

# Does Statutory Incidence Matter? Earnings Responses to Social Security Contributions

Enda Patrick Hargaden\*  
Barra Roantree†

June 6, 2020

## Abstract

Economic textbooks posit that it is irrelevant which side of a market statute formally levies taxes on. This paper provides evidence that statutory incidence can play an important role in explaining taxpayer behaviour, at least in the case of discontinuous taxes. Using administrative data from Ireland, we show that earnings responses to ‘notches’ in social security contributions are far stronger when statutorily levied on employees compared to employers. This is despite being levied on the same base and remitted by the employer in both cases, suggesting that the irrelevance of statutory incidence may be more limited than previously thought.

*JEL classification:* H22, H26, H55, J38

*Keywords:* statutory incidence, bunching, notches, social security contributions.

---

The authors are grateful for comments and suggestions from Don Bruce, Celeste Carruthers, Georg Schaur, Joel Slemrod, Michael Stimmelmayer, and from seminar participants at the University of Michigan, University of Oxford, National Bureau of Economic Research, Society of Labor Economists, National Tax Association, International Institute of Public Finance, and the Irish Economics Association. This research was generously supported by the University of Tennessee’s Boyd Center for Business and Economics Research.

\*University of Tennessee, 702 Stokely Management Center, Knoxville, TN 37996-0570, USA. Email: enda@utk.edu.

†Economic and Social Research Institute, Whitaker Square, Dublin 2, Ireland. Email: barra.roantree@esri.ie

# 1 Introduction

Classical economic theory as taught to most students predicts that firms’ and consumers’ responses to taxes emerge from the relative sizes of the elasticities of supply and demand. As firms have the ability to pass taxes onto consumers through higher prices, it is irrelevant whether statute formally levies taxes on the consumer or the producer. In this model “statutory incidence is not a useful economic concept” (Salanié, 2011).

Recent work has suggested meaningful departures from this ‘irrelevance of statute’ proposition in both labour and product markets. Chetty et al. (2009) find tax-exclusive price labelling increases retailer profitability and thus that tax salience affects real outcomes. Kopczuk et al. (2016) provide convincing evidence that statutory incidence affects tax collection if some agents are better at evasion than others. Saez et al. (2012) shows that employers compensate workers affected by a cohort-based payroll tax reform for employer but not employee contributions. In a lab experiment, Weber and Schram (2017) find differential extensive-margin labour supply responses to income taxes and dollar-equivalent payroll taxes.

This paper provides a similar point of departure for *notches*, thresholds where tax liabilities increase discontinuously. Crossing these thresholds does not just increase marginal rates, but also triggers substantial infra-marginal/lump sum liabilities. These discrete jumps provide extremely strong incentives to report earnings just below the thresholds. Tax notches have been used extensively to estimate behavioural responses (e.g. Kleven, 2016) but with some notable exceptions (e.g. Blinder and Rosen, 1985; Kleven and Waseem, 2013) less work has been done on their theoretical properties.

The first contribution of this paper is to show that the standard theoretical results on statutory incidence do not hold in the neighbourhood of notches. We formalise the conditions when the irrelevance of statute applies to tax notches, showing that if a notch threshold falls between the gross and net price created by the tax wedge, then liabilities do not arise when the tax is levied on the lower price. In the context of labour markets, there are circumstances where tax liabilities for wages near notch thresholds differ for payroll (employer) taxes

compared to income (employee) taxes.

Secondly, we exploit quasi-experimental variation to evaluate whether earnings respond equally to notches statutorily levied on employees and employers. The social security tax in Ireland (Pay Related Social Insurance, or PRSI) splits contributions into a portion that “shall be payable [by the] employer”, and a portion that will be “to the exclusion of” the employee.<sup>1</sup> In economic terms, this means the tax schedule contains notches which discontinuously increase statutory liabilities on either employees (via decreased pay) or employers (via increased gross costs). For example, crossing a €339 per week earnings threshold in 2007 increased employee taxes by €8.48 per week (€440 per year, 2.5% of income) while crossing a €356 per week threshold increased employer taxes by €8.01 per week (€416 per year, 2.25% of income). Textbook economic theory predicts responsiveness is determined by elasticities, not statutory liability, and so this setup facilitates a test of statutory incidence.

We investigate responses to these notches using administrative linked employer-employee earnings data. We find clear evidence of earnings responding to this incentive, but only when the tax is levied formally on employees. This finding is striking because the tax is levied on the same base and remitted fully by the employer in both cases. Citing differential ability to evade taxes across the supply-chain, Kopczuk et al. (2016) find tax collections for state diesel taxes increased when the physical requirement of remitting the tax shifted from distributors to wholesalers. There is no difference in evasion possibilities between the taxes in our setting. Employers calculate the combined employee and employer contributions due and remit this to the tax authority. The only difference is who the tax is statutorily incident on: the employee, or the employer.

This is precisely the situation where one might expect that the canonical invariance of incidence result should hold. We show it does not. If tax schedules are discontinuous, as they are with notches, then taxes formally levied on employees do not necessarily result in the same equilibrium as dollar-equivalent taxes formally levied on employers.

To investigate mechanisms underpinning this empirical finding, we use ma-

---

<sup>1</sup>Section 13(2)(d) of the Social Welfare Consolidation Act, 2005.

chine learning techniques to decompose the determinants of responsiveness. We analyze if earnings behaviour can be predicted by the characteristics of the employee, of the employer, or both. Our findings again indicate differences between employee versus employer taxes. There are some intuitive results for responsiveness to employee taxes: having self-employment income, or working in the construction sector, for example, are good predictors of reporting a “tax-advantaged” income<sup>2</sup> just below the notch.<sup>3</sup> However, we find that the predictors differ between employee and employer taxes. While a suite of variables are relevant for the employee tax, only a single variable (the form of incorporation) is a robust predictor of responses to the employer tax. Thus earnings are not only less sensitive to employer-focused taxes, they appear to be determined by different factors.

As well as adding to the growing literature establishing significant deviations from the textbook predictions of classical public finance theory, this paper also relates to the literature exploring how employers and employees respond to payroll taxes. Like Lehmann et al. (2013) and Adam et al. (2019), we find that earnings respond more to changes in employee than employer SSCs. This is not to say employer SSCs cannot or do not impact on the labour market. Rather, their effects may instead be felt more at the firm rather than individual worker level, as suggested by Saez et al. (2019) who find that Swedish firms use payroll tax cuts to expand employment but do not increase wages.

In providing evidence that statutory incidence can play an important role in explaining taxpayer behaviour at notches, this paper has potentially important implications for teaching and policy. Firstly, most textbooks for both graduate and undergraduate courses teach the irrelevance of incidence proposition as a key prediction of classical economic theory (e.g. Salanié, 2011). Our results suggest that this proposition may be more limited than previously thought and that results from linear taxation should not ex-ante be assumed generalisable to other settings. Second, as Kleven (2016) notes, notches are “ubiquitous across

---

<sup>2</sup>Tax-advantaged incomes avoid the notch, resulting in a higher *after-tax* income. More detail is available in Section 4.

<sup>3</sup>Working in cash-based sectors like construction increases the opportunities for misreporting, and there is a literature that finds self-employed people are relatively more responsive to tax incentives (e.g. Adam et al., 2017).

a wide range of tax and non-tax settings.” The results here may therefore have important implications for policy areas other than payroll taxation. For example, property sales taxes are in many countries statutorily levied on the purchaser. Whether the economic incidence of these taxes may differ if they were statutorily levied on sellers — as briefly considered by the British government<sup>4</sup> — is therefore an interesting question for future research.

Section 2 outlines the theoretical component of the paper, formally applying statutory incidence analysis to notches. Section 3 discusses the institutional details of PRSI, with Section 4 providing an overview of the administrative dataset used. Section 5 comprises the empirical analysis which find differential magnitudes of responses and mechanisms underpinning them, while Section 6 concludes.

## 2 Theory

Economists have known since at least Cournot (1838) that taxes are partly shifted from the remitter to another person in the transaction. The standard incidence results comprise several inter-related predictions: that the dollar-value of the tax defines the extent of the response, not whether it is administered on a gross cost or net price basis; that the identity of the remitter does not affect equilibrium prices; and that the elasticities of supply and demand allocate relative burdens (Gruber, 2005). This section outlines the extent to which these results apply to notches, i.e. discontinuous increases in tax liabilities. Our empirical analysis applies specifically to labor markets, and so we use that as our example. The theoretical result is applicable to goods and services in general.<sup>5</sup>

Let  $W$  denote the larger, tax-inclusive cost of hiring a worker, let  $w$  denote the smaller after-tax income of a worker, and let  $\eta$  denote an elasticity of supply ( $\eta^S$ ) or demand ( $\eta^D$ ). The theory of statutory irrelevance says equilibrium wages

---

<sup>4</sup>See <https://www.ft.com/content/9b75e93e-c1b0-11e9-a8e9-296ca66511c9>, accessed on 2 June 2020

<sup>5</sup>One difference is the base for income taxes is the product of hourly wage  $w$  and hours  $L$ , whereas commodity taxes are typically defined per unit sold. These are equivalent except in ‘peculiar’ cases of backward-bending labor supply curves, etc.

are unaffected by whether the tax is statutorily placed on the employer or the employee.

**Standard Tax Theory:** If tax liabilities are continuous in price, the incidence of a employee-borne income tax  $t$  is equivalent to that of a dollar-equivalent employer-borne payroll tax  $\tau$ .

When  $t = \tau$ ,<sup>6</sup> an employee-borne income tax  $t$  will see the gross cost of hiring a worker  $W$  increase by  $\left(\frac{\eta^S}{\eta^S - \eta^D}\right) t$ . As  $W - t = w$ , the worker's after tax benefit falls by  $\left(\frac{\eta^D}{\eta^S - \eta^D}\right) t$ . For a employer-borne tax  $\tau$ , the benefits to the worker  $w$  falls by  $\left(\frac{\eta^D}{\eta^S - \eta^D}\right) \tau$ . Adding the tax  $\tau$  onto  $w$ , the cost to the employer increases by  $\left(\frac{\eta^S}{\eta^S - \eta^D}\right) \tau$ , i.e. the same amount as above. This is a textbook result (e.g. Gruber, 2005) and so we refer the reader to e.g. Fullerton and Metcalf (2002) for a fuller derivation.

**Proposition 1:** Statutory incidence can affect equilibrium outcomes in the neighbourhood of notches.

We provide proof by contradiction with a numerical example. Incidence is statutorily on the gross price  $W$ . Prior to the introduction of a tax, there is no difference between the gross price ( $W$ ) and benefit to the worker ( $w$ ), so that  $W = w$  and in this numerical example  $W = w = 105$ . For simplicity, the magnitudes of elasticities of supply and demand are equal (i.e. that  $|\eta^S| = |\eta^D|$ ). A tax of  $t = 20$  must be paid if the gross price  $W$  is strictly greater than the notch threshold  $N = 100$ . Formally:

$$\text{Tax Liability } t = \begin{cases} 0 & \text{if } W \leq 100 \\ 20 & \text{otherwise} \end{cases}$$

Equal elasticities imply  $|\eta^D / (\eta^S - \eta^D)| = \frac{1}{2}$  and so the burden of the tax will be shared evenly. Then  $w = 105 - \frac{1}{2}(t) = 95$  and  $W = w + t = 115$ . This

---

<sup>6</sup>Or dollar-equivalent. This analysis focuses on a unit tax case because the presentation is simpler. For an ad valorem tax, the rate  $\tau$  on a high price needs to scaled down to match a dollar-equivalent rate  $t$  on a lower price, i.e.  $\tau \equiv t/(1 + t)$ . The unit/lump-sum case permits the simpler  $t = \tau$  formulation.

represents a coherent equilibrium where the tax is the difference between  $W$  and  $w$ , and the departures from the original price (of 105) are consistent with the underlying elasticities.

Now consider the almost-equivalent case with the change that the tax  $\tau$  is levied on the net benefit to the worker  $w$  rather than the gross cost  $W$ . Formally,

$$\text{Tax Liability } \tau = \begin{cases} 0 & \text{if } w \leq 100 \\ 20 & \text{otherwise} \end{cases}$$

As we have not changed the baseline prices or elasticities, the ‘irrelevance of statutory incidence’ prediction is a continuation of the previous equilibrium: a wage pair  $\{W, w\}$  equal to  $\{115, 95\}$ .

This does not constitute an equilibrium. As  $w \leq 100$ , no tax is due. With no tax due, there can be no difference between  $W$  and  $w$ . Thus, for the case of notches, whether the tax is statutorily incident on the consumer or the producer affects equilibrium outcomes. ■

We provide the intuition of the result with a graphical analysis in Figure 1 below. It depicts a standard market equilibrium with a tax.

Without a notch, the equilibrium is quantity  $L_1$  with a wedge between gross wages and net wages equal to the amount of the tax. The notch threshold at  $N$  changes the analysis. If the tax is levied on the employer-side ( $W$ ), then the price is above the notch threshold, and the textbook case prevails. If the tax is levied on the employee’s side of things, then the relevant price ( $w$ ) is below the threshold and the tax is not due. For any notch and elasticities, there are wages close enough to the cut-off point that the net wage is below the threshold.

Note the crucial distinction between when the tax increases gross costs versus decreases net pay. In the gross cost case (and in the net pay case where wages are far above the notch), the standard incidence results constitute a coherent equilibrium. Although mutually beneficial bargains may still exist, the standard equilibrium results are a reasonable starting-point for predicting behaviour. However, in the net pay case when the wage is close enough to the notch threshold, the standard equilibrium will not pertain. The distribution of after-tax burdens predicted by standard economic theory simply does not hold,

Figure 1: Graphical depiction of the theoretical result

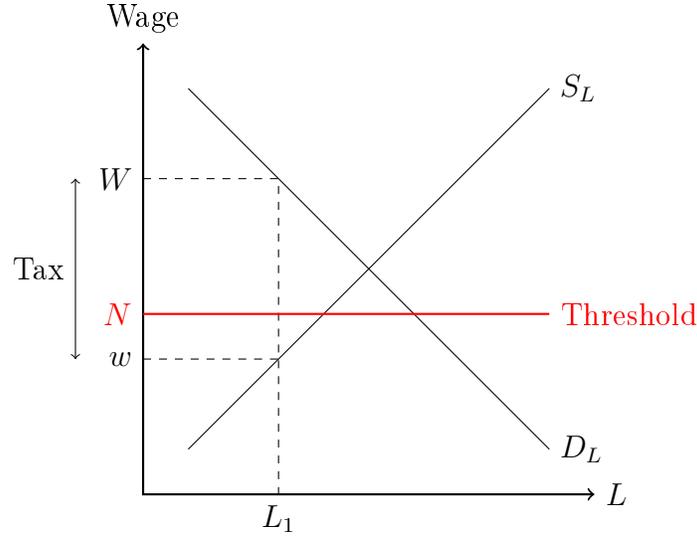


Figure shows the (in)applicability of standard theory to the introduction of a notch. Taxes are due if the transaction price is above  $N$ . Administering the tax system on the lower/net price  $w < N$  means no tax is due. Tax liabilities accrue if the tax is administered on the higher/gross price  $W > N$ .

as no tax is owed. We thus do not expect equivalent adjustments in the net pay (Employee notch) case as the gross cost (Employer notch) case. To the extent that markets tend to adjust to equilibrium conditions, we expect the gross cost case to adjust as standard theory predicts, but expect ‘new’ adjustment behaviour (such as the Pareto-improving bargain  $w = W = N$ )<sup>7</sup> to occur in the net pay case.

For wages far enough above the notch threshold, the standard incidence results continue to hold. It is only in the neighbourhood of the notch threshold where the possibility of divergence occurs. We formalize this in Proposition 2.

**Proposition 2:** For given elasticities of supply  $\eta^S$  and demand  $\eta^D$ , there are points in the income distribution ‘close enough’ to a notch threshold  $N$  where the incidence of a notched income tax  $\tau$  does not equal the incidence of a dollar-equivalent notched payroll tax  $t$ .

<sup>7</sup>Of course, this is just one such agreement from a class of potential agreements.

Proof For a pre-tax equilibrium wage of  $W_0 \in (N, N + \tau)$  one of two possible scenarios must prevail. Noting that  $W_0 - N$  is how far above the notch threshold the initial wage was, it clearly must be the case that either:

$$W_0 - N > \left| \left( \frac{\eta^D}{\eta^S - \eta^D} \right) \tau \right|$$

or

$$W_0 - N \leq \left| \left( \frac{\eta^D}{\eta^S - \eta^D} \right) \tau \right|$$

In the first case, there is a relatively large gap between the pre-tax wage  $W_0$  and the notch threshold  $N$ . This is a sufficient condition for  $w > N$ . That implies that even the relatively low price  $w$ , the wage that accrues to the worker, is above the notch threshold. Then the tax is due, and the regular incidence results hold.

However in the second case, there is a relatively small gap between the pre-tax wage  $W_0$  and the notch threshold  $N$ . We refer to this pre-tax wage  $W_0$  as existing “in the neighbourhood of” or “close enough to” the notch threshold. In this case,  $w \leq N$ . This means the wage accruing to the worker  $w$ , on which the tax is based, falls weakly below the threshold  $N$ . No tax is due.

Indeed, for any given elasticities (and thus any given  $\left| \left( \frac{\eta^D}{\eta^S - \eta^D} \right) \tau \right| = \delta$ ) we can find a  $W_0$  close enough to the threshold such that  $W_0 - N = \epsilon < \delta$ , where the tax is not due. ■

Proposition 2 shows that when a tax is statutorily on the employee, there is always a range of pre-tax wages when the laws of invariance do not hold. In the range depicted in the first scenario, where the initial wage is adequately above  $N$ , the after-tax wage  $w$  remains above the notch threshold and the standard formulae apply.

However in the second scenario, where the initial wage is ‘close enough’ to  $N$ , tax notches can force such a divergence from original prices that the notches are self-defeating in a revenue sense. This is the key difference between discontinuous taxes (notches) and the usual incidence results we find in textbooks. Notches’ discontinuous nature are different.

Theory thus suggests there may be non-equivalence in responsiveness between notches on employees and notches on employers. The empirical section that follows investigates this theoretical possibility. We analyze whether earnings responses differ between the gross cost and net pay cases.

### 3 Institutional Details

Social security in Ireland is funded primarily funded through the Pay Related Social Insurance (PRSI) system. PRSI is a tax with legal obligations on both employees and employers to contribute. Contributions entitle workers to a number of benefits such as increased unemployment insurance. Eligibility for these benefits is based on the duration that the tax is paid, rather than the number of euros paid. Thus, although taxes increase with income, because PRSI is largely an ‘in or out’ system, benefits are essentially independent of income. In this respect PRSI has elements of redistribution between workers rather than an actuarially fair insurance system.

The legislation specifies the shares of total contributions (“statutory incidence”) that are to be borne by the employer and the employee. Section 13 of the Social Welfare Consolidation Act states:

*The employer shall, in relation to any employment contribution, be liable in the first instance to pay both the employer’s contribution comprised therein and also, on behalf of and to the exclusion of the employed contributor, the contribution comprised therein payable by the contributor.*

This legal text asserts that remittance of both employer and employee contributions is the responsibility of the employer, but that the employer may subtract the employee contribution from gross wages. The implication is that the employee contribution is administered as a reduction in the worker’s net pay while the employer contribution, which has a separate schedule of rates, is to be added to the employer’s gross costs.

Like most social insurances taxes the PRSI system is progressive, with marginal rates increasing as income crosses thresholds from one weekly pay band (or

“subclass”) to another. Two unique features of this system are at the heart of this paper, however. Firstly, crossing these thresholds does not just increase marginal rates, but also triggers substantial infra-marginal/lump sum liabilities. These discrete jumps in tax liability — what the literature commonly calls ‘notches’ — provide extremely strong incentives to report earnings just below these thresholds. Following Saez (2010) and Chetty et al. (2011), a large literature has investigated the extent to which agents ‘bunch’ near these thresholds. The amount of bunching, the excess mass of agents reporting incomes just below these thresholds relative to just above, reveals the extent of the responsiveness of taxpayers to the tax (Kleven and Waseem, 2013). Bunching estimators can be used to infer the elasticity of earnings, with the extent of the bunching positively related to the elasticity of earnings.

Figure 2: Accounting effects of notches.

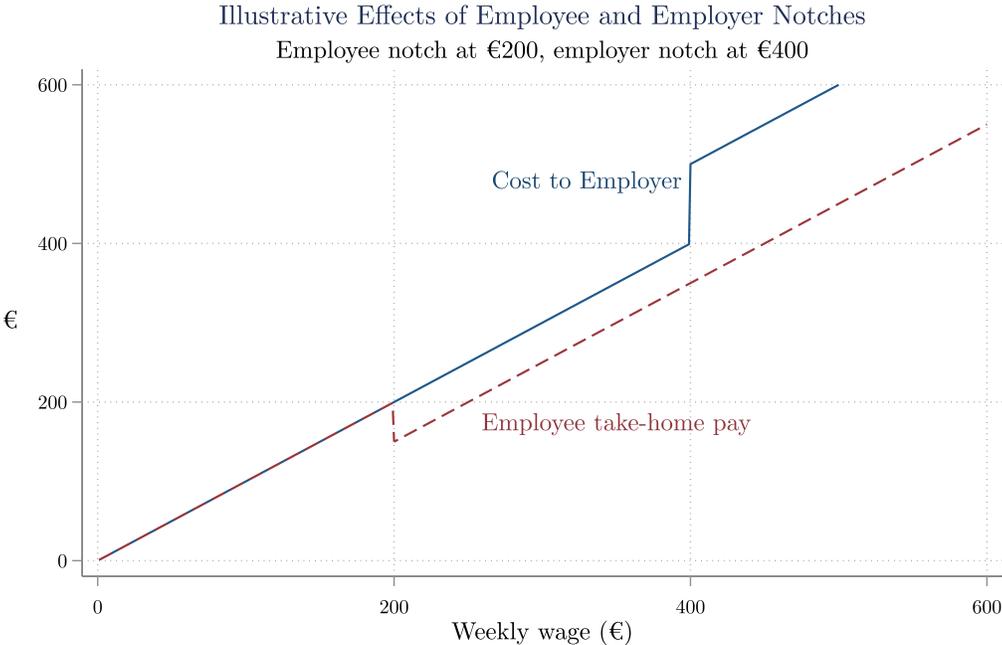


Figure depicts a stylized case of payroll- and income-tax notches. In this example, the notch is €50 for the income (employee) tax and €100 for the payroll (employer) tax.

Secondly, these notches apply differentially to the employee and employer

shares of total PRSI contributions. That is, statute not only induces these notches, but also specifies whether the notch should be administered through an increase in gross costs (employer contribution) or through a decrease in net pay (employee contribution). Figure 2 illustrates a hypothetical example of the effects of this differential treatment. The figure shows the relationship between wage payments made by an employer and the take-home pay of the employee. At an employee notch, which in the figure is located at €200, the take-home pay of the employee drops by the amount of the tax (€50), but the cost to the employer is unaffected. In contrast, the notch at €400 sees the cost to the employer increase by the amount of the tax (here, €100) but does not change the net pay of the employee. Of course this scenario depicts an accounting exercise, not behavioural responses.

The experimental design of this paper relies on the fact that the notched element of these thresholds apply differentially to the employers' tax and the employees' tax. In particular, this paper investigates if responsiveness differs between employee notches and employer notches. The null hypothesis motivated by standard theory is that statutory incidence does not affect behaviour.

The full list of notches and liabilities associated with them is outlined in Table 1. An example helps clarify the information presented in this table. In 2007, earning €0.01 over €339 per week (€17,628 per year) pushed an employee into "Subclass AX". This increased the PRSI marginal tax rate from 8.5% to 12.5%, but more importantly triggered a €8.48 per week (€440 per year) lump-sum penalty in the employee share. Notice that the tax increase applied to the employee share. The legal requirement of the firm was to implement this tax through a decrease in net pay. In accounting terms, their gross costs were entirely unaffected.<sup>8</sup>

In contrast, earning anything above €356 per week (€18,512 per year) pushed an employee into "Subclass AL". This increased the PRSI marginal tax rate from 12.5% to 14.75%, and triggered an €8.01 per week (€416 per year) lump-sum penalty — but this time on the employer share without any change in social security benefit entitlements.<sup>9</sup> That is, unlike the previous example,

<sup>8</sup>Beyond the notional increase in liability from a 1 cent pay increase, of course.

<sup>9</sup>Any employee with pay exceeding €38 per week over this period was deemed to have

Table 1: Outline of notches and tax penalties for crossing threshold

Year	<i>AX (Employee)</i>		<i>AL (Employer)</i>	
	Threshold	Notch amount	Threshold	Notch amount
2005	287	6.40	356	8.01
2006	300	6.92	356	8.01
2007	339	8.48	356	8.01
2008	352	9.00	356	8.01
2009	352	9.00	356	8.01

Table shows the income thresholds and liabilities (in €) from crossing the thresholds. All amounts relate to weekly frequency. AX and AL are the names of the tax subclasses entered by crossing a threshold.

this notch was statutorily incident on the firm. In legal terms, this tax was to be paid through an increase in gross costs, not a reduction in net pay.

Textbook analysis suggests taxes on employers and employees should be equivalent. The equilibrium will be an adjustment of prices and quantities independent of whether the tax is accounted as an increase in gross cost or as a reduction in net pay. Section 2 demonstrated that notched taxes on the employee and employer need not be equivalent. The invariance of statutory incidence does not hold in the neighbourhood of notches. As the tax is less likely to owed in the net pay case, ‘new’ adjustment behaviour is expected.

## 4 Data Description

The data in this paper are an administrative panel of employee tax returns for Ireland, with access provided under a confidentiality agreement with the Central Statistics Office (CSO). CSO acts as an intermediary for collating relevant data from various state agencies. The primary source are tax returns from the Irish Revenue Commissioners. These data contain the details from the P35 tax form. This is comparable to a W-2 in the United States in that it is

---

made a PRSI contribution entitling them to social security benefits later in life, including unemployment benefit and the (contributory) state pension. Neither eligibility for nor the amount of benefit entitlement changes at or around any of the PRSI thresholds we consider in our analysis.

the firm’s statement of payments made to an employee, and the amount of tax withheld and remitted to the IRS. This is an advantage of the dataset, as it is income reported by a third-party (the employer, who faces additional legal ramifications for mis-reporting) rather than data populated by the employees themselves. Self-employed people are subject to self-assessment. The income figure is formally called “Taxable Pay”. The P35 form includes both the firm’s ID and the individual’s social security number, and with these CSO merges in firms’ form of incorporation and four-digit industrial sector; the year of birth, sex, and nationality of individuals; and firms’ number of employees, number of hires, and number of separations. The data come in the form of a random sample of 10% of all individual tax returns, which are reported by the employer annually. It is a representative sample of the universe of workers.

Table 2: Summary statistics from administrative data sources

Variable	Obs	Mean	Std Dev	Min	Max
Employee Dom. Region	27,326	0.65	0.48	0	1
Employer Dom. Region	28,119	0.76	0.43	0	1
Age	934,171	35.59	12.71	16	85
Irish	934,171	0.67	0.47	0	1
Male	934,171	0.47	0.50	0	1
EU 2004	934,171	0.14	0.35	0	1
52 Weeks	934,171	0.46	0.50	0	1
Construction industry	934,171	0.07	0.25	0	1
Hotels and Restaurants	934,171	0.11	0.32	0	1
Public Sector	934,171	0.19	0.39	0	1
Agriculture	934,171	0.02	0.13	0	1
Public body	934,171	0.10	0.30	0	1
Sole Proprietorship	934,171	0.13	0.33	0	1
Any self-employment income	934,171	0.03	0.18	0	1

Units are person-years. The Employee and Employer Dominated Regions are the ‘tax-disadvantaged’ region: defined as reporting a pre-tax income above a notch threshold, triggering a liability such that lowering pre-tax income to below the notch threshold would increase post-tax income. The Irish and EU 2004 dummy variables indicate Irish citizenship, and citizenship of the ten countries that joined the EU in 2004. The base category is thus non-Irish, non-EU 2004 expansion citizenship. The 52 Weeks variable records working for the same employer for 52 weeks of the year, e.g. not switching employers. ‘Public body’ includes state-owned agencies, such as the electricity utility company and the public health service. ‘Public sector’ relates to government/public administration as an economic activity rather than corporate form.

Two variables generated from the dataset are the employee and employer dominated regions dummies. Consider the employee who faces an additional €8 weekly tax liability by earning one cent above a certain notch threshold. It should be immediately clear that the employee is strictly better off reporting earnings just below that threshold than by earning any income in the range of the threshold and €8 above. In the language of Kleven and Waseem (2013), this is a “dominated region” as the employee could increase both leisure and after-tax income by working less and reporting lower earnings. We thus define the dominated region dummy variable equal to one if the employee reports an income in that €8 interval.<sup>10</sup> Earning just below that threshold is a tax-advantageous income. We define a tax-advantaged income as earning within €3 of the threshold per week without crossing it.<sup>11</sup> Formally,

$$\text{Dominated Region} = \begin{cases} 1 & \text{if income in dominated region} \\ 0 & \text{if income} \in (\text{Threshold}-\text{€3}, \text{Threshold}] \end{cases}$$

The dominated region variable is generated analogously for the employer notch. Though analogous, we note a subtle difference between the clearly-dominated region from the employee’s perspective and the likely-suboptimal region from the employer’s perspective. Due to the increased tax liability, an employee whose earnings are just above the notch threshold is substantially more expensive than the same worker who simply works marginally fewer minutes. It seems unlikely that the marginal product of the worker is high enough to recoup the additional hundreds of euro in taxes in those few minutes. Though unlikely, we cannot say with certainty that this represents suboptimal behaviour on the part of the employer. Without knowledge of the firm’s costs, it is possible this is still profitable for the firm if the marginal product of labour is extremely high. For convenience we ignore this possibility and continue to refer to crossing this threshold as suboptimal behaviour.

---

<sup>10</sup>If crossing a notch threshold  $N$  increases one’s average tax rate from  $t_1$  to  $t_2$ , the dominated region is  $\left(N, \frac{1-t_1}{1-t_2}N\right]$ .

<sup>11</sup>The €3 cut-off is the largest possible amount that avoids overlap when the notches are €4 apart. Smaller bands (e.g. within €2 of the threshold) are feasible, but with reduced statistical power.

## 5 Empirical Analysis

The primary empirical question in this paper is if taxpayer behaviour depends on whether the tax statutorily falls on the employee or the employer. This question is teased out in three separate approaches below.

Principally, we investigate the extent of bunching just below the thresholds. This approach, pioneered by Saez (2010) and others, measures if there is an excess mass of earnings just below the notch thresholds. We will do this both in terms of the absolute number of people reporting earnings at the threshold and in terms of changes in the number. The latter is a method to alleviate concerns about a preference for round-numbers. As alluded to above, we do indeed find differential empirical responses between employee notches and employer notches. We see evidence of bunching in both levels and differences for employee notches, and essentially no evidence for employer bunching in either levels or differences.

Secondly, to explore channels that predict the different levels of responsiveness, we investigate whether the characteristics of ‘bunchers’ differ significantly from those in the dominated region. This approach comprises regressions predicting whether an individual reports earnings just above versus just below a notch threshold. We confirm that the determinants (e.g. nationality, age, sector, firm size) of earning just below the threshold are different from those just above.

Thirdly, we compare these determinants across notches. We will find that employee characteristics (e.g. age, nationality, any self-employment income) are predictors of employee-focused notches, but not employer-focused notches. The dataset spans from 2006–2013 but the introduction of the Universal Social Charge in 2010 reformed PRSI rules. To avoid comparing PRSI under these different regimes, we focus our attention on the pre-reform window.<sup>12</sup>

---

<sup>12</sup>The introduction of the Universal Social Charge (USC) at the height of the Great Recession was politically controversial. Thus, separate to changes in tax law, it may be difficult to assume everything else stayed constant post-USC. For example, Acheson and Lynch (2017) note potential salience effects of the USC extracting the Health Contribution tax away from PRSI. As an empirical matter, Hargaden (2020) finds little evidence of bunching during the recession.

## 5.1 Bunching estimates

Our first empirical analysis on statutory incidence investigates if the extent of bunching differs between employee and employer notches. The work on bunching near kink/notch thresholds is now very large, for example Ramnath (2013), Bastani and Selin (2014), Kleven and Waseem (2013), Saez (2010), Sallee and Slemrod (2012), Best and Kleven (2018), Mortenson and Whitten (2020).

Below we plot figures of the income distribution near the notch thresholds. In particular, these figures represent the weekly earnings in €2 bins for each year of our analysis. The dotted line to the left represents the threshold for crossing into Subclass AX, which causes a discrete jump in employee contribution. We thus call this the Employee notch. The dashed line to the right is at the threshold for Subclass AL, crossing which triggers an increased liability for the employer, and thus we call this the Employer notch.

Figure 3: Excess bunching graph in the first full year of data

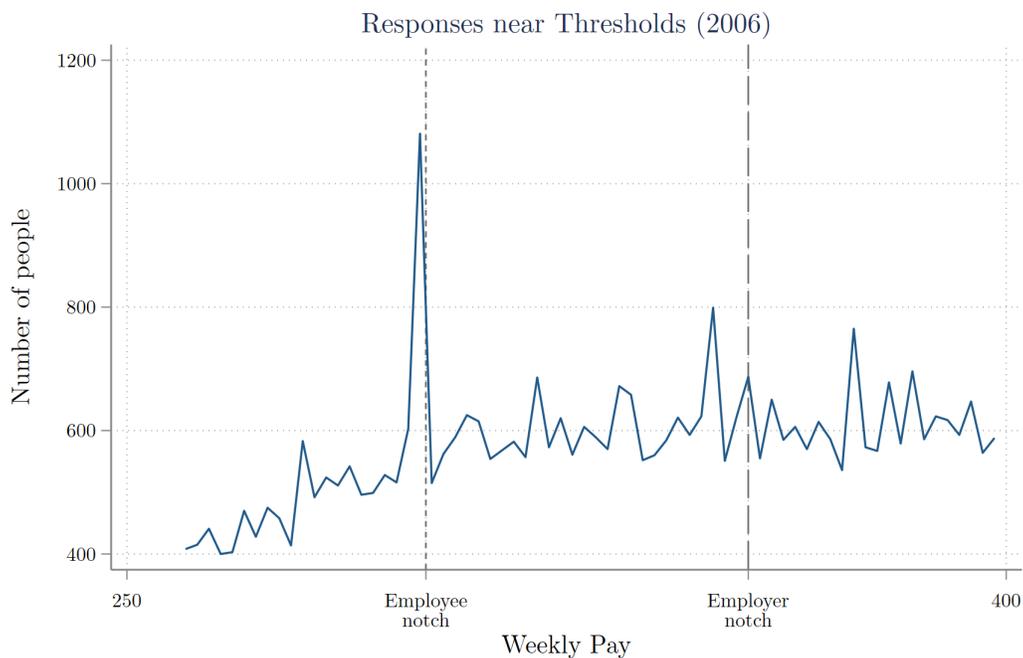


Figure 3 shows large spikes, or bunching, at several points in the income

distribution. There is clear evidence of bunching just below the employee notch. We see approximately 1100 people reporting an income that avoids the penalty associated with crossing that threshold, whereas the income distribution would suggest closer to 500 would be expected to earn within that €2 band, implying an excess mass of approximately  $1100 - 500 = 600$  people responding to the tax incentives. Similarly, we see a considerable (but smaller) spike to the left of the employer notch, with approximately  $800 - 600 = 200$  more people apparently reporting earnings just below the threshold than would be expected looking at comparable bins.

Figure 4: Bunching (Levels)

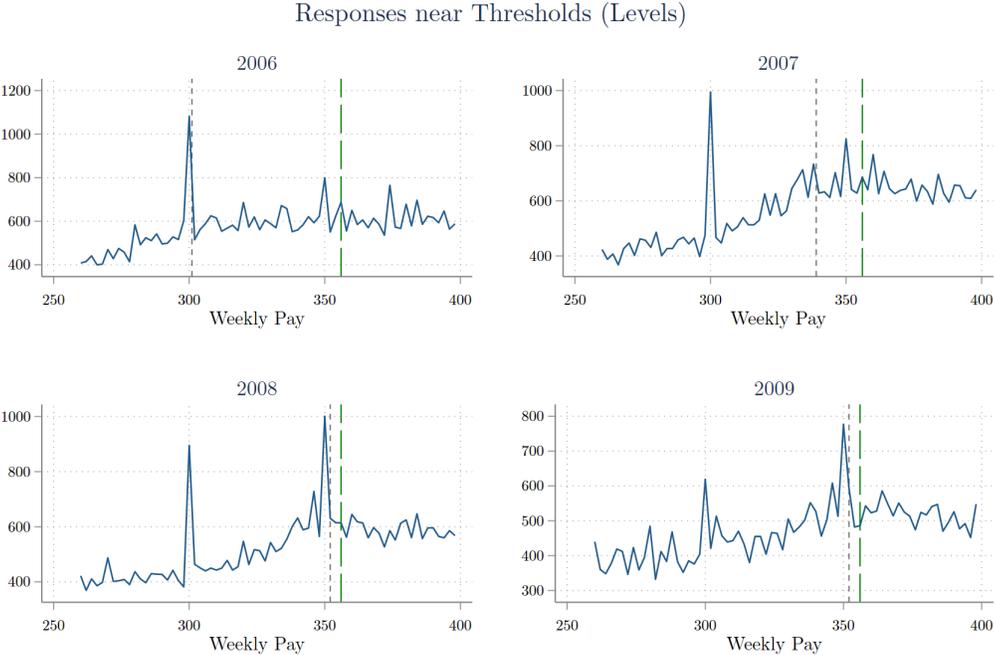


Figure 4 shows bunching graphs for all years between 2006 and 2009 inclusive. The results are weak for 2007, but quite striking for 2008. In that year, the AX and AL notches converged to within €4 per week of each other. In this case, both the employer and employee could lower that statutorily defined contribution with a relatively small adjustment in earnings. We see a large response to the notch thresholds in this year. This effect persists in 2009, as demonstrated in

the bottom right-hand panel.

However, it is impossible with this approach to compellingly disentangle tax-inspired bunching from round-number bunching when looking at figures like those above. For example, the large spike at €300 could be simply a preference for payments in multiples of one hundred. To overcome this confounding problem, we employ a different identification strategy. In contrast to Figure 4, the panels in Figure 5 use a difference-in-bunching approach.<sup>13</sup> This approach combines the benefits of the bunching estimator with a difference-in-differences framework. Rather than plot the level (or number) of people in a particular income bin, the difference-bunching estimator looks at the change in the number of people in that income bin. This is a cleaner form of identification for the tax effect, as opposed to a round number effect. The identification assumes that the taste for round numbers is constant through time. Conditional on the taste for round numbers not shifting between e.g. 2006 and 2007, we can attribute bunching near notch thresholds to tax incentives. Of course, just like regular differencing techniques, this approach comes at the cost of us losing the initial time-period’s observation.<sup>14</sup> Figure 5 demonstrate the bunching in differences rather than levels. Absent changes in the size of the labour force, the difference-bunching estimator should be mean zero.

In 2006 the spike at the employee threshold persists, suggesting the levels-bunching evidence is tax-inspired and not a round number effect. There is no spike at the employer threshold that year, which coheres with the lack of bunching in the levels-framework. In 2007 we see a spike of (approximately 180) people responding to the tax treatment at the employee threshold. The absence of bunching near the employer notch continues to hold in the differences framework in this year. We see spikes near the employee notch but not near the employer notch.<sup>15</sup> We conclude that there is evidence of earnings responses at

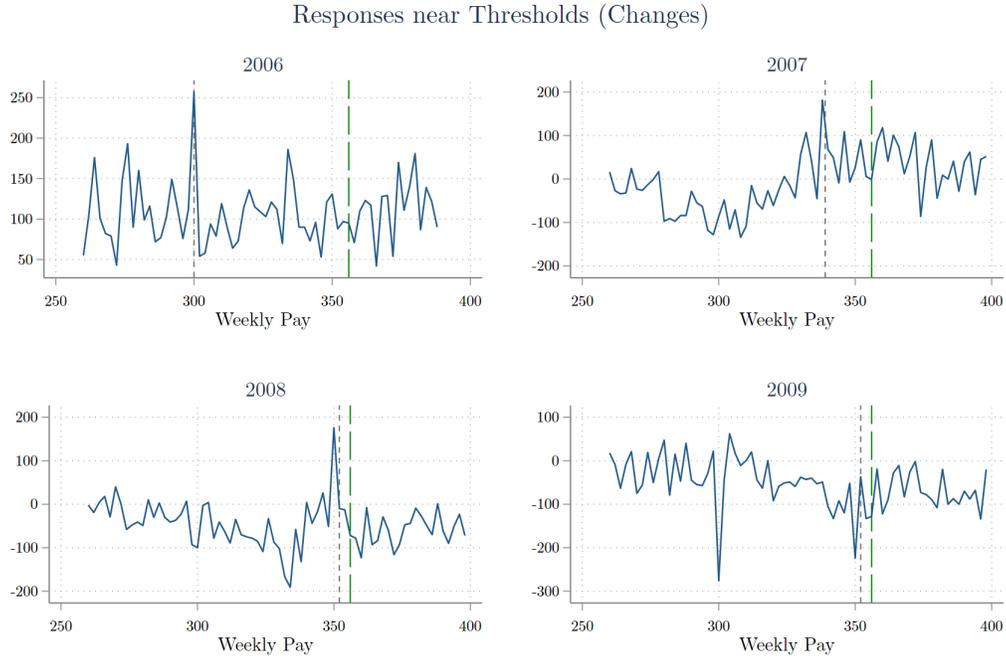
---

<sup>13</sup>An alternative but less flexible approach would be to scale down bunching where bins are multiples of ten. Identification from such an approach would require the taste for round numbers to be constant across different round numbers.

<sup>14</sup>For this paper, that year is 2005. Although not strictly included in the dataset, previous year’s earnings are available for the 80% of workers are continue employment with the same firm in both 2005 and 2006. The differences graphs for 2006 are calculated using this subsample.

<sup>15</sup>More formal analysis that generates a counterfactual density and estimates the statistical

Figure 5: Bunching-in-Differences



the employee threshold in both the levels- and differences-frameworks, but no evidence of responses at the employer threshold in either framework.

In 2008, when the notches are only €4 apart, we see yet another large spike in extra people (this time close to 200) responding to the tax incentives. This is quite strong evidence, and it occurs when a single small response could avoid both employer and employee tax notches. However the spike of approximately 200 people above expectations is only of the same magnitude as at the employee spikes in 2006 and 2007. That is, there is little evidence of additional bunching when the employer notch is also relevant. Specifically, this is consistent with a simple continuation of employee tax-oriented responsiveness.

We see little evidence of responsiveness in 2009 when the recession hit Ireland's labour market. This result mimics the findings of Hargaden (2020), which finds substantial cyclicity in responsiveness at other notches in the Irish income tax system.

---

significance of the bunching is included in Appendix C.

## 5.2 Determinants of Bunching

Given that we observe differences in the extent of bunching between employee and employer notches, it is pertinent to ask what is driving it. Who is doing the bunching? Is it possible to decompose the determinants of responsiveness into sensible predictors?

While we cannot interpret the determinants of bunching causally, correlational evidence can suggest mechanisms to better understand our headline results. This approach has been taken in the literature previously, e.g. Slemrod et al. (2001) and Advani (2017). Suppose we find workers in cash-based industries such as construction are more responsive to the notches than others. If we believe construction workers' real responses (e.g. their elasticity of labour supply) are similar to workers in other sectors, then our results would be more consistent with reporting behaviour than real responses.<sup>16</sup>

Inspired by the determinants of tax responsiveness found in Slemrod et al. (2001) and Advani (2017), we focus analysis on a list of plausible predictors of reporting an income just below the notch. We start with a relatively large list of covariates and will later use LASSO techniques developed by Belloni et al. (2014) to narrow the list down. We will see if the LASSO chooses similar variables for responding to the employee notch and the employer notch. These initial variables are listed in Table 3 and are broken down by whether the characteristics relate to the employee or employer. Summary statistics were presented in Table 2 above.

The variables are age, sex, Irish national, national of the EU 2004 accession states, a dummy for any self-employment income, a dummy for whether the individual worked for the same firm for fifty-two weeks of the year, construction, agriculture, hotels and restaurants, and public sector dummies, the legal form of incorporation of the firm (sole-proprietorship or other), and a dummy variable for whether the employer is a semi-state company. The majority of these indicators are self-explanatory, but some may require justification. Sex is

---

<sup>16</sup>As income is the product of wages  $w$  and hours  $L$ , it can be hard to disentangle real responses from reporting effects without direct data on either  $w$  or  $L$ . In Appendix B we decompose the response of people we do have this data on: workers who get paid the minimum wage. We find negligible evidence of real responses for minimum wage workers.

Table 3: Suite of potential determinants of tax avoidance

Individual Characteristics	Firm Characteristics
Age	Construction sector
Sex	Agriculture sector
Irish citizen	Hospitality sector
EU-2004 citizen	Public Sector
Any self-employment income	Sole-proprietorship
Same-firm 52 weeks a year	Public body

included as earlier studies have found women have more elastic labour supply responses than men (Bargain et al., 2014). The base category for citizenship is non-Irish/non-EU 2004 citizen, the majority of which are UK citizens. The prior expectation is that Irish citizens have better knowledge of the tax code than UK citizens, and that citizens of newly admitted EU countries (who are almost certainly recent migrants to Ireland) have less knowledge. The base category for industry is retail; it is expected that cash-based sectors like construction will have less tax compliance (and thus more avoidance) than retail, and that public sector bodies will have less responsiveness to tax incentives.

Following Hargaden (2020) we suspect the flexibility of labour market conditions (such as the availability of overtime hours) could vary over the business cycle, and thus we want to be sure the determinants are similar in the early periods (2006–2007) and the Great Financial Crisis recessionary period (2008–2009). Table 4 tests this, by estimating regressions for the early 2006–2007 period and for the complete 2006–2009 period. The table shows the coefficients from an OLS regression on whether an individual reports an income below ( $Y = 0$ ) or above ( $Y = 1$ ) the relevant threshold. The first two columns relate to the employee (AX) notch, and the third and fourth to the employer (AL) notch. Odd columns, that is columns 1 and 3, refer to years 2006–2007 inclusive; and the even columns refer to the full 2006–2009 period. As the regressions on the full sample have approximately twice as many observations as those for just 2006–2007, it is unsurprising that the statistical significance changes.

The most significant predictors for the AX notch in the early years are working for the same firm fifty-two weeks of the year (about 4.8% less likely to cross

Table 4: Determinants of reporting earnings above notch thresholds

	Employee (AX) notch		Employer (AL) notch	
	(1)	(2)	(3)	(4)
Age (Decade)	-0.007 (0.0052)	-0.003 (0.0035)	-0.014** (0.0049)	-0.015*** (0.0035)
Irish	0.005 (0.016)	0.017 (0.011)	0.017 (0.015)	0.011 (0.011)
Male	-0.011 (0.013)	-0.014 (0.0086)	-0.011 (0.011)	-0.008 (0.0082)
EU 2004	0.035 (0.019)	0.047*** (0.012)	0.027 (0.017)	0.021 (0.012)
Fifty-two weeks	-0.048*** (0.014)	-0.039*** (0.0089)	-0.032** (0.012)	-0.009 (0.0086)
Construction sector	-0.070** (0.022)	-0.088*** (0.015)	0.027 (0.020)	0.017 (0.015)
Hotels & Restaurants	0.000 (0.017)	-0.01 (0.011)	-0.011 (0.015)	-0.017 (0.011)
Public Sector	0.077*** (0.021)	0.069*** (0.014)	-0.012 (0.019)	-0.022 (0.013)
Agriculture	-0.025 (0.044)	-0.076** (0.029)	0.00073 (0.041)	-0.008 (0.030)
Public Body	0.022 (0.029)	-0.003 (0.020)	0.051* (0.026)	0.061** (0.019)
Sole-proprietorship	-0.13*** (0.015)	-0.14*** (0.011)	-0.027 (0.015)	-0.012 (0.011)
Self-employment income	-0.086* (0.044)	-0.11*** (0.027)	-0.15*** (0.038)	-0.11*** (0.028)
Constant	0.72*** (0.024)	0.70*** (0.018)	0.77*** (0.022)	0.80*** (0.017)
Year FEs	Yes	Yes	Yes	Yes
Observations	6,503	13,994	7,195	13,515
Adjusted $R^2$	0.037	0.036	0.007	0.005

Table reports coefficients from a linear probability model of crossing the notch threshold specified in the header. The unit of observation is person-year. The base category for the industry dummies is retail. Columns 1 and 3 refer to years 2006–2007 inclusive. Columns 2 and 4 are for 2006–2009 inclusive.

threshold), working in the construction sector (also less likely, by 7%), working in the public sector (7.7% more likely), working for a sole-proprietorship (13% less likely), and having self-employment income (8.6% less likely).

The marginal effects are not markedly different through time. This alleviates concerns about the Great Recession fundamentally changing responsiveness. For example, the statistical significance of the estimated effect of working in the construction sector has changed from the 5% level to the 1% level. However the coefficient shift (from -0.07 to -0.088) is relatively small, and is a statistically insignificant change, and so the broader narrative/conclusions from the first column in Table 4 continue.

In terms of the employer side, we find less statistically significant results in Columns 3 and 4. We find although many variables are insignificant, certain individual characteristics (e.g. age, any self-employment income) and firm characteristics (e.g. form of incorporation) are good predictors of tax responsiveness/avoidance behaviour.

Reassuringly, the signs on the coefficients across specifications are comparable, e.g. working fifty-two weeks of the year (usually a pre-condition for being salaried) lowers the probability of paying the early AX (employee) tax by 4.8% and the AL (employer) tax by 3.2%. However, there is less precision for the employer specification than for the employee. This is not surprising, as we have already noted that there appears to be greater responsiveness around the employee notch.

### 5.3 LASSO analysis of predictors

The evidence presented above suggest that the mechanisms for greater responsiveness are different for employer and employee taxes. However, it is plausible that the statistical insignificance could be driven by having too many explanatory variables in the regression. Just as there is an argument for including a full suite of variables when trying to understand the determinants of tax avoidance, there is an argument to be made for keeping regression equations parsimonious.

How sensitive are the results to including statistically redundant variables? In this section we use LASSO to find ‘robust determinants’ of responding, and

compare these determinants across employer and employee notches. This can give insight if the mechanisms for responsiveness differ across employee and employer notches, and serve as a robustness check of our specifications above,

We exploit the lassoShooting procedure in Stata to apply the Double-Lasso method (Belloni et al., 2014; Urminsky et al., 2016). The Double-Lasso method uses LASSO to select the ‘best’ set of covariates, then OLS to obtain point estimates. This approach isolates the variables that provide the most robust statistical significance given an atheoretic/flexible approach to prediction.

Table 5: Double-Lasso variable selection, Employee (AX) notch

Variable	Post-Lasso Marginal Effect
Male	-.0206
EU 2004	.0395
Construction sector	-.0632
Public sector	.0371
Sole-proprietorship	-.1920
Self-employment income	-.0924

Table 6: Double-Lasso variable selection, Employer (AL) notch

Variable	Post-Lasso Marginal Effect
Sole-proprietorship	-.0379

Tables 5 and 6 show the variables (and associated coefficients/marginal effects) from the Double-LASSO prediction algorithm for the Employee notch (Table 5) and Employer notch (Table 6). We see that the selected variables are notably different across notches.

Table 5 shows the covariates chosen by the Double-Lasso method as the most robust predictors of the Employee (AX) notch, and their associated marginal effects. We can see that both the signs and coefficients on the variables are consistent with this suggested by the Linear Probability Model, for example the construction sector indicating about a 6.3% decrease in the probability of crossing the notch threshold. We can also see that Double-Lasso, a flexible approach that is not driven by theoretical priors, focuses in on six variables to predict employee responses. These variables include both characteristics of the

individual (e.g. sex, nationality) and also characteristics of the firm (e.g. sector, form of incorporation).

Table 6 performs the identical procedure as Table 5 but on the Employer (AL) notch. It is immediately apparent that the variable selection varies enormously from that suggested in Table 5. Unlike the results there, which indicate a relatively large number of variables (both firm-based and employee-based) that predict responsiveness to the notches, Table 6 suggests that only a single variable — the form of incorporation (sole-proprietorship vs. other) — robustly predicts responsiveness. No characteristic of the individual, such as nationality or even their self-employment status, predicts reporting earnings below the notch threshold. The channels by which avoidance occurs differs between employee- and employer-focused incidence.

We also investigated the determinants of responsiveness of all notches in a pooled sample, and the inter-temporal stability of the estimates. The results cohere with the existing tables and so we provide them for the interested reader in Appendix A. A key takeaway is that while the determinants of responsiveness vary across notches, the determinants within notches are consistent through time. We take the inter-temporal stability as evidence that the differences across notches are systematic in our data.

## 6 Conclusion

This paper provides theoretical and empirical evidence that statutory incidence matters in ways not previously thought. The textbook case implicitly assumes continuity. We investigate incidence for tax notches, thresholds where taxes increase discontinuously. We show that the textbook incidence results do not hold for the case of notches. Notches are common: they have been studied in settings including pensions, social insurance, welfare programs, education, labor regulation, minimum wages, fuel economy policies, electricity prices, cellular service prices, and mortgage interest rates (Kleven, 2016).

For notches, taxes on net pay differ from dollar-equivalent taxes on gross costs. When notch thresholds are defined by the employee net pay, some workers'

net pay would fall below the notch threshold. Falling below the notch threshold means no tax is due. This implies there cannot be a seamless transition to the standard post-tax equilibrium, because those conditions do not constitute an equilibrium. Absent this default transition, both sides of the market can improve outcomes by adjusting earnings to avoid taxes.

Economists have arguably been too quick to assume the generality of classical results on statutory incidence. A growing literature has provided considerable empirical evidence questioning the applicability of these results. Chetty et al. (2009) introduced the importance of tax salience; Saez et al. (2019) found that Swedish firms used payroll tax cuts to expand employment but not increase wages; Lehmann et al. (2013) suggest sticky wages mitigate equilibrium responses; and Kopczuk et al. (2016) found differential evasion possibilities imply incidence is a function of the remitter. We show that incidence matters when tax schedules contain discontinuities.

Exploiting a natural experiment in Ireland where notch thresholds differ for employee and employer contributions to the social insurance tax, we find earnings respond to both employer-focused and employee-focused taxes, but not equivalently. There is a stronger response to taxes that are statutorily placed on the employee.

These responses may be driven by either real or reporting effects. Earnings are the product of hours and hourly wages, and adjusting hours is one channel of real response. To investigate if responses are real or reporting effects, we look at one set of workers with little flexibility on hourly wages: those on the minimum wage. We find no evidence that the hours of minimum wage workers systematically responded to the notch thresholds.<sup>17</sup> With all the caveats that looking at minimum-wage workers entails, we take this finding as suggestive evidence that the earnings responses documented in this paper are driven by reporting effects rather than real responses.

---

<sup>17</sup>Details are available in Appendix B.

## References

- Acheson, Jean and Donal Lynch**, “Implications of behavioural economics for tax policy,” *Irish Government Economic and Evaluation Service*, 2017, pp. 1–27.
- Adam, Stuart, David Phillips, and Barra Roantree**, “35 years of reforms: A panel analysis of the incidence of, and employee and employer responses to, social security contributions in the UK,” *Journal of Public Economics*, 2019, *171*, 29–50.
- , **James Browne, David Phillips, and Barra Roantree**, “Frictions and taxpayer responses: evidence from bunching at personal tax thresholds,” Technical Report, IFS Working Papers 2017.
- Advani, Arun**, “Who does and doesn’t pay taxes?,” *Institute for Fiscal Studies Briefing Note*, 2017.
- Bargain, Olivier, Kristian Orsini, and Andreas Peichl**, “Comparing labor supply elasticities in Europe and the United States: New results,” *Journal of Human Resources*, 2014, *49* (3), 723–838.
- Bastani, Spencer and Håkan Selin**, “Bunching and Non-Bunching at Kink Points of the Swedish Tax Schedule,” *Journal of Public Economics*, January 2014, *109*, 36–49.
- Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen**, “Inference on Treatment Effects after Selection among High-Dimensional Controls,” *The Review of Economic Studies*, 2014, *81* (2), 608–650.
- Best, Michael and Henrik Kleven**, “Housing Market Responses to Transaction Taxes: Evidence from Notches and Stimulus in the UK,” *The Review of Economic Studies*, 2018, *85*, 157–193.
- Blinder, Alan S and Harvey S Rosen**, “Notches,” *American Economic Review*, 1985, *75* (4), 736–747.

- Chetty, Raj, Adam Looney, and Kory Kroft**, “Salience and Taxation: Theory and Evidence,” *American Economic Review*, September 2009, 99 (4), 1145–77.
- , **John N. Friedman, Tore Olsen, and Luigi Pistaferri**, “Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records,” *The Quarterly Journal of Economics*, May 2011, 126 (2), 749–804.
- Cournot, Antoine A.**, *Recherches sur les Principes Mathématiques de la Théorie des Richesses*, Chez L. Hachette, Paris, 1838. Translated as *Researches on the Mathematical Principles of the Theory of Wealth* by Bacon, Nathaniel T. in 1897 and published by Macmillan, London.
- Fullerton, Don and Gilbert E Metcalf**, “Tax incidence,” *Handbook of Public Economics*, 2002, Volume 4, 1787–1872.
- Gruber, Jonathan**, *Public Finance and Public Policy*, Macmillan, 2005.
- Hargaden, Enda Patrick**, “Taxpayer Responses in Good Times and Bad,” *Journal of Economic Behavior and Organization*, 2020, *Forthcoming*.
- Kleven, Henrik Jacobsen**, “Bunching,” *Annual Review of Economics*, 2016, 8 (1), 435–464.
- and **Mazhar Waseem**, “Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan,” *The Quarterly Journal of Economics*, May 2013, 128 (2), 669–723.
- Kopczuk, Wojciech, Justin Marion, Erich Muehlegger, and Joel Slemrod**, “Does Tax-Collection Invariance Hold? Evasion and the Pass-Through of State Diesel Taxes,” *American Economic Journal: Economic Policy*, May 2016, 8 (2), 251–86.
- Lehmann, Etienne, François Marical, and Laurence Rioux**, “Labor income responds differently to income-tax and payroll-tax reforms,” *Journal of Public Economics*, 2013, 99, 66–84.

- Mortenson, Jacob A. and Andrew Whitten**, “Bunching to Maximize Tax Credits: Evidence from Kinks in the U.S. Tax Schedule,” *American Economic Journal: Economic Policy*, 2020, *Forthcoming*.
- Ramnath, Shanthi**, “Taxpayers’ Responses to Tax-Based Incentives for Retirement Savings: Evidence from the Saver’s Credit notch,” *Journal of Public Economics*, May 2013, *101*, 77–93.
- Saez, Emmanuel**, “Do Taxpayers Bunch at Kink Points?,” *American Economic Journal: Economic Policy*, August 2010, *2* (3), 180–212.
- , **Benjamin Schoefer, and David Seim**, “Payroll taxes, firm behavior, and rent sharing: Evidence from a young workers’ tax cut in Sweden,” *American Economic Review*, 2019, *109* (5), 1717–63.
- , **Manos Matsaganis, and Panos Tsakloglou**, “Earnings determination and taxes: Evidence from a cohort-based payroll tax reform in Greece,” *The Quarterly Journal of Economics*, 2012, *127* (1), 493–533.
- Salanié, Bernard**, *The Economics of Taxation*, MIT Press, 2011.
- Sallee, James M. and Joel Slemrod**, “Car Notches: Strategic Automaker Responses to Fuel Economy Policy,” *Journal of Public Economics*, December 2012, *96* (11), 981–999.
- Slemrod, Joel, Marsha Blumenthal, and Charles Christian**, “Taxpayer response to an increased probability of audit: evidence from a controlled experiment in Minnesota,” *Journal of Public Economics*, 2001, *79* (3), 455–483.
- Urminsky, Oleg, Christian Hansen, and Victor Chernozhukov**, “Using Double-Lasso Regression for Principled Variable Selection,” *Available at SSRN: <https://ssrn.com/abstract=2733374>*, 2016.
- Weber, Matthias and Arthur Schram**, “The Non-equivalence of Labour Market Taxes: A Real-effort Experiment,” *The Economic Journal*, 2017, *127* (604), 2187–2215.

## A Further tables on the determinants of responsiveness

The fact that predictor variables differ between employer- and employee-taxes provides initial evidence that mechanism of avoidance differ between employers and employees. However, what about the determinants of response in a pooled sample? Taking both AX and AL notches together, does the Double-Lasso method provide a sensible variable selection algorithm? Tables 7 and 8 investigate this. Firstly, Table 7 presents regression results on the pooled AX and AL sample, and also on a larger sample that includes other notches. As a major focus of this paper is the effect of statutory incidence, these additional notches (A1 and S subclasses) are not directly comparable to the AX and AL notches. However, if we are simply finding robust determinants of responsiveness, then including these notches improves the precision and power of our estimates.

Applying the Double-Lasso method on this broader set of notches that includes over 60,000 observations, we find a similar set of robust determinants of earnings response. In particular, the variables include both individual characteristics (age, and self-employment status) and firm characteristics (sector, and form of incorporation) as the strongest predictors of response. As this list of variables is chosen from the largest set of notches, we proceed with some further analysis taking the choice of these variables as the ‘best’ predictors of tax avoidance.

With this set of ‘best’ predictors, we now re-run the determinants of crossing the AX (employee) and AL (employer) thresholds with these predictors as covariates. Tables 9 and 10 are similar regressions to those presented earlier, but with two key differences. Firstly, the set of covariates is determined algorithmically by the Double-Lasso operator on the full set of available notches. Secondly, the tables include an explicit test of whether the coefficients in these regressions systematically differ from each other. This is achieved via Seemingly Unrelated Estimation. This procedure is comparable to a Hausman test comparing fixed effects models to random effects models. In the Hausman test, one checks if the coefficients in the different models are systematically different and thus if the

Table 7: Determinants of reporting above various notch thresholds

	(1)	(2)
	AX or AL	AX, AL, A1, or S
Age (Decade)	-0.0081*** (0.0020)	-0.0067*** (0.0017)
Irish	0.0087 (0.0063)	0.0082 (0.0052)
Male	-0.0073 (0.0049)	0.0000033 (0.0041)
EU 2004	0.036*** (0.0073)	0.026*** (0.0062)
Fifty-two weeks	-0.017*** (0.0050)	-0.0036 (0.0041)
Construction	-0.029** (0.0098)	-0.035*** (0.0079)
Hotels & Restaurants	-0.019** (0.0066)	-0.031*** (0.0059)
Public Sector	0.018* (0.0076)	0.021*** (0.0064)
Agriculture	-0.044** (0.017)	-0.048** (0.015)
Public Body	0.029** (0.011)	0.037*** (0.0087)
Sole-proprietorship	-0.13*** (0.0067)	-0.12*** (0.0058)
Self-employment income	-0.088*** (0.015)	-0.13*** (0.011)
Constant	0.78*** (0.011)	0.76*** (0.0096)
Year FEs	Yes	Yes
Observations	43,119	60,279
Adjusted $R^2$	0.037	0.028

Table 8: Double-Lasso variable selection, any notch

Variable	Post-Lasso Marginal Effect
Age (decade)	-.0088
Public Sector	.0270
Public Body	.0370
Sole-proprietorship	-.1190
Self-employment income	-.1434

RE model varies from the FE model. Here, we start in Table 9 by checking if the pre-recession AX coefficients are different from the 2008–2009 coefficients. The test is summarized by the  $\chi^2$  statistic displayed towards the bottom of the table, with its associated  $p$ -value. A high  $\chi^2$  (and thus low  $p$ -value) would reject the null of equivalent coefficients over the two time periods.

Tables 9 and 10 demonstrate that the robust, flexibly-selected determinants of crossing any notch are consistent, within notch, over time. Table 9 shows the determinants for the Employee AX notch in both the pre- and during-recession periods, and although the coefficients are not identical, there is not much evidence from the  $\chi^2$  that the determinants are systematically different, namely a  $p$ -value of 0.34 fails to reject a null that the determinants are statistically equivalent. Table 10 presents comparable information for the Employer AL notch. Again, the determinants are largely similar in both direction and magnitude, and a formal test of equivalent coefficients is not rejected ( $p = 0.36$ ). These null results are reassuring, as there does not seem to be systematic differences within notches through time. The structural relationship appears consistent regardless of the time period.

However, we can also test whether the coefficients from Tables 10 and 11 are different from each other. Just as we found that the variables chosen by the Double-Lasso method differed between notches, testing if the coefficients between Tables 10 and 11 are different is inherently a test of whether the determinants of responsiveness differ between notches. As above, this procedure will produce a test-statistic that follows a  $\chi^2$  distribution. The results overwhelmingly reject the null of equivalent coefficients. Comparing within-notch

Table 9: Determinants of crossing AX threshold on Lasso-selected variables, by time period

	(1) 2006 and 2007	(2) 2008 and 2009
Age (Decade)	-0.012* (0.0049)	-0.0023 (0.0046)
Public Sector	0.084*** (0.020)	0.076*** (0.018)
Public Body	0.016 (0.029)	-0.032 (0.028)
Sole-proprietorship	-0.14*** (0.015)	-0.18*** (0.015)
Self-employment income	-0.12*** (0.043)	-0.16*** (0.033)
Constant	0.72*** (0.017)	0.70*** (0.017)
Year FEs	Yes	Yes
Observations	6,503	7,491
Adjusted $R^2$	0.033	0.023
$\chi^2$ on null of equivalent determinants		5.68
$p$ -value		0.34

coefficients produced test-statistics around 5.5 and  $p$ -values around 0.35. Comparing between-notch coefficients produces a test-statistic of 218.4 and a  $p$ -value of less than 0.0000: the determinants are hugely different. Even when using the list of variables algorithmically chosen from a large set of notches, the channels that determine earnings responses are enormously different between the employee-notch and employer-notch.

Table 10: Determinants of crossing AL threshold on Lasso-selected variables, by time-period

	(1) 2006 and 2007	2 2008 and 2009
Age (Decade)	-0.018*** (0.0046)	-0.014** (0.0048)
Public Sector	-0.012 (0.019)	-0.028 (0.018)
Public Body	0.048 (0.026)	0.074** (0.028)
Sole-proprietorship	-0.024 (0.014)	0.0074 (0.016)
Self-employment income	-0.16*** (0.038)	-0.062 (0.040)
Constant	0.79*** (0.016)	0.80*** (0.018)
Year FEs	Yes	Yes
Observations	7,195	6,320
Adjusted $R^2$	0.006	0.003
$\chi^2$ on null of equivalent determinants		5.44
$p$ -value		0.36

## B Analysis of Minimum Wage Workers

This paper presented evidence that earnings responses are stronger for taxes on employees than employers.

Earnings responses can be broadly classed as ‘real’ (such as adjusting hours) or ‘reporting’ (such as avoidance or evasion). As earnings are the product of hours and hourly pay, it is not possible to separate hours adjustment without additional data on hourly pay.

We investigate hourly responses for workers with a set hourly rate: those who earn the minimum wage. If the earnings adjustments of this paper are caused by real responses, then the hours worked by minimum-wage workers will adjust to fall under the notch threshold. As the minimum wage is the prevailing wage for a non-trivial fraction of the labour force, this serves as a useful case to see how much of the results are driven by the real responses of minimum wage workers.

The Survey of Income and Living Conditions (SILC) is an EU-wide survey on income and inequality. It is administered by the Irish census bureau (Central Statistics Office) and collection is required by EU law. It has been collected since 2003, with a sample size of approximately 13,000 individuals annually. It includes hours worked. Table 11 shows the distribution of usual number of hours worked for people earning (close to) the National Minimum Wage.

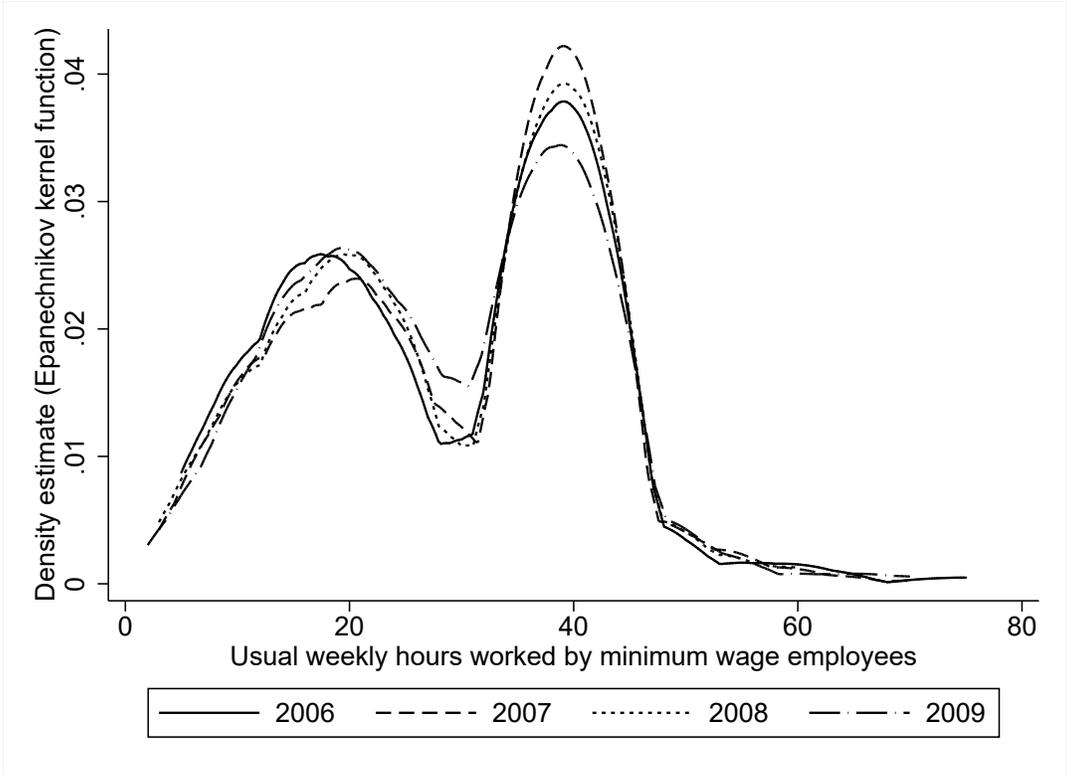
Table 11: Distribution of hours for people at/near the minimum-wage

Year	Mean	P10	P25	P50	P75	P90
2006	28.7	10	19	30	39	40
2007	29.3	12	20	32	39	40
2008	28.9	11	20	30	39	40
2009	28.6	12	18	30	39	40
Total	28.9	11	19	30	39	40

In 2006, the minimum wage of €7.65 meant a person working 40 hours a week was close to the notch threshold of €300. This raises questions about minimum wage-induced bunching, and so we investigated the hours-response of these workers to the changing notch threshold. SILC shows the majority of

workers (55%) work less than 36 hours per week. A substantial minority work 39 hours (18%) or 40 hours (14%). This almost certainly constitutes some of the bunching in 2006 at €300.

Figure 6: Kernel Density Estimates of Minimum Wage Workers



Graph shows the distribution of hours of minimum-wage workers over the four years in our sample. If bunching were caused by hours adjustment, we should see a 13% drop in hours of minimum-wage workers in 2007 (orange line) who had previously worked 39-40 hours a week.

If anything, we see an increase in the number of workers at 39-40 hours in 2007.

Figure 4 showed the bunching at €300 per week persisted in 2007. If the bunching was driven by minimum-wage workers, then their hours must adapt to changes in the hourly rate. As the minimum wage increased twice in 2007, (to €8.30 in January and €8.65 in July), the only way minimum-wage workers could be driving the bunching they would have needed to decrease their hours by 13% to avoid the notch threshold. We do not see any evidence of this reduction in the data. In fact, we do not see any bunching in 2008 or 2009 consistent

with 39 hours at the minimum wage (€337.35 per week). The distribution of hours for those reporting an hourly wage of within €0.10 of the minimum wage is effectively constant over the years. It does not appear that the bunching is driven by hours manipulation of minimum-wage workers.

## C Bunching Graphs with Polynomial of Best Fit

The figures in the paper do not include formal estimates of statistical significance. This section includes that, using the method from Chetty, Friedman, Olsen and Pistaferri (2011). We thank Tore Olsen for sharing his code. This approach provides a reduced form estimate for bunching below the notch.<sup>18</sup> Table 12 summarizes estimates of excess bunching and associated standard errors. The point estimates for bunching below the employer threshold is negative, indicating less mass than expected, though this is insignificant. As we did not see any evidence of earnings responses in the levels graphs, we see this as a continuation of no evidence of earnings responses at the employer notch. The point estimates are positive and significant for bunching near the employee threshold, except at the onset of the Great Recession in 2009.

Table 12: Difference bunching estimates and standard errors, by year

Year	Employee Bunching	Employer Bunching
2006	174.0 (40.9)	-37.1 (38.3)
2007	146.5 (73.1)	-61.0 (77.4)
2008	310.5 (62.4)	-35.2 (60.4)
2009	-142.7 (65.8)	-38.0 (72.0)

---

<sup>18</sup>The algorithm developed in Kleven and Waseem (2013) which calculates elasticities (as opposed to reduced form bunching estimates) did not converge for our dataset.

Figure 7: Employee Bunching (Differences) 2006

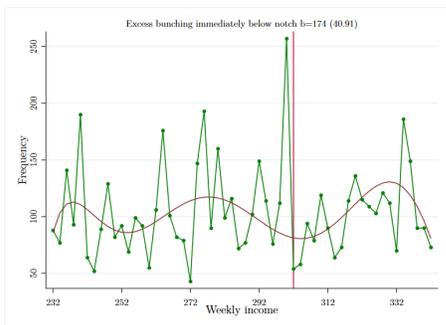


Figure 8: Employer Bunching (Differences) 2006

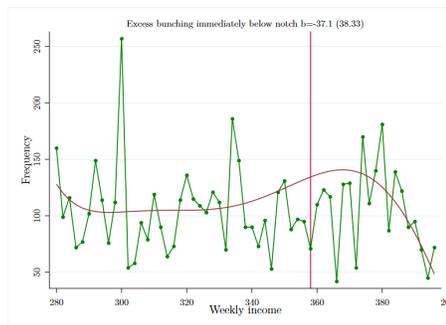


Figure 9: Employee Bunching (Differences) 2007

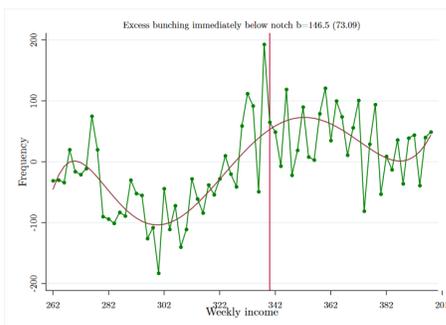


Figure 10: Employer Bunching (Differences) 2007

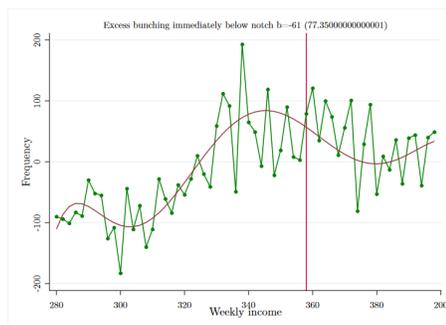


Figure 11: Employee Bunching (Differences) 2008

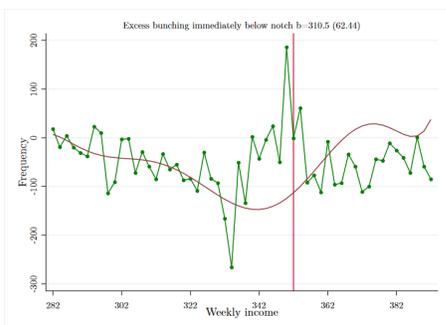


Figure 12: Employer Bunching (Differences) 2008

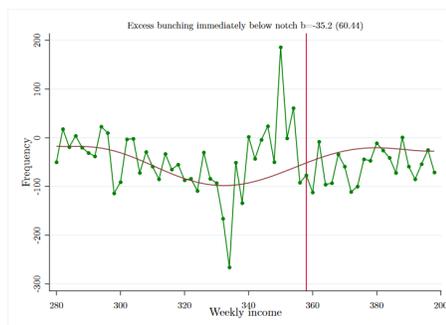


Figure 13: Employee Bunching (Differences) 2009

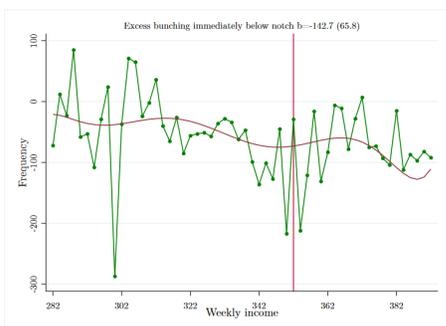


Figure 14: Employer Bunching (Differences) 2009

