged 22-34	30.2	29.6
% aged 35-49	48.3	45.1
% aged 50-59	21.5	25.3
Observations	250,706	227,145
People-jobs	97,810	92,593
People	78,924	77,114
People with >1 job	15,803	13,329

Notes: This table shows the characteristics of my two analysis samples. Both samples contain 22- to 59-year-old private-sector employees with annual gross earnings between £30,000 and £70,00 in real (2019) terms. The first sample contains observations between 2005 and 2012, while the second sample contains observations between 2013 and 2019. Average DC employee and employer contributions are measured as a percentage of gross earnings among employees with strictly positive contributions. People-jobs refers to the amount of unique employee-employer observations, people refers to the amount of unique employees in the sample, and people with >1 job refers to the amount of unique employees in the sample observed working for two or more different employers.

3.1 Calculating the tax price of retirement saving

In the empirical analysis the main independent variable of interest will be the upfront tax price of retirement saving. This is how much contemporaneous disposable income the employee forgoes by contributing one more pound to their retirement account, calculated as:

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

where $T(y)$ the total income tax paid by an individual with annual taxable earnings y, and with $\Delta s = 1$. In other words, as employee contributions to retirement accounts are exempt from income taxes, one extra pound contributed to a retirement account will reduce an employee's disposable income by £$1-\tau_I$ today, where τ_I is the marginal income tax rate.

In the UK, the tax price of retirement saving can also be affected by how employee contributions are made, that is, whether they are made through 'salary sacrifice' or not.

Under a salary sacrifice arrangement, employees agree with their employer to reduce their earnings by an amount equal to their desired employee contributions to the retirement account, and for these contributions to be made as employer contributions instead. This is advantageous as, unlike employee contributions to retirement accounts, employer contributions are not subject to another payroll tax (National Insurance). Therefore, for employees with salary sacrifice arrangements, one pound contributed to their retirement account 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. 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. Specifically, I calculate the tax price using Equation 3.1, but where $T(y)$ denotes the total income tax and national insurance paid by someone with earnings y. For earlier years, I show how the results differ when I assume no one saves using salary sacrifice, when everyone saves using salary sacrifice, and when a randomly selected half of employees save using salary sacrifice (my preferred approach given that in 2013 50% of savers into DC retirement accounts used salary sacrifice).

In essence, individuals in the sample have one of four possible (upfront) tax prices. Basic rate taxpayers without salary sacrifice have a tax price of £0.80 or £0.78 (depending on the year); basic rate taxpayers with salary sacrifice have a tax price of £0.68; higher rate taxpayers without salary sacrifice have a tax price of £0.6; and higher rate taxpayers with salary sacrifice have a tax price of £0.58.[4] Therefore, employees saving without salary sacrifice face a 25% drop in the price of retirement saving at the higher rate threshold, while those saving with salary sacrifice face approximately a 15% drop. Either way, clearly, there is a large discontinuity in the tax price at the higher rate threshold whether employees save using salary sacrifice or not.

4 Graphical evidence

I start by showing graphically how participation in and contributions to DC retirement accounts vary around the higher rate tax threshold. I analyse the periods 2005 to 2012 and 2013 to 2019 separately to explore whether the introduction of automatic enrolment

[4] Taxpayers in Scotland can have slightly different tax prices after 2017/18 when the tax schedules diverge.

(rolled out from late 2012 onwards) has any discernable effect on saving behaviour.

Figure 4.1 plots the proportion of employees with strictly positive contributions to DC retirement accounts 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 DC retirement account increases discontinuously above the higher rate threshold. For example, for those without a salary sacrifice agreement 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. 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 DC retirement account above the higher rate threshold in either period. This suggests a low extensive-margin responsiveness to changes in the tax price of retirement saving.

Figure 4.1: Participation in DC retirement accounts around the higher rate threshold, 2005-12 vs. 2013-19

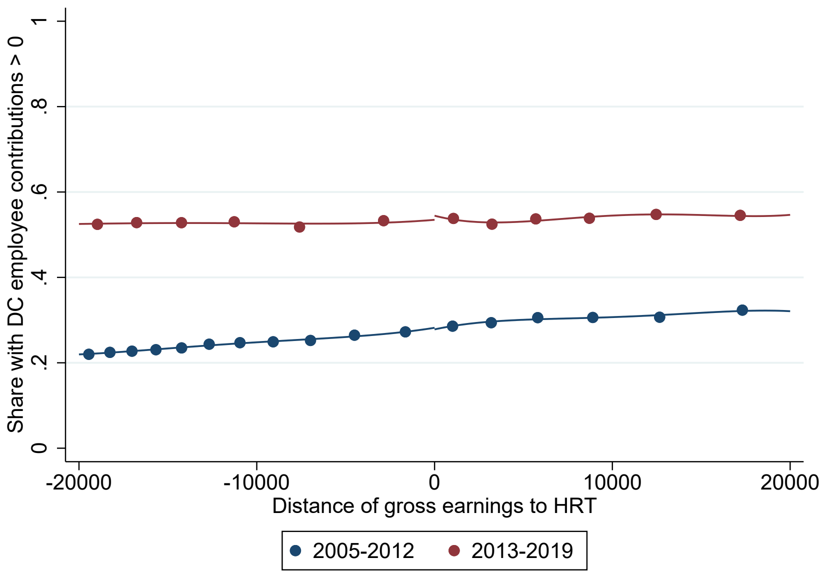

Notes: This figure shows the proportion of employees with strictly positive employee contributions to a DC retirement account around the higher rate threshold. The higher rate threshold is the level of taxable income at which the marginal tax rate increases from the 'basic' rate to the 'higher' rate. Dots show these proportions by bins of gross earnings (normalised relative to the HRT), separately for two sample periods. Bins are quantile-spaced and selected to minimise mean-squared error as outlined in Cattaneo et al. (2020). Lines show an estimated fourth-order polynomial fit to the data separately either side of the HRT for both periods.

In Appendix Figure C.1 I show that there is also no discontinuity in the share of employees making a positive contribution to any (DC or DB) workplace retirement account in either period.[5] In addition, Appendix Figure C.2 demonstrates that there is also no discontinuity in the share of employees making positive contributions to DC retirement accounts among those employees whose gross earnings have been persistently over or under the higher rate threshold for at least one year.

To analyse whether DC 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 does not save via salary sacrifice. To start with, each pound contributed to their DC retirement account 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. The same logic holds for salary sacrifice savers (although the exact numbers differ).

In Figure 4.2, I therefore investigate the degree to which employees are making DC 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 retirement account contributions around the higher rate threshold for employees with strictly positive DC contributions. Of course, there could also be bunching in this variable at the threshold for reasons unrelated to retirement 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' retirement 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 retirement saving (and gross earnings) to the tax price.

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

Figure 4.2: Bunching of earnings around the higher rate threshold, 2005-12

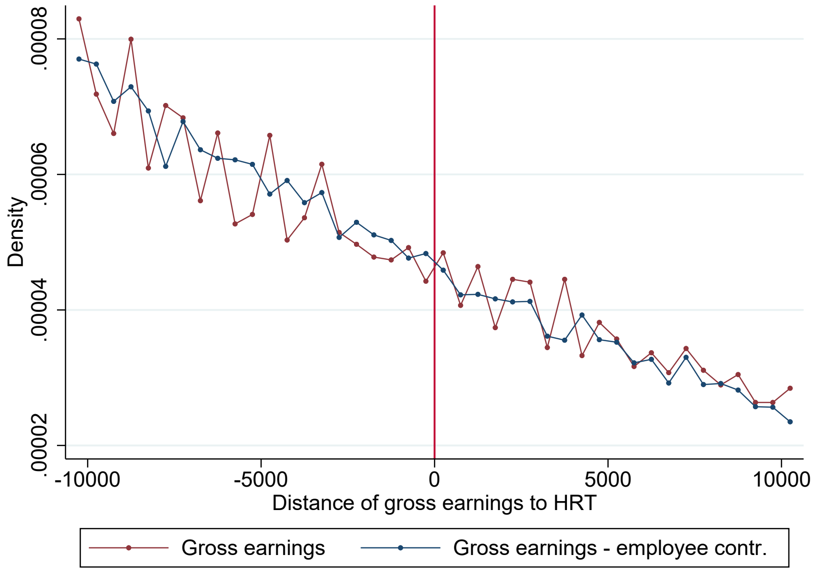

Notes: This figure shows the extent of bunching of gross earnings (in red) and gross earnings after subtracting employee contributions to DC retirement accounts (in blue) around the higher rate threshold. The higher rate threshold is the level of taxable income at which the marginal tax rate increases from the 'basic' rate to the 'higher' rate. The lines plot the frequency density of each variable in bins with width of £500. The sample is the 2005-12 analysis sample restricted to employees with gross earnings (after subtracting DC contributions) of at most £10,500 away from the level of the higher rate threshold, and with strictly positive DC employee contributions.

The equivalent figure for the 2013-19 period, Figure 4.3, perhaps shows a small degree of bunching in gross earnings just below the higher rate threshold. However, importantly, there is no evidence of bunching in gross earnings after subtracting employee retirement account contributions, 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.

In Appendix Figures C.3 and C.4, I repeat these figures for the sample of employees whose gross earnings were above the higher rate threshold in both year t and year $t-1$. In other words, I restrict my attention to those employees who have been over the higher rate threshold for at least a year and have therefore had more time to change their pension contributions in response to tax incentives. There is a slightly lower density of period t gross earnings just above the higher rate threshold in both graphs as most employees experience positive growth in real earnings and so typically those with earnings just above the threshold in year t had earnings below the threshold in year $t-1$. There is no evidence, however, that these employees with more time to respond to tax incentives use retirement account contributions in order to bunch their taxable income.

Figure 4.3: Bunching of earnings around the higher rate threshold, 2013-19

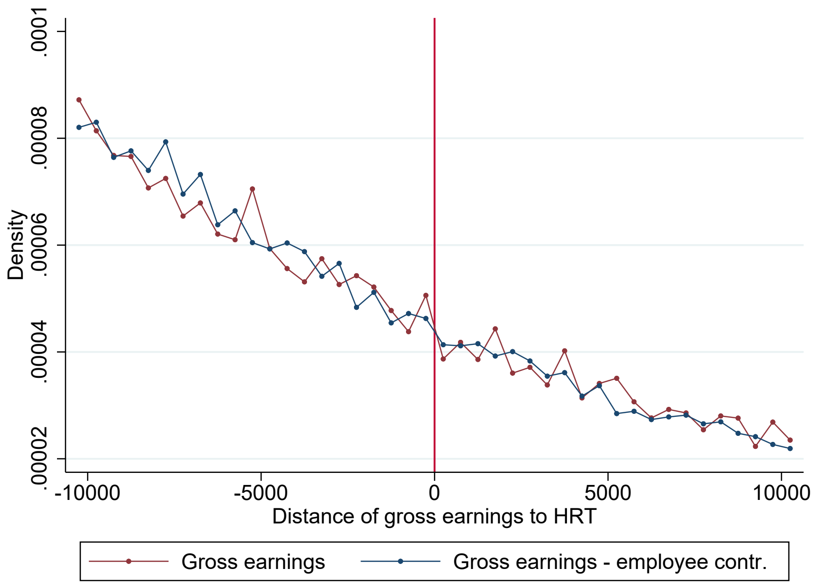

Notes: This figure shows the extent of bunching of gross earnings (in red) and gross earnings after subtracting employee contributions to DC retirement accounts (in blue) around the higher rate threshold. The higher rate threshold is the level of taxable income at which the marginal tax rate increases from the 'basic' rate to the 'higher' rate. The lines plot the frequency density of each variable in bins with width of £500. The sample is the 2013-19 analysis sample restricted to employees with gross earnings (after subtracting DC contributions) of at most £10,500 away from the level of the higher rate threshold, and with strictly positive DC employee contributions.

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 retirement saving of individual i working for employer j in year t depends both on the upfront tax price of retirement saving, p_{it}, and post-tax earnings, y_{it}.

To estimate the intensive-margin responsiveness of employees' DC contributions to the tax price of retirement saving, I estimate the following equation when DC 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 DC contribution of employee 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 and the employee's job tenure in discrete bins in X_{it}.[6] Then, ε_{INT} is the intensive-margin elasticity of retirement saving to the upfront tax price, and η_{INT} is the intensive-margin income elasticity of retirement saving. Note that y_{it} is post-tax earnings for employee i if they made no contribution to their DC account, i.e. $y_{it} = T(z_{it})$, so it does not depend on s_{it}.

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 DC 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 DC employee contributions are strictly positive. In the appendix, I also show how the results change if I redefine the outcome variable as whether they have a positive employee contribution to any type of retirement saving account (either DB or DC). This is because some employees might be offered a DB pension by their employer as opposed to a DC retirement account. I show that this leads to a smaller estimated elasticity.

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 retirement saving 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 retirement saving, p_{it}, with the "first-pound" price of retirement-saving, p^f_{it} (Feldstein and Taylor, 1976). In this setting, this means the instrument is the

[6] The bins for job tenure denote whether the employee's job tenure is less than 3 months, between 3 months and one year, between one and two years, between two and five years, between five and ten years, or over ten years.

tax price of retirement saving that the employee would have faced had he or she not done any retirement saving:

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

This instrument is highly correlated with the "last-pound" price of retirement saving, since most employees do not contribute so much to their DC retirement account as to change their marginal tax band. Indeed, all first-stage regressions (for both the intensive and the extensive margin, and in both periods) have an F-statistic of over 1,500, indicating that the instrument is strong. 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 retirement account (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 have to assume employees are not suddenly more motivated to earn more because they decide they want to increase the amount they save into their DC retirement account. 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 contributions to retirement accounts, which could potentially affect employee 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 DC retirement account, given their post-tax earnings and the relative tax price of retirement 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 DC retirement account contributions with respect to the upfront tax price of retirement 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 DC contributions for the years 2005 to 2012, either by OLS or IV, and including year fixed effects and either employee or employee-employer fixed effects. The estimates are shown in Table 5.1, with the coefficient on the log price of retirement saving 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 without controls for the square of age and job tenure, and I find a positive estimated elasticity of around 0.25. 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 retirement contributions and the price of retirement saving.

In column (2) I instrument the log of the "last-pound" price of retirement saving using the log of the "first-pound" price of retirement saving, as defined in Equation 5.3. The estimated elasticity becomes negative at around -0.16. Including employee-employer fixed effects, rather than employee fixed effects, as in column (3), increases the magnitude of the estimate slightly to around -0.17. Also including controls for job tenure and the square of age again very slightly increases the magnitude of the estimated elasticity to -0.18.

In column (5), I show how the estimated elasticity changes if, instead of assuming that no employee saves using salary sacrifice, I assume that all employees save via salary sacrifice. As explained in Section 3.1, employees who save via salary sacrifice face a 15% drop in the tax price of retirement saving at the higher rate threshold, compared to a drop of 25% for employees who do not use salary sacrifice. As expected, the estimated elasticity increases from column (4) to column (5), reaching -0.27.

Column (6) is my preferred specification, where I randomly assign half my sample to save with salary sacrifice and half to save without. This is motivated by the fact that 50% of employees saving in a DC retirement account in 2013 use salary sacrifice. Unsurprisingly, this leads to an estimated elasticity of -0.21, approximately halfway between the estimated elasticities in columns (4) and (5). The 95% confidence interval is (-0.27, -0.16), showing that the estimate is fairly precise.

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

	(1) OLS	(2) IV	(3) IV	(4) IV	(5) IV	(6) IV
Log price of retirement saving	0.245***	-0.167***	-0.173***	-0.178***	-0.265***	-0.212***
	(0.029)	(0.051)	(0.048)	(0.048)	(0.071)	(0.057)
Log post-tax earnings	0.739***	0.583***	0.517***	0.491***	0.513***	0.500***
	(0.022)	(0.027)	(0.026)	(0.027)	(0.023)	(0.025)
Observations	53761	53761	50725	50725	50725	50725
R^2	0.838	0.039	0.033	0.036	0.037	0.034
Controls	No	No	No	Yes	Yes	Yes
Employee FE	Yes	Yes	No	No	No	No
Employee-employer FE	No	No	Yes	Yes	Yes	Yes

Notes: This table shows the estimated intensive-margin price and income elasticity of employee contributions to DC retirement accounts. The cells show the ε_{INT} and η_{INT} coefficients from estimating equation 5.1 on my 2005-12 sample with strictly positive employee DC contributions. All columns contain year fixed effects. Columns (1) and (2) include employee fixed effects, while columns (3) to (6) include employee-employer fixed effects. Columns (4) to (6) also include controls for the square of age and the employee's job tenure. Column (1) is estimated using OLS, while columns (2) to (6) instrument $\log(p_{it})$ with $\log(p_{it}^f)$. Column (4) assumes no employee saves using salary sacrifice, column (5) assumes all employees save using salary sacrifice, and column (6) assumes a randomly-selected 50% of employees save using salary sacrifice. Robust standard errors clustered at the employee level are in parentheses. *, **, 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 retirement saving and the log of post-tax earnings correspond to the estimates of β and η in Equation 5.2, and at the bottom of the table 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, using the "first-pound" instrument gives us a negative elasticity. Adding in employee-employer fixed effects in column (3) and controls for age and job tenure in column (4) again only has a small effect on the estimated elasticity. In column (4), the estimated coefficient is -0.03 and statistically significant at the 10% level, corresponding to a price elasticity of -0.11. In my preferred specification, column (6), where I randomly assign half of my sample to save via salary sacrifice, I estimate a price elasticity of -0.12.

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

	(1) OLS	(2) IV	(3) IV	(4) IV	(5) IV	(6) IV
Log price of retirement saving	0.140***	-0.041**	-0.023	-0.029*	-0.041*	-0.032*
	(0.013)	(0.017)	(0.016)	(0.016)	(0.023)	(0.018)
Log post-tax earnings	0.125***	0.057***	0.026***	0.012	0.016**	0.013
	(0.009)	(0.010)	(0.009)	(0.009)	(0.008)	(0.009)
Observations	226741	226741	208923	208923	208923	208923
R^2	0.704	0.000	-0.000	0.004	0.004	0.005
Controls	No	No	No	Yes	Yes	Yes
Employee FE	Yes	Yes	No	Yes	Yes	Yes
Employee-employer FE	No	No	Yes	No	No	No
Price elasticity	0.535	-0.157	-0.085	-0.108	-0.154	-0.119
Income elasticity	0.478	0.219	0.097	0.044	0.059	0.047

Notes: This table shows the estimated extensive-margin price and income elasticity of employee contributions to DC retirement accounts. The cells show the β and γ coefficients from estimating equation 5.2 on my 2005-12 sample. The price and income elasticities are then obtained by dividing the estimated coefficients by the proportion of employees with strictly positive employee DC contributions. All columns contain year fixed effects. Columns (1) and (2) include employee fixed effects, while columns (3) to (6) include employee-employer fixed effects. Columns (4) to (6) also include controls for the square of age and the employee's job tenure. Column (1) is estimated using OLS, while columns (2) to (6) instrument $\log(p_{it})$ with $\log(p^f_{it})$. Column (4) assumes no employee saves using salary sacrifice, column (5) assumes all employees save using salary sacrifice, and column (6) assumes a randomly-selected 50% of employees save using salary sacrifice. Robust standard errors clustered at the employee level are in parentheses. *, **, and *** denote significance at 0.10, 0.05 and 0.01 levels, respectively.

To summarise, in the preferred specifications (column (6) in both Tables 5.1 and 5.2), I estimate an intensive-margin elasticity of DC retirement contributions to the upfront tax price of retirement saving of -0.21, and a corresponding extensive-margin elasticity of -0.12 for the time period 2005 to 2012. Aggregating, I estimate a total elasticity of around -0.33.

These elasticities are small. At the higher rate threshold, where the tax price of retirement saving falls by 25% after 2008 (from 0.8 to 0.6) for those not saving via salary sacrifice, these estimated elasticities imply an increase in participation of 3%, and an increase in contributions of 5.2%. Given that the share of employees contributing into a DC retirement account is approximately 30% around and above the higher rate threshold in this period, the estimated elasticity implies that participation is only around 1 percentage point higher above the threshold due to the change in the price of retirement saving. This extensive margin effect is therefore significantly smaller than those estimated in Duflo et al. (2006) and Engelhardt and Kumar (2007) in the US, who both find that a 25 percentage point increase in the match rate increases 401(k) participation by 5 percentage points. The intensive margin effect is also economically small: taking an employee earning £60,000 a year and contributing £3,000 into their DC retirement account (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 £150 more into their DC retirement account each year due to the change in the tax price at the higher rate threshold.

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

I now move on to estimating the responsiveness of employees' retirement 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 participation in DC retirement accounts among private-sector employees, with many contributing the default rates under automatic enrolment (Cribb and Emmerson, 2020). In this subsection, I explore whether these changes, in particular the large increase in the number of savers enrolled in workplace pensions, have an effect on the responsiveness to tax incentives. Throughout this section, when including controls in the regression, I also include a dummy for whether the employee's firm has passed the date at which they must automatically enrol their employees (this date is a function of the firm size; see Cribb and Emmerson (2020) for more details).

Table 5.3 presents my preferred IV estimates of the intensive- and extensive-margin results from the 2013-19 sample period, including all controls and employee-employer fixed effects. Note that I now calculate the log price of retirement saving accounting for whether the actually employee uses salary sacrifice, as I can observe this in my data from 2013 on. Tables C.1 and C.2 in the appendix show the full set of results akin to Tables 5.1 and 5.2. The estimated intensive-margin price elasticity for 2013-19 is actually slightly positive, at 0.010, with a 95% confidence interval of (-0.10, 0.12). The estimated coefficient on $\ln p_{it}$ in the extensive-margin regression is negative but statistically insignificantly different from zero at conventional significance levels. It corresponds to an extensive-margin elasticity of just -0.08.

Overall, compared to the 2005-12 period, the estimated elasticities are much smaller in magnitude and now insignificantly different from zero. This suggests that automatic enrolment might have led to a significant fall in the tax responsiveness in retirement saving among UK employees. I explore the mechanisms behind this result in Section 5.6.

Table 5.3: The responsiveness of DC retirement saving to tax incentives, 2013-19

	(1) Intensive	(2) Extensive
Log price of retirement saving	0.003	-0.037
	(0.051)	(0.026)
Log post-tax earnings	0.633***	-0.007
	(0.027)	(0.015)
Observations	94249	186905
R^2	0.864	0.732
Price elasticity	0.003	-0.067
Income elasticity	0.633	-0.013

Notes: This table shows the estimated intensive- and extensive-margin price and income elasticity of DC retirement saving for the 2013-19 period. The intensive-margin cells show the ε_{INT} and η_{INT} coefficients from estimating equation 5.1 on my 2013-19 sample with strictly positive employee DC contributions, while the extensive-margin cells show the β and γ coefficients from estimating equation 5.2 on the full 2013-19 sample. The extensive-margin price and income elasticities are then obtained by dividing the estimated β and γ coefficients by the proportion of employees with strictly positive employee DC contributions. All regressions contain year fixed effects, employee-employer fixed effects and controls for the square of age, the employee's job tenure, and whether the employee's firm has passed the deadline for automatically enrolling employees. In both regressions $\log(p_{it})$ is instrumented with $\log(p_{it}^f)$. Robust standard errors clustered at the employee level are in parentheses. *, **, and *** denote significance at 0.10, 0.05 and 0.01 levels, respectively.

5.5 Robustness

In Table C.3, I show the effects of including DB pension saving in the regression analysis. Some employers facilitate membership in DB pensions rather than DC retirement accounts: employees working for these employers would not be able to join a workplace DC retirement account. The change at the higher rate threshold in the extensive-margin tax incentive to save in a DB pension is identical to the change in the tax incentive to save in a DC retirement account. To determine whether including saving in DB pensions affects the estimated extensive-margin elasticity, I redefine the outcome variable as whether an employee has a positive employee contribution to either a DC retirement account or a DB pension and then repeat the specifications in column (6) of Table 5.2 and in column

(2) of 5.3. The estimated elasticity is slightly lower than when estimated using only DC contributions, indicating that my preferred extensive-margin results may well be an upper bound on the elasticity. This table also shows that the intensive-margin elasticity also falls when including contributions to DB pensions in the analysis, consistent with the fact that DB pension savers have little autonomy over their employee contributions.

In Table C.4, I show how employer contributions to DC retirement accounts respond to the change in the tax price at the higher rate threshold in both periods. This is because employer contributions could also change at the higher rate threshold, and employees could change their own contributions in response to this. For example, it is possible that a given employer might change their compensation package in response to a change in the employee's marginal tax rate (e.g. they might increase their contribution to the employee's retirement account once the employee's earnings reach the higher rate threshold, rather than increasing their take-home pay). If employee contributions then responded to this change in employer contributions, the regression results earlier in this section would not be solely capturing how tax incentives directly affect employee DC contributions. Of course, employer contributions can also respond to employee contributions if employers have a matching arrangement. Given I find a small responsiveness of employee contributions to the price of retirement saving, it is therefore likely that there will be some responsiveness of employer contributions too; a large responsiveness would, however, be more indicative of employer contributions driving the response. In the earlier part of my sample period, the estimated price elasticities are statistically significant; however, they are smaller in magnitude than the corresponding elasticities for employee contributions. For the latter part of my sample period, the estimated price elasticities of employer contributions are small and statistically insignificantly different from zero. Overall, this suggests that changes in employer contributions are not driving changes in employee contributions at the higher rate threshold.

In Tables C.5 and C.6, I test how the estimates change when adjusting the restriction that employees in the sample have real annual gross earnings between £30,000 and £70,000. Changing this sample restriction will indicate how sensitive my results are to the functional form with which I control for post-tax earnings. I rerun the specifications from columns (6) of Tables 5.1, 5.2 and from columns (1) and (2) of Table 5.3 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. Including all employees with gross annual earnings between £20,000 and £80,000 has a larger effect on the estimated elasticities, but generally makes them more positive.

Finally, to assess whether there are longer-run responses of retirement saving to crossing the higher rate threshold, in Table C.7 I include three lags of the log price of retirement saving, and of log post-tax earnings as regressors. In both periods, on both the intensive and extensive margin, the total estimated price elasticities (obtained by summing the lagged coefficients) are no more negative than the contemporaneous price elasticities estimated previously in this section without including lags. These results therefore suggest that there are also not large delayed effects of retirement saving to crossing the higher rate threshold.

5.6 Why does the elasticity fall after the introduction of automatic enrolment?

There are two possible explanations for the fall in the intensive-margin elasticity between the 2005-12 period and the 2013-19 period. The first is that many of the employees that automatic enrolment brought into retirement saving in the latter period are passive savers, who do not respond to tax incentives. Second, it is also possible that automatic enrolment itself had an effect on the responsiveness of existing savers to tax incentives. This may be the case if the new default contribution rates acted as a focal point, or an implicit recommendation, for some of these savers, and they chose to set their contributions at this level rather than responding to the tax incentive.

To test between these explanations, I run regressions across the whole period separately for a sample of "existing" and a sample of "new" savers in Table 5.4. To construct these groups, I first drop from the 2005-2019 sample any individuals who I do not observe saving in a DC retirement account with a positive employee contribution for at least two years after they become eligible for automatic enrolment. Then, existing savers are defined as those employees who I also observe saving in a DC retirement account for at least two years before 2012. New savers are defined as those employees who I observe in at least two years before they become eligible for automatic enrolment, and who do not have positive employee contributions in either DC or DB accounts in any of those years.

For the sample of new savers, I estimate Equation 5.1, while for the sample of existing savers, I also interact whether they are eligible for automatic enrolment with the log price of retirement saving and log post-tax earnings.

Table 5.4: Comparing the intensive-margin responsiveness of old and new savers

	(1) Existing savers	(2) New savers
Log price of retirement saving * AE-eligible	-0.232***	0.045
	(0.083)	(0.108)
Log price of retirement saving * not AE-eligible	-0.178***	
	(0.060)	
Log post-tax earnings * AE-eligible	0.742***	0.643***
	(0.042)	(0.061)
Log post-tax earnings * not AE-eligible	0.589***	
	(0.029)	
Observations	49902	15363
R^2	0.789	0.851

Notes: This table shows how the estimated intensive-margin elasticity of DC employee contributions varies before and after automatic enrolment for existing savers, and how this compares with the estimated elasticity after automatic enrolment for new savers. The first regression includes all observations of 'existing savers', who are employees I observe saving in a workplace DC retirement account for two years before 2012, and also saving in a workplace DC retirement account in two years once they are automatically enrolled. I then estimate equation 5.1, but interacting $\ln p_{it}$ and $\ln y_{it}$ with an indicator variable for whether they are eligible for automatic enrolment, and display the estimated coefficients on these interactions in the table. The second regression includes all observations of 'new savers', who are employees who are not saving in any years before automatic enrolment (minimum 2 years), but are saving in at least two years once they have been automatically enrolled. I then estimate equation 5.1 and display the estimated coefficients ε_{INT} and η_{INT}. Both regressions contain year fixed effects, employee-employer fixed effects and controls for the square of age, and the employee's job tenure. The regression in column (1) also controls for whether the employee's firm has passed the deadline for automatically enrolling employees. In both regressions $\log(p_{it})$ is instrumented with $\log(p^f_{it})$. Robust standard errors clustered at the employee level are in parentheses. *, **, and *** denote significance at 0.10, 0.05 and 0.01 levels, respectively.

The results demonstrate that the lower intensive-margin price elasticity in the 2013-19 period is entirely driven by the influx of passive savers into retirement saving in the latter period. The existing savers have very similar intensive-margin elasticities whether they are eligible for automatic enrolment or not, at around -0.2, while the new savers' elasticity is statistically insignificantly different from zero and in fact slightly positive.

6 Conclusion

Financial incentives and nudges are two ways in which governments around the world incentivise individuals to save for retirement. Existing evidence has shown that nudges, in particular automatic enrolment into retirement accounts, generally do lead to a significant increase in saving (Madrian and Shea, 2001; Cribb and Emmerson, 2020), while the effectiveness of financial incentives seems to vary much more depending on the context (Friedman, 2017).

This paper is the first to provide evidence on how the responsiveness of saving to financial incentives differs in a world with and without nudges. Before the introduction of automatic enrolment, I find a small responsiveness of UK employees' retirement 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 retirement saving. I precisely estimate an intensive margin tax price elasticity of -0.21 and an extensive margin elasticity of -0.12 for the period 2005 to 2012. These elasticities are smaller than in other papers in the literature, and imply that if employees earning £60,000 per year received 20% upfront income tax relief on their retirement saving (rather than 40%), then participation would only fall by 1 percentage point, from 30% to 29%, and contributions (conditional on participation) would only fall by £150, from £3,000 to £2,850.

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. I show that existing savers' responsiveness did not fall with the introduction of automatic enrolment, despite evidence from other contexts that many employees perceive the automatic enrolment default rate as investment advice (Madrian and Shea, 2001). Rather, the lower elasticity in the 2013-19 period is because the new savers brought into retirement saving by automatic enrolment are "passive" savers, who are less likely to respond to tax incentives than existing savers (who, it should be emphasised, are already quite passive).

These results imply that this tax incentive, which is an implicit part of the UK tax system, is not particularly effective at encouraging retirement saving overall. However, it is even less effective at encouraging retirement saving among individuals brought into saving by automatic enrolment, which might be precisely the group that policymakers worry are undersaving. Interpreting these results through the framework of Choukhmane and Palmer (2023) would imply that this tax incentive likely does not have a large positive effect on welfare. On the other hand, Mirrlees et al. (2011) highlight that a benefit of the EET nature of retirement savings accounts in the UK is that it allows individuals to choose between TEE and EET savings vehicles and therefore potentially ‘smooth out’ their tax rate over different periods’ consumption. The results in this paper suggests that taxpayers around the higher rate threshold do not make much use of this tax smoothing advantage of retirement savings accounts.

A Appendix: Further details on dataset construction

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 on employees' earnings, their employee and employer contributions to DC retirement accounts, an employer ID, and limited information on employee characteristics such as age and gender. I calculate the (first- and last-) pound price of retirement saving based on information on whether the employee uses salary sacrifice as well as their earnings. 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 approximately the middle of 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 DC 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 approximate annual earnings well. The reason why I do not use the annual earnings data in the main analysis is that there is no data on DC contributions for the same time frame.

B Appendix: Further details on institutional setting

B.1 Net pay vs. relief-at-source arrangements for DC retirement accounts

In the UK there are two arrangements through which tax relief on contributions to retirement accounts can be administered, which will have an effect on the measurement of these contributions. In 'net pay' plans, contributions to retirement accounts are deducted before tax is calculated on the employee's pay. Conversely, in 'relief at source' plans, the contribution is deducted after tax is calculated, and HMRC (the UK tax authority) then sends an additional 25% contribution to the retirement account 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 DC contributions.

I cannot observe whether a plan is 'net pay' or 'relief at source'. For 'relief at source' plans, I observe contributions to DC retirement accounts before the additional top-up from HMRC, and so underestimate the total amount entering the employee's DC retirement account by 20%. However, this underestimation will be captured in the employer-employee fixed effects I include in the main specification. This is because the intensive margin outcome is *log* employee DC contributions. In other words, if measured retirement contributions at an employer offering a relief at source plan are $\hat{s}_{ijt} = 0.8 \times s_{ijt}$, my outcome variable will be $\ln \hat{s}_{ijt} = \ln(0.8 \times s_{ijt}) = \ln 0.8 + \ln s_{ijt}$, showing that the mismeasurement is a constant that will be absorbed by the employee-employer fixed effect. I will also mismeasure the taxable income of relief-at-source employees, and therefore potentially their retirement 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 retirement saving i.e. before deducting any contributions to DC retirement accounts.

B.2 The effect of the benefit system on the tax price

Parts of the benefit system can also affect the upfront price of retirement saving. Eligibility for most means-tested benefits is assessed against a measure of income after deducting contributions to DC retirement accounts. This means that the full tax price of retirement saving is also affected by whether or not individuals would gain extra entitlement to benefits as a result of their saving. Since I am exploiting a non-linearity in the income tax schedule that occurs around the 90th percentile of the population 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.[7] 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. These people would then not only save £0.40 of income tax by contributing one extra pound to their workplace retirement account, but they also get back £0.175 of child benefit. The tax price of retirement saving therefore falls from £0.80 below the higher threshold to £0.425 above the higher rate threshold for these people. 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.

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

C Appendix tables and figures

Figure C.1: Participation in (DB and DC) retirement accounts around the higher rate threshold, 2005-12 vs. 2013-19

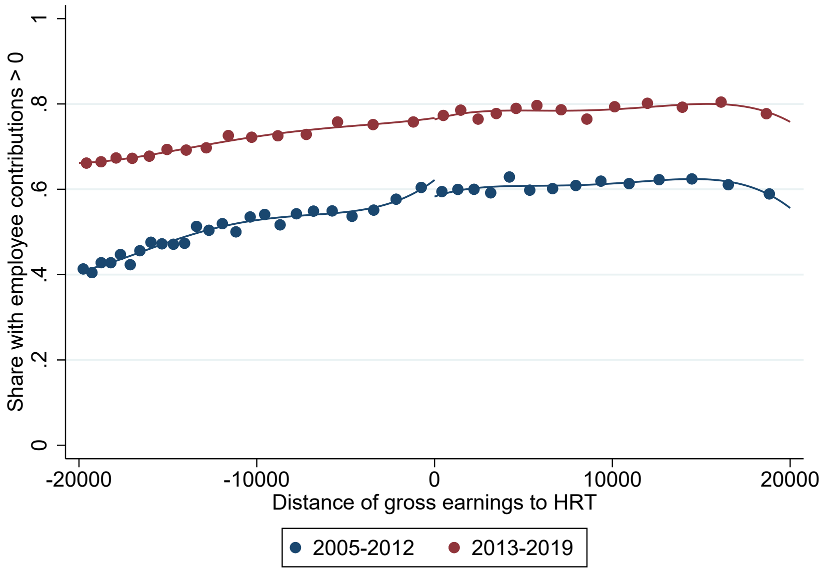

Notes: This figure shows the proportion of employees with strictly positive employee contributions to either a DB or DC retirement account around the higher rate threshold. The higher rate threshold is the level of taxable income at which the marginal tax rate increases from the 'basic' rate to the 'higher' rate. Dots show these proportions by bins of gross earnings (normalised relative to the HRT), separately for two sample periods. Bins are quantile-spaced and selected to minimise mean-squared error as outlined in Cattaneo et al. (2020). Lines show an estimated fourth-order polynomial fit to the data separately either side of the HRT for both periods.

Figure C.2: Participation in DC retirement accounts around the higher rate threshold, 2005-12 vs. 2013-19, for the sample of employees whose (first-pound) marginal tax rate has remained unchanged for at least a year

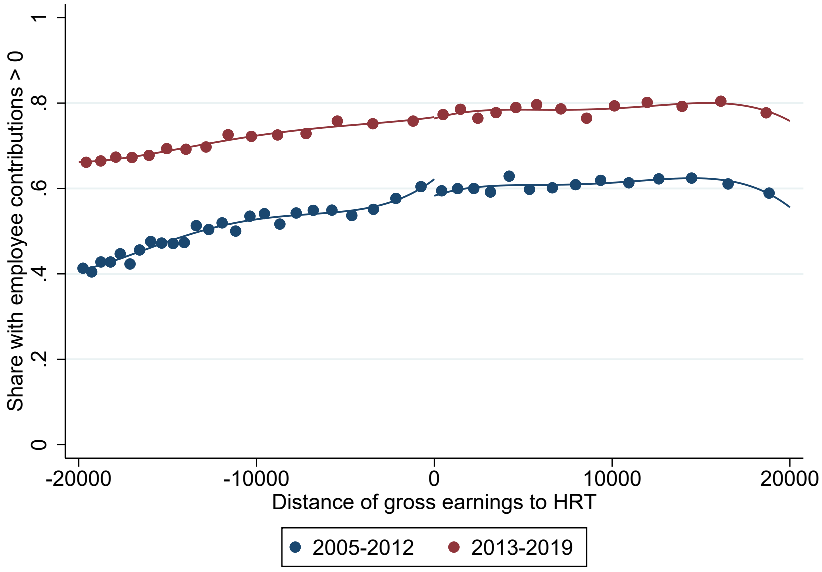

Notes: This figure shows the proportion of employees with strictly positive employee contributions to a DC retirement account around the higher rate threshold. The higher rate threshold is the level of taxable income at which the marginal tax rate increases from the 'basic' rate to the 'higher' rate. I drop from the sample any individuals whose gross earnings are over (under) the HRT in the current year but under (over) the HRT in the previous year. Dots show these proportions by bins of gross earnings (normalised relative to the HRT), separately for two sample periods. Bins are quantile-spaced and selected to minimise mean-squared error as outlined in Cattaneo et al. (2020). Lines show an estimated fourth-order polynomial fit to the data separately either side of the HRT for both periods.

Figure C.3: Bunching of earnings around the higher rate threshold among persistent higher-rate taxpayers, 2005-12

Notes: This figure shows the extent of bunching of gross earnings (in red) and gross earnings after subtracting employee contributions to DC retirement accounts (in blue) around the higher rate threshold. The higher rate threshold is the level of taxable income at which the marginal tax rate increases from the 'basic' rate to the 'higher' rate. The lines plot the frequency density of each variable in bins with width of £500. The sample is the 2005-12 analysis sample restricted to employees with gross earnings (after subtracting DC contributions) of at most £10,500 away from the level of the higher rate threshold, and with strictly positive DC employee contributions. I drop from the sample any individuals whose gross earnings are less than the higher rate threshold in either year t or $t-1$.

Figure C.4: Bunching of earnings around the higher rate threshold among persistent higher-rate taxpayers, 2013-19

Notes: This figure shows the extent of bunching of gross earnings (in red) and gross earnings after subtracting employee contributions to DC retirement accounts (in blue) around the higher rate threshold. The higher rate threshold is the level of taxable income at which the marginal tax rate increases from the 'basic' rate to the 'higher' rate. The lines plot the frequency density of each variable in bins with width of £500. The sample is the 2013-19 analysis sample restricted to employees with gross earnings (after subtracting DC contributions) of at most £10,500 away from the level of the higher rate threshold, and with strictly positive DC employee contributions. I drop from the sample any individuals whose gross earnings are less than the higher rate threshold in either year t or $t-1$.

Table C.1: The intensive-margin responsiveness of DC retirement saving to tax incentives, 2013-19

	(1) OLS	(2) IV	(3) IV	(4) IV	(5) IV	(6) IV
Log price of retirement saving	0.454***	0.009	0.011	0.008	0.013	0.010
	(0.027)	(0.047)	(0.052)	(0.052)	(0.069)	(0.054)
Log post-tax earnings	1.045***	0.823***	0.701***	0.636***	0.635***	0.635***
	(0.024)	(0.032)	(0.027)	(0.027)	(0.028)	(0029)
Observations	101588	101588	93872	93872	93872	93872
Controls	No	No	No	Yes	Yes	Yes
Employee FE	Yes	Yes	No	No	No	No
Employee-employer FE	No	No	Yes	Yes	Yes	Yes

Notes: This table shows the estimated intensive-margin price and income elasticity of employee contributions to DC retirement accounts. The cells show the ε_{INT} and η_{INT} coefficients from estimating equation 5.1 on my 2013-19 sample with strictly positive employee DC contributions. All columns contain year fixed effects. Columns (1) and (2) include employee fixed effects, while columns (3) to (6) include employee-employer fixed effects. Columns (4) to (6) also include controls for the square of age and the employee's job tenure. Column (1) is estimated using OLS, while columns (2) to (6) instrument $\log(p_{it})$ with $\log(p^f_{it})$. Column (4) assumes no employee saves using salary sacrifice, column (5) assumes all employees save using salary sacrifice, and column (6) uses actual information on whether employees save using salary sacrifice. Robust standard errors clustered at the employee level are in parentheses. *, **, and *** denote significance at 0.10, 0.05 and 0.01 levels, respectively.

Table C.2: The extensive-margin responsiveness of DC retirement saving to tax incentives, 2013-19

	(1) OLS	(2) IV	(3) IV	(4) IV	(5) IV	(6) IV
Log price of retirement saving	0.130***	-0.024	-0.038	-0.037	-0.058	-0.042
	(0.015)	(0.023)	(0.023)	(0.023)	(0.037)	(0.027)
Log post-tax earnings	0.116***	0.040***	0.011	-0.011	-0.009	-0.010
	(0.013)	(0.016)	(0.016)	(0.016)	(0.015)	(0.015)
Observations	201108	201108	186029	186029	186029	186029
Price elasticity	0.239	-0.043	-0.069	-0.067	-0.107	-0.078
Income elasticity	0.214	0.073	0.020	-0.021	-0.017	-0.018

This table shows the estimated extensive-margin price and income elasticity of employee contributions to DC retirement accounts. The cells show the β and γ coefficients from estimating equation 5.2 on my 2013-19 sample. The price and income elasticities are then obtained by dividing the estimated coefficients by the proportion of employees with strictly positive employee DC contributions. All columns contain year fixed effects. Columns (1) and (2) include employee fixed effects, while columns (3) to (6) include employee-employer fixed effects. Columns (4) to (6) also include controls for the square of age and the employee's job tenure. Column (1) is estimated using OLS, while columns (2) to (6) instrument $\log(p_{it})$ with $\log(p^f_{it})$. Column (4) assumes no employee saves using salary sacrifice, column (5) assumes all employees save using salary sacrifice, and column (6) uses actual information on whether employees save using salary sacrifice. Robust standard errors clustered at the employee level are in parentheses. *, **, and *** denote significance at 0.10, 0.05 and 0.01 levels, respectively.

Table C.3: The responsiveness of employee contributions to DB and DC retirement accounts to tax incentives

	2005-12		2013-19	
	(1)	(2)	(3)	(4)
	Intensive	Extensive	Intensive	Extensive
Log price of retirement saving	-0.126***	-0.025	0.017	0.016
	(0.036)	(0.017)	(0.044)	(0.023)
Log post-tax earnings	0.462***	0.034***	0.640***	0.038***
	(0.014)	(0.008)	(0.023)	(0.013)
Observations	112509	208923	138472	186905
R^2	0.856	0.839	0.879	0.700
Price elasticity	-0.126	-0.081	0.017	0.028
Income elasticity	0.462	0.109	0.640	0.065

Notes: This table shows the estimated intensive- and extensive-margin price and income elasticity of employee contributions to both DB and DC retirement accounts, separately for the 2005-12 and the 2013-19 period. The intensive-margin cells show the ε_{INT} and η_{INT} coefficients from estimating equation 5.1 on the samples with strictly positive DC or DB contributions, while the extensive-margin cells show the β and γ coefficients from estimating equation 5.2 on the full samples, where s_{it} instead denotes contributions to DB or DC retirement accounts. All regressions contain year fixed effects, employee-employer fixed effects and controls for the square of age and the employee's job tenure. The 2013-19 regressions also control for whether the employee's firm has passed the deadline for automatically enrolling employees. In all regressions $\log(p_{it})$ is instrumented with $\log(p^f_{it})$. Robust standard errors clustered at the employee level are in parentheses. *, **, and *** denote significance at 0.10, 0.05 and 0.01 levels, respectively.

Table C.4: The responsiveness of employer contributions to tax incentives

	2005-12		2013-19	
	(1)	(2)	(3)	(4)
	Intensive	Extensive	Intensive	Extensive
Log price of retirement saving	-0.111**	-0.035**	0.047	-0.008
	(0.044)	(0.018)	(0.043)	(0.026)
Log post-tax earnings	0.515***	0.017**	0.652***	-0.005
	(0.022)	(0.009)	(0.024)	(0.015)
Observations	60871	208852	101153	186886
R^2	0.892	0.797	0.909	0.733
Price elasticity	-0.111	-0.111	0.047	-0.013
Income elasticity	0.515	0.054	0.652	-0.009

Notes: This table shows the estimated intensive- and extensive-margin price and income elasticity of DC employer contributions separately for the 2005-12 and the 2013-19 period. The intensive-margin cells show the ε_{INT} and η_{INT} coefficients from estimating equation 5.1 on the samples with strictly positive employer DC contributions, while the extensive-margin cells show the β and γ coefficients from estimating equation 5.2 on the full samples, where s_{it} instead denotes DC employer contributions. All regressions contain year fixed effects, employee-employer fixed effects and controls for the square of age and the employee's job tenure. The 2013-19 regressions also control for whether the employee's firm has passed the deadline for automatically enrolling employees. In all regressions $\log(p_{it})$ is instrumented with $\log(p^f_{it})$. Robust standard errors clustered at the employee level are in parentheses. *, **, and *** denote significance at 0.10, 0.05 and 0.01 levels, respectively.

Table C.5: Testing the sensitivity of the intensive margin estimates to the sample earnings range

		2005-12			2013-19	
	(1)	(2)	(3)	(4)	(5)	(6)
	30-70K	35-65K	25K-75K	30-70K	35-65K	25K-75K
Log price of retirement saving	-0.212***	-0.231***	-0.214***	0.003	-0.025	0.090**
	(0.055)	(0.069)	(0.050)	(0.051)	(0.064)	(0.045)
Log post-tax earnings	0.500***	0.484***	0.504***	0.633***	0.609***	0.685***
	(0.023)	(0.037)	(0.018)	(0.027)	(0.045)	(0.019)
Observations	50725	34981	67326	94249	62042	134918
R^2	0.866	0.849	0.879	0.864	0.859	0.874

Notes: This table shows how the estimated intensive-margin price and income elasticity of DC employee contributions for the 2005-12 and 2013-19 period change when altering the sample restriction on gross earnings. The cells show the ε_{INT} and η_{INT} coefficients from estimating equation 5.1 on observations with strictly positive employee DC contributions. In each column, the sample for the regression is restricted to observations with real gross earnings in the range specified in the column title. All regressions contain year fixed effects, employee-employer fixed effects and controls for the square of age, and the employee's job tenure. The 2013-19 regressions also control for whether the employee's firm has passed the deadline for automatically enrolling employees. In all regressions $\log(p_{it})$ is instrumented with $\log(p^f_{it})$. Robust standard errors clustered at the employee level are in parentheses. *, **, and *** denote significance at 0.10, 0.05 and 0.01 levels, respectively.

Table C.6: Testing the sensitivity of extensive margin estimates to the sample earnings range

	2005-12			2013-19		
	(1)	(2)	(3)	(4)	(5)	(6)
	30-70K	35-65K	25K-75K	30-70K	35-65K	25K-75K
Log price of retirement saving	-0.032*	-0.038*	-0.018	-0.037	-0.064**	0.011
	(0.018)	(0.021)	(0.016)	(0.026)	(0.032)	(0.022)
Log post-tax earnings	0.013	0.005	0.028***	-0.007	-0.028	0.016
	(0.009)	(0.013)	(0.006)	(0.015)	(0.024)	(0.010)
Observations	208923	140986	288717	186905	123529	266213
R^2	0.774	0.775	0.775	0.732	0.746	0.721
Price elasticity	-0.119	-0.140	-0.073	-0.067	-0.118	0.019
Income elasticity	0.047	0.018	0.110	-0.013	-0.051	0.029

Notes: This table shows how the estimated extensive-margin price and income elasticity of DC employee contributions for the 2005-12 and 2013-19 period change when altering the sample restriction on gross earnings. The cells show the β and γ coefficients from estimating equation 5.2. The price and income elasticities are then obtained by dividing the estimated coefficients by the proportion of employees in the sample with strictly positive employee DC contributions. In each column, the sample for the regression is restricted to observations with real gross earnings in the range specified in the column title. All regressions contain year fixed effects, employee-employer fixed effects and controls for the square of age, and the employee's job tenure. The 2013-19 regressions also control for whether the employee's firm has passed the deadline for automatically enrolling employees. In both regressions $\log(p_{it})$ is instrumented with $\log(p_{it}^f)$. Robust standard errors clustered at the employee level are in parentheses. *, **, and *** denote significance at 0.10, 0.05 and 0.01 levels, respectively.

Table C.7: Checking for lagged responses to changes in the price of retirement saving

	2005-12		2013-19	
	(1)	(2)	(3)	(4)
	Intensive	Extensive	Intensive	Extensive
Log price of retirement saving	-0.169*	-0.017	0.088	-0.028
	(0.099)	(0.037)	(0.070)	(0.037)
L.Log price of retirement saving	0.026	-0.019	0.103*	0.086***
	(0.087)	(0.034)	(0.055)	(0.031)
L2.Log price of retirement saving	-0.131	0.040	0.104**	0.058**
	(0.094)	(0.035)	(0.050)	(0.029)
L3.Log price of retirement saving	0.140*	-0.010	0.017	0.135***
	(0.078)	(0.033)	(0.048)	(0.028)
Log post-tax earnings	0.429***	0.012	0.594***	-0.033
	(0.041)	(0.018)	(0.037)	(0.021)
L.Log post-tax earnings	0.089***	-0.003	0.136***	0.033***
	(0.022)	(0.011)	(0.021)	(0.012)
L2.Log post-tax earnings	0.029	0.005	0.108***	0.007
	(0.019)	(0.010)	(0.019)	(0.010)
L3.Log post-tax earnings	0.073***	0.007	0.057***	0.022**
	(0.018)	(0.009)	(0.016)	(0.009)
Observations	15244	55305	49329	97405
R^2	0.906	0.836	0.882	0.770

Notes: This table tests for lagged effects of the price of retirement saving and post-tax earnings on DC employee contributions. For the intensive-margin cells, I estimate 5.1, but including three lags for all right-hand-side variables in the regression, and display the estimated coefficients on $\ln pit$, $\ln y_{it}$, and the corresponding lags. For the extensive-margin cells, I estimate 5.2, but including three lags for all right-hand-side variables in the regression, and display the estimated coefficients on $\ln pit$, $\ln y_{it}$, and the corresponding lags. All regressions contain year fixed effects, employee-employer fixed effects and controls for the square of age and the employee's job tenure. The 2013-19 regressions also control for whether the employee's firm has passed the deadline for automatically enrolling employees. In all regressions $\log(p_{it})$ and lags thereof are instrumented with $\log(p_{it}^f)$ and lags thereof. Robust standard errors clustered at the employee level are in parentheses. *, **, and *** denote significance at 0.10, 0.05 and 0.01 levels, respectively.

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.

Attanasio, O., J. Banks, and M. Wakefield (2004). Effectiveness of tax incentives to boost (retirement) saving: theoretical motivation and empirical evidence. *Oecd Economic Studies*.

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

Bernheim, B. D., A. Fradkin, and I. Popov (2015). The welfare economics of default options in 401(k) plans. *The American Economic Review 105*(9), 2798–2837.

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.

Choi, J. J., E. Haisley, J. Kurkoski, and C. Massey (2017). Small cues change savings choices. *Journal of Economic Behavior & Organization 142*, 378–395.

Choukhmane, T. and C. Palmer (2023). How do consumers finance increased retirement savings? Working paper.

Christensen, C. S. and B. E. Ellegaard (2023). Do tax subsidies for retirement saving affect total private saving? New evidence on middle-income workers*. *The Scandinavian Journal of Economics 125*(4), 933–955.

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.

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.

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.

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.

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

OECD (2021). *Pensions at a Glance 2021.* OECD Publishing, Paris.

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.

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