

Does Statutory Incidence Matter? Earnings Responses to Social Security Contributions

Enda Patrick Hargaden*
Barra Roantree†

October 31, 2018

Abstract

This paper provides evidence that statutory incidence is a determinant of taxpayer behaviour. For the case of tax notches, we show that taxes on net pay differ from dollar-equivalent taxes on gross costs. Irish social security taxes contain separate notches on employers' and employees' contributions. Exploiting administrative records, we find short-run earnings responses are stronger when subtracted from employees' wages than when added to employers' gross costs. This challenges the statutory invariance result of standard economic models. By decomposing responses by employer and employee characteristics, we explore mechanisms to explain the differential levels of responsiveness.

JEL classification: H22, H26, H55, J38

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

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

†Economic and Social Research Institute, Whitaker Square, Sir John Rogerson's Quay, Dublin 2, Ireland. Email: barra.roantree@esri.ie

1 Introduction

In the textbook model, firms’ and consumers’ responses to taxes emerge from the relative sizes of the elasticities of supply and demand. These elasticities predict both the magnitude of the response and the distribution of burden. The economic incidence of a tax — who actually bears the burden of the tax — is independent of the statutory incidence, the party which the law declares responsible for the tax’s payment. As firms have the ability to pass taxes onto consumers through higher prices, it is irrelevant whether taxes are levied on the consumer price or the producer price. In this model, “statutory incidence is not a useful economic concept” (Salanié, 2011).

Recent work has suggested meaningful departures from the textbook model. Chetty et al. (2009) find tax-exclusive price labelling increases retailer profitability, and thus that the salience of taxes affects real outcomes. Kopczuk et al. (2016) provide convincing evidence that statutory incidence matters if some agents are better at evasion than others. In a lab experiment, Weber and Schram (2017) find differential extensive-margin labour supply responses for income taxes and dollar-equivalent payroll taxes.

This paper provides a similar point of departure for tax notches. Firstly we show that the standard theoretical results on statutory incidence do not hold when tax liabilities have discontinuous jumps. Secondly, we exploit quasi-experimental variation in employer and employee tax notches to test our theoretical predictions. The social security tax schedule in Ireland contains multiple notches which discontinuously increase statutory liabilities on *either* employees (via decreased pay) *or* employers (via increased gross costs). Consistent with our model, we find that statutory incidence itself affects responsiveness. We find clear evidence of earnings responses to this incentive, but only when the channel is decreased net pay. We find very little evidence of earnings responses when the tax increases gross costs.

This is striking because in both cases, the tax is remitted fully by the employer. There is no shift in evasion possibilities. The only difference is whether the tax represents an increase in gross costs or a decrease in net pay, a situation where one might suspect that textbook statutory invariance results would hold. We show they do not. If tax schedules are discontinuous, as they are with notches, then taxes on employees’ net pay do not result in the same equilibrium as dollar-equivalent taxes on employers’ gross costs.

The specifics of this policy facilitate a novel test of the invariance of outcomes to statutory incidence and, we believe, are unique to this setting. The Irish social security tax (Pay Related Social Insurance, or PRSI) splits contributions¹ into a portion that “shall be payable [by the] employer”, i.e. administered through an increase in gross costs, and a portion that will be “to

¹Section 13(2)(d) of the Social Welfare Consolidation Act, 2005

the exclusion of” the employee, i.e. remitted by the employer and administered via a decrease in net pay. These separate schedules face discontinuous notches at different thresholds. For example, crossing a €339 per week earnings threshold in 2007 decreased an employee’s net pay by €8.48 per week (€440 per year, 2.5% of income); and crossing a €356 threshold increased an employer’s gross costs by €8.01 per week (€416 per year, 2.25% of income). We find earnings are much more responsive to increases in ‘employee taxes’ than to comparable increases in ‘employer taxes’.

To investigate further, we empirically analyze the determinants of responsiveness. Using administrative tax return data, we decompose earnings responses to see if behaviour is driven by the characteristics of the employee, of the employer, or both. Self-employment income, or working in the construction sector, for example, are good predictors of reporting a tax-advantaged income. However, we find that these predictors systematically differ between employee- and employer-taxes. A suite of variables are found to be relevant for the employee tax, but only a single predictor — the form of incorporation — is a robust determinant of employer tax responsiveness. Thus earnings are not only less sensitive to employer-focused taxes, they appear to be determined by employee characteristics only when the statutory burden falls on that employee.

The most similar paper to ours is perhaps Lehmann et al. (2013) which finds differential responses to income taxes than to payroll taxes, suggesting sticky posted wages as a mechanism. We do not rely on market frictions as an explanation but rather show the circumstances when the textbook results do not hold for notches. In an important contribution, Saez et al. (2012) finds firms do not pass payroll taxes onto employees in the manner suggested by the standard theory. More recently, Weber and Schram (2017) finds non-equivalence between payroll and income taxes in a laboratory setting.

Similar papers investigating statutory incidence have generally relied on considerably more fundamental changes in tax law than this paper. For example Kopczuk et al. (2016) find that tax collections for state diesel taxes increased when the physical requirement of remitting the tax shifted from distributors to wholesalers, citing differential ability to evade taxes across the supply-chain. For the tax investigated in this paper, in all cases the responsibility for remittance remains entirely on the employer. There is no physical shift from distributors to wholesalers, or from employers to employees. As the point of the supply-chain does not change, and indeed the employer is continuing to remit using the same tax form (the “P35”), changes in evasion ability are not a relevant consideration. The only major difference between these notches is whether the incidence statutorily falls on the employer or the employee, i.e. whether in accounting terms it increases gross costs or decreases net pay.

The paper carries some considerable implications for both policy and the existing litera-

ture. From a policy perspective, it highlights the importance of administrative concerns for substantive economic outcomes. To a lesser degree, it provides support for real-world cases where statute formally states the desired legislative effect of taxes. In terms of existing academic work, the paper adds to the growing literature underscoring deviations from classical tax theory results. We believe this is the first time the conditions when statutory incidence applies to tax notches has been formalized. We provide empirical evidence supporting the result of non-equivalence, and highlight specific ways responses differ between notches on employees and employers.

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 shifted, in part, from the remitter to another person in the transaction. The standard incidence results comprise several inter-related predictions: that the identity of the remitter does not affect revenue; 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; and that the elasticities of supply and demand allocate relative burdens.

This section outlines the extent to which these results apply to ‘notches’, discontinuous increases in tax liabilities. Notches have been used primarily as a convenient tool for estimating important behavioural elasticities (e.g. Kleven and Waseem, 2013; Sallee and Slemrod, 2012) but with some notable exceptions (e.g. Blinder and Rosen, 1985) less has been done on their theoretical properties.

To match the empirical section that follows we will focus on a labour market example.² We will see that the textbook equilibrium predictions pertain when the tax notch is administered on a gross cost basis, but will not in general pertain when calculated on administered on a net pay basis. Intuitively, the discontinuous nature of notches can push net pay below the tax threshold. This implies the standard after-tax predictions do not constitute a well-defined equilibrium, as no tax is due. No such effects occur when the notches are administered on a gross cost basis, as the tax can only further increase the gross cost.

Denote W as the gross cost of wages to an employer, and w as the net wage to the

²The analysis applies equally to regular commodities.

employee. Absent taxation, $W = w$. Now consider a notch threshold, N , exceeding which triggers a lump-sum tax T , and thus for wages in excess of the threshold $W = w + T$. For now, specify that the tax-relevant price is the gross cost W .

$$\text{Tax Due} = \begin{cases} 0 & \text{if } W \leq N \\ T & \text{otherwise} \end{cases}$$

This setup closely resembles a standard payroll tax. It is instructive at this point to consider the standard incidence results for taxation, based on Taylor approximations of deviations from a prevailing pre-tax wage of W_0 . If the tax is owed, the fraction of the tax burden that is borne by employers is determined by the elasticities of supply (η_s) and demand (η_d). Assuming that the elasticities take standard³ values, the after-tax benefit w and gross cost to employer W adjust to split the incidence between the two sides. Specifically, under standard assumptions the after-tax gross cost of employment W will equal:

$$W = W_0 + \underbrace{\left(\frac{\eta_s}{\eta_s - \eta_d} \right)}_{>0} T$$

and thus the net wage (w) is:

$$\begin{aligned} w = W - T &= W_0 + \left(\frac{\eta_s}{\eta_s - \eta_d} \right) T - T \\ &= W_0 + \underbrace{\left(\frac{\eta_d}{\eta_s - \eta_d} \right)}_{<0} T \end{aligned}$$

These are the standard results per any textbook. In cases where $W \leq N$ then the lump-sum tax T is not owed, it is clear that the transaction price is $W_0 = w = W$. In a world where the tax due is based on gross costs (employer side) W , the equilibrium conditions are that prices adjust as above.

However, consider $W_0 \in (N, N + T)$. This is where the prevailing wage W_0 strictly exceeds the notch threshold, but not ‘by much’. We show that in this situation, whether the tax is defined by the gross cost (W) or net wage (w) matters. In the setting where the tax due formula is defined not by the employer cost W but by the employee benefit w , the formula above ($w = W_0 + \left(\frac{\eta_d}{\eta_s - \eta_d} \right) T$) generates circumstances where no tax is due. That is, holding remittance obligations constant, a form of statutory incidence affects real outcomes.

Specifically, one of two possible scenarios must prevail. Noting that $W_0 - N$ is how far

³That is strictly positive and finite supply elasticities, strictly negative and finite demand elasticities.

above the notch threshold the initial wage was, it must be the case that either:

$$\left| \left(\frac{\eta_d}{\eta_s - \eta_d} \right) T \right| \begin{cases} < W_0 - N & \Rightarrow w > N \Rightarrow \text{standard incidence formulae apply} \\ \geq W_0 - N & \Rightarrow w \leq N \Rightarrow \text{no tax is due} \end{cases}$$

We can see that when tax law is based on the lower wage, namely the one that accrues to the employee, there are circumstances when the laws of statutory invariance do not hold. 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 only slightly above N , the after-tax wage can fall below the notch threshold. Intuitively, if the tax wedge pushes the employee benefit below the notch threshold, then no tax is due. In this case, the standard formula do not constitute an equilibrium.. This is the key difference between discontinuous taxes (notches) and the usual ad valorem incidence results we find in textbooks. Notches' discontinuous nature are different.

In this second scenario the large, discontinuous nature of tax notches are self-defeating in a revenue sense. When statutory incidence is based on the employee side, tax notches can force such a divergence from original prices that the taxes themselves are not due. This breakdown in equilibrium cannot happen when statutory incidence is on the employer side, as taxes cannot decrease the total cost of labour.

Absent the standard equilibrium conditions pertaining, it is natural to ask what predictions can be made. In this second scenario, where tax liability is statutorily on the employee-side and where wages are only slightly above the notch threshold, mutually beneficial (Pareto improving) agreements/bargains exist for the employee and employer. The most obvious such bargain is where gross wages adjust to just avoid the tax, i.e. $w = W = N$. Relative to the case where a tax must be paid, this bargain would leave both employer and employer better off.⁴

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 equilibrium conditions will hold by default. 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, 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

⁴Of course, this is just one such agreement from a class of potential agreements.

theory predicts, but expect ‘new’ adjustment behaviour (such as the Pareto-improving bargain $w = W = N$) to occur in the net pay case.

3 Institutional Details

Social security in Ireland is funded primarily 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:

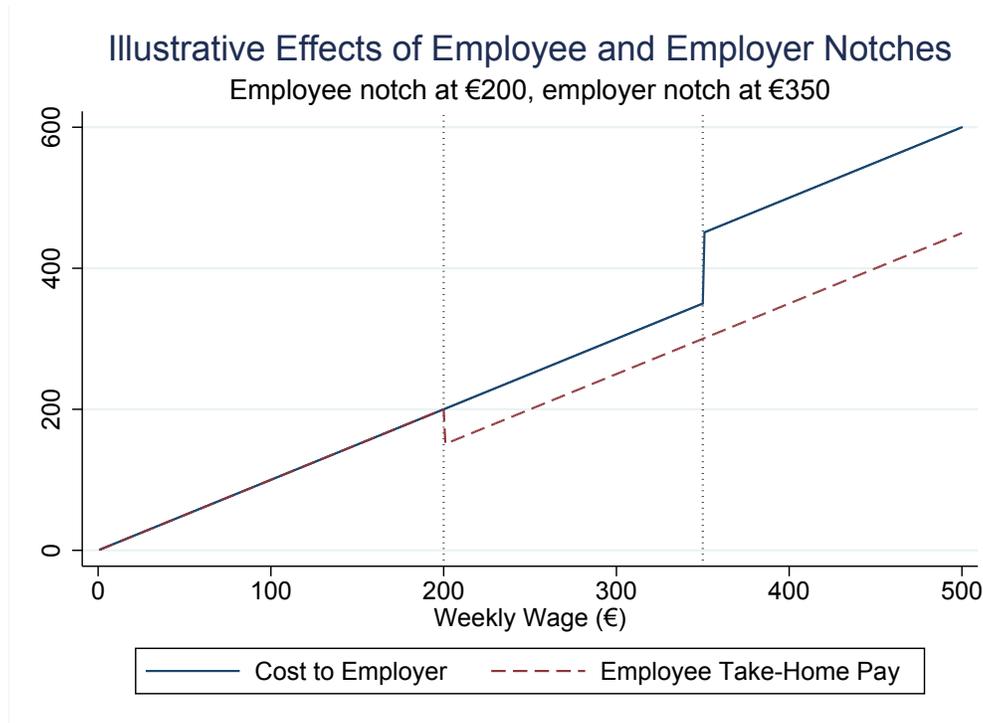
The employer shall, in relation to any employment contribution, be liable in the first instance to pay both the employer’s contribution comprised therein and also, on behalf of and to the exclusion of the employed contributor, the contribution comprised therein payable by the contributor.

This legal text asserts that remittance of both employer and employee contributions is the responsibility of the firm, but that the firm 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 firm’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 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), 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 the elasticity of earnings.

Figure 1: Graphical depiction of accounting effects of hypothetical €50 employee and €100 employer notches.



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 1 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 €350 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 will investigate if responsiveness differs between employee notches and employer notches. The null hypothesis motivated by standard theory is that statutory incidence does not affect behaviour. Of course the prediction of Section 2 is that there will be a greater

response to taxes based on reductions in net pay.

The full list of notches and liabilities associated with them is outlined in Table 1.

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

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.⁵

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. That is, unlike the previous example, the statutory incidence of this threshold fell entirely on the firm. In legal terms, this tax was to be paid through an increase in gross costs, not a reduction in net pay.

Standard Marshallian analysis would suggest that the statutory incidence of these notches should have zero impact on the behaviour of firms. In the Marshallian framework, 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. The theoretical analysis in Section 2 showed that this does not hold in the neighbourhood of notches, and additional ‘new’ adjustment behaviour is expected for the net pay case.

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).

⁵Beyond the notional increase in liability from a 1 cent pay increase, of course.

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 contains the details from the P35 tax form. This is comparable to a W-2 in the United States in that it is 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

Two variables generated from the dataset are the employee and employer “dominated regions” dummy. Consider an employee who faces an additional €8 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. Earning in the region is clearly in a tax disadvantaged portion of the income distribution. We thus define the dominated region dummy variable equal to one if the employee reports an income in that €8 interval. Earning below that threshold is a tax-advantageous income. We define, somewhat arbitrarily, reporting a tax-advantaged

income as earning within €3 of the threshold per week without crossing it. 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, it is important to 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.

5 Empirical Analysis

The primary empirical question in this paper is whether 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.

Firstly, we investigate the extent of bunching just below the thresholds. This approach, pioneered by Saez (2010) and others, measures if there exists 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 responses between employee notches and employer notches.

Secondly, to investigate potential channels to explain the different levels of avoidance, we investigate whether the characteristics of bunchers differ significantly from those who do not. This approach comprises regressions predicting whether an individual reports earnings just above versus just below a notch threshold. We find that the determinants (e.g. nationality, age, sector, firm size) of earning just below the threshold are different from those just above, as expected.

Thirdly, we will 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 large reforms introduced the Universal Social Charge in 2010 which affected PRSI structures. Furthermore notes Hargaden (2018) overall taxpayer responsiveness declined considerably during the recession. Consequently, we focus our attention on a tight window of 2006–2009 to minimize the effect of cyclical changes or policy reforms on behaviour.

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), Onji (2009), Mortenson and Whitten (2015).

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 solid red line 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 green line 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 2: Excess bunching graph in the first full year of data

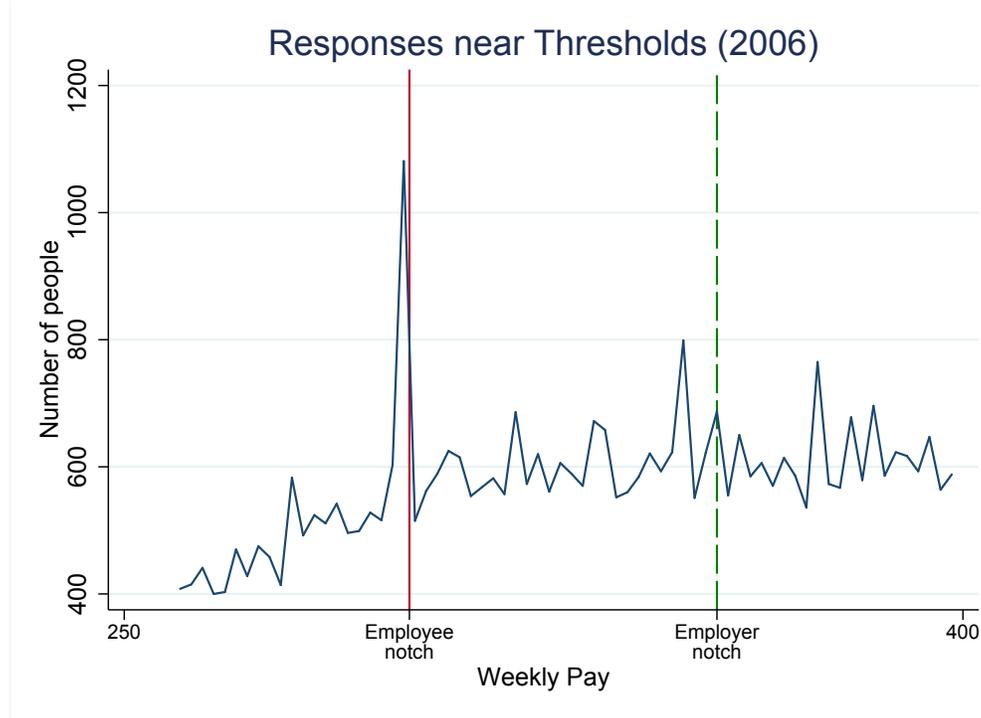


Figure 2 shows large spikes, or bunching, at several points in the income distribution. There is clear evidence of bunching just below the employee notch. We approximately 1100 people reporting an income that avoid the penalty associated with crossing that threshold, whereas the income distribution would 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 3: Bunching (Levels)

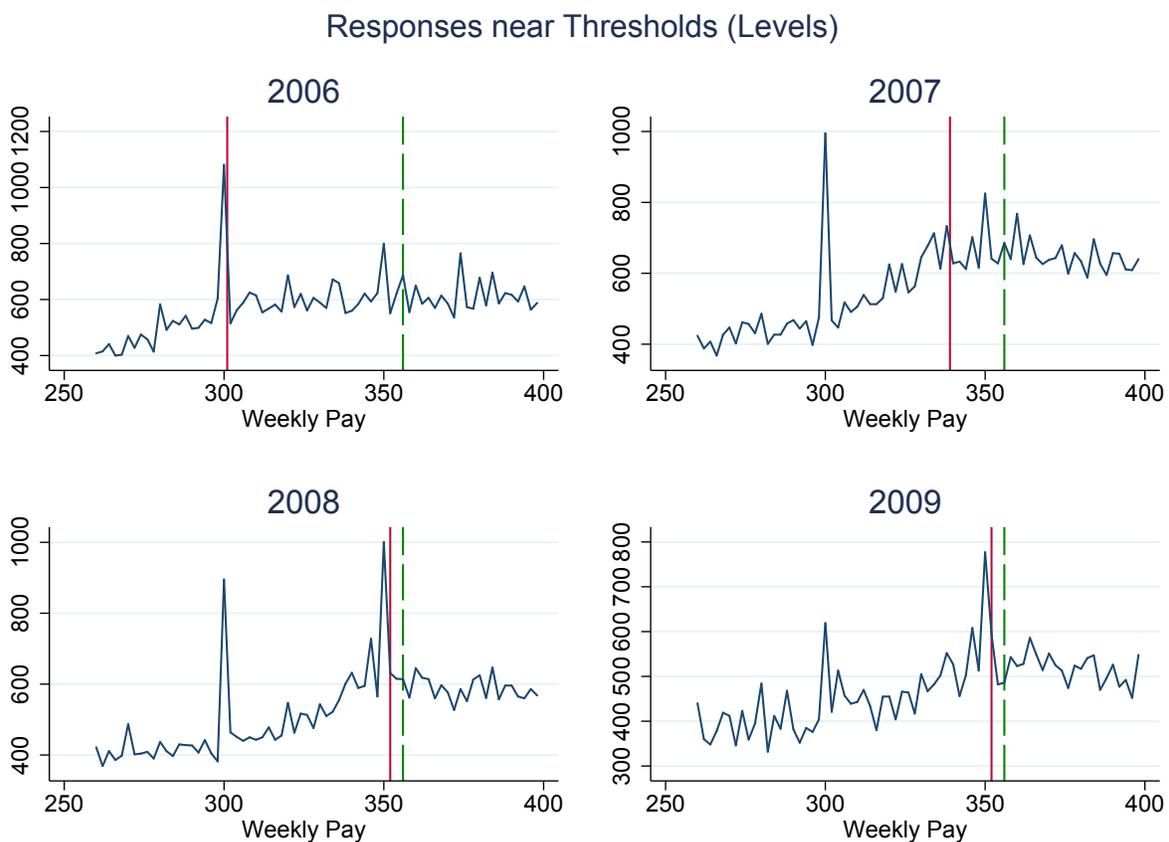


Figure 3 shows bunching graphs for all years between 2006 and 2009 inclusive. The results are weak for 2007, but quite striking for 2008. In that year, the AX and AL notches converged to within €4 per week of each other. In this case, both the employer and employee could lower that statutorily defined contribution with a relatively small adjustment in earnings. It is thus not surprising that 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 paying in multiples of one hundred. To overcome this confounding problem, we employ a different identification strategy. In contrast to Figure 3, the panels below use a difference-in-bunching estimator. 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. Of course, just like regular differencing techniques, this approach comes at the cost of us losing the initial time-period's observation.⁶ Figure 4 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. The implication of this approach, somewhat like difference-in-differences, is that any deviations from zero near tax thresholds are attributable to the tax avoidance and not round-number bunching.

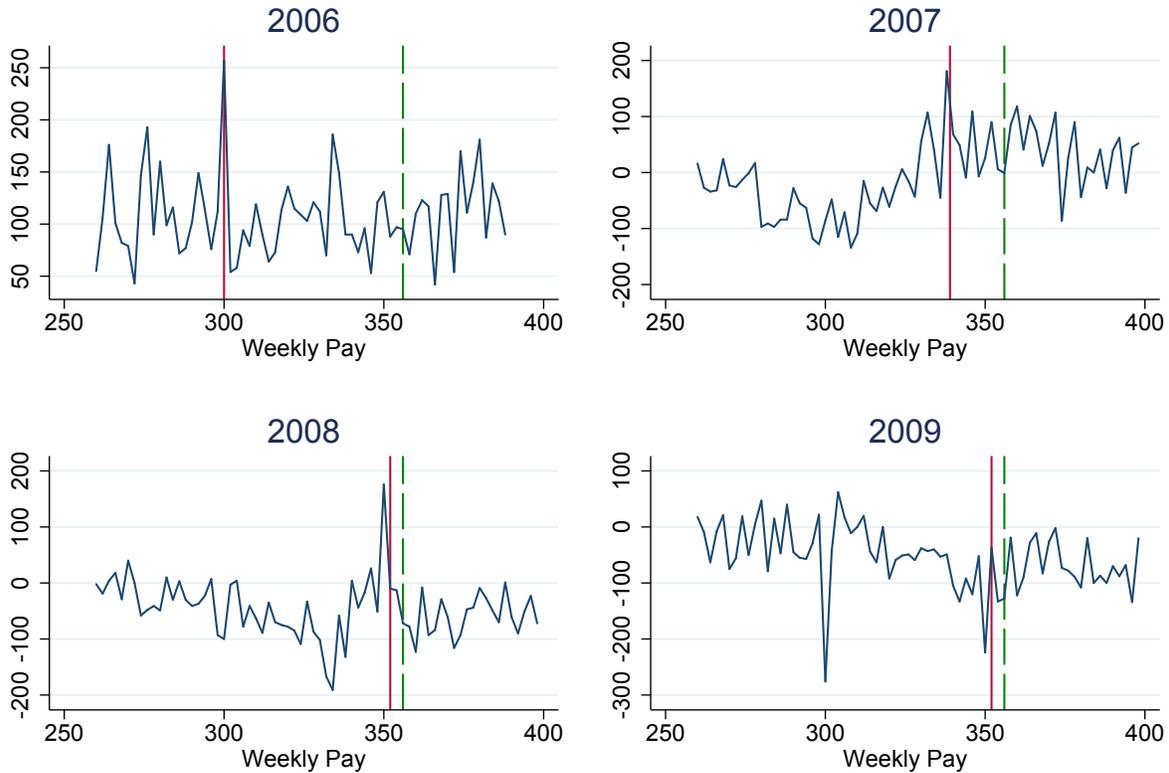
In 2007 we see a spike of (approximately 180) people responding to the tax treatment. This is a cleaner form of identification for the tax effect, as opposed to a round number effect. The identification assumes that the taste for round numbers is constant through time. Conditional on the taste for round numbers not shifting between 2006 and 2007, we can attribute these changes to tax incentives. Notice that unlike the bunching evidence in levels, we do not see much of a reaction for the employer threshold in 2007. We conclude that what was previously interpreted as a tax effect is better considered a round number effect.

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 (2018), which finds substantial cyclicity in responsiveness at other notches in the Irish income tax system.

⁶For 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.

Figure 4: Bunching-in-Differences
Responses near Thresholds (Changes)



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? For example, are workers in cash-based industries such as construction better at avoiding taxes than equivalent workers whose pay comes directly from government? Such analysis has been done before in other countries, e.g. Slemrod et al. (2001), Advani (2017). However these studies are based on ex post audits of tax returns, whereas this analysis is predicting avoidance behaviour separately for employee- and employer-targeted notches.

Inspired by the determinants found in Slemrod et al. (2001) and Advani (2017), we focus analysis on a list of plausible predictors of reporting a tax-advantaged income. We start with a relatively large list of covariates and will later use techniques developed by Belloni et al. (2014) to narrow the list down. 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.

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

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 included 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 (2018) 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). The analysis of Tables 4 and 5 test this, by estimating regressions for the early 2006–2007 period and for the complete 2006–2009 period. The first table of results 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 column of results relates to the employee notch, and the second to the employer notch. Unsurprisingly, given our finding of more bunching at the employee notch, the predictors of responding to this notch appear more significant than the employer notch.

The most significant predictors for the AX notch are working for the same firm fifty-two weeks of the year (about 4.8% less likely to cross threshold), working in the construction sector (also less likely, by 7%), working in the public sector (7.7% more likely), working for

Table 4: Determinants of reporting above notch thresholds, 2006–2007

	(1) Employee (AX) notch	(2) Employer (AL) notch
Age (Decade)	-0.0070 (0.0052)	-0.014** (0.0049)
Irish	0.0048 (0.016)	0.017 (0.015)
Male	-0.011 (0.013)	-0.011 (0.011)
EU 2004	0.035 (0.019)	0.027 (0.017)
Fifty-two weeks	-0.048*** (0.014)	-0.032** (0.012)
Construction sector	-0.070** (0.022)	0.027 (0.020)
Hotels & Restaurants	0.00019 (0.017)	-0.011 (0.015)
Public Sector	0.077*** (0.021)	-0.012 (0.019)
Agriculture	-0.025 (0.044)	0.00073 (0.041)
Public Body	0.022 (0.029)	0.051* (0.026)
Sole-proprietorship	-0.13*** (0.015)	-0.027 (0.015)
Self-employment income	-0.086* (0.044)	-0.15*** (0.038)
Constant	0.72*** (0.024)	0.77*** (0.022)
Year FEs	Yes	Yes
Observations	6,503	7,195
Adjusted R^2	0.037	0.007

a sole-proprietorship (13% less likely), and having self-employment income (8.6% less likely). 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 AX (employee) tax by 4.8% and the AL (employer) tax by 3.2%. However, there is less significance for the employer specification than for the employee. This is not surprising, as we have already noted that there appears to be greater manipulation around the employee notch.

Table 5: Determinants of reporting above notch thresholds, 2006–2009

	(1) Employee (AX) notch	(2) Employer (AL) notch
Age (Decade)	-0.0030 (0.0035)	-0.015*** (0.0035)
Irish	0.017 (0.011)	0.011 (0.011)
Male	-0.014 (0.0086)	-0.0082 (0.0082)
EU 2004	0.047*** (0.012)	0.021 (0.012)
Fifty-two weeks	-0.039*** (0.0089)	-0.0088 (0.0086)
Construction	-0.088*** (0.015)	0.017 (0.015)
Hotels & Restaurants	-0.0100 (0.011)	-0.017 (0.011)
Public Sector	0.069*** (0.014)	-0.022 (0.013)
Agriculture	-0.076** (0.029)	-0.0079 (0.030)
Public Body	-0.0027 (0.020)	0.061** (0.019)
Sole-proprietorship	-0.14*** (0.011)	-0.012 (0.011)
Self-employment income	-0.11*** (0.027)	-0.11*** (0.028)
Constant	0.70*** (0.018)	0.80*** (0.017)
Year FEs	Yes	Yes
Observations	13,994	13,515
Adjusted R^2	0.036	0.005

Table 5 presents a similar picture to Table 4, but this time the period of interest is expanded

to include the whole 2006–2009 period. As the regressions in Table 5 have approximately twice as many observations as in Table 4, it is not surprising that the likelihood of discovering a statistically significant result is increased. What is crucial is that the marginal effects are not markedly different through time. 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, but the coefficient change (from -0.07 to -0.088) is within one standard deviation and so the conclusions from the first column in Table 5 continue the narrative from Table 4. We find although many variables are insignificant, certain individual characteristics (e.g. nationality, number of weeks worked, any self-employment income) and firm characteristics (e.g. sector, and form of incorporation) are good predictors of tax avoidance behaviour. Further, we find less statistically significant results on the employer side.

5.3 Seemingly Unrelated Estimation

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. The search for ‘robust determinants’ of behaviour is particularly appropriate if one is comparing across different outcomes: how sensitive are the results to the variable selection method? Thus in addition to the determinants shown above, we also exploit the lassoShooting procedure in Stata to apply the Double-Lasso method (Belloni, Chernozhukov and Hansen, 2014; Urmitsky, Hansen and Chernozhukov, 2016). The Double-Lasso method, as suggested by the name, is a LASSO estimator that iterates through multiple possible specifications until it settles on what it considers the ‘best’ set of covariates. This approach isolates the variables that provide the most robust statistical significance given an atheoretic/flexible approach to prediction.

Table 6: 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 7: Double-Lasso variable selection, Employer (AL) notch

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

Table 6 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 7 performs the identical procedure as Table 6 but on the Employer (AL) notch. It is immediately apparent that the variable selection varies enormously from that suggested in Table 6. 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 7 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.

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 8 and 9 investigate this. Firstly, Table 8 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

Table 8: Determinants of reporting above various notch thresholds

	(1) AX or AL	(2) 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 9: 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

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 10 and 11 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 RE model varies from the FE model. Here, we start in Table 10 by checking if the pre-recession AX coefficients are different from the 2008–2009 coefficients. The test is summarized by the χ^2 statistic displayed towards the bottom of the table, with its associated p -value. A high χ^2 (and thus low p -value) would reject the null of equivalent coefficients over the two time periods.

Table 10: 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
χ^2 on null of equivalent determinants		5.68
p -value		0.34

Tables 10 and 11 demonstrate that the robust, flexibly-selected determinants of crossing any notch are consistent, within notch, over time. Table 10 shows the determinants for

Table 11: 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
χ^2 on null of equivalent determinants		5.44
p -value		0.36

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 χ^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 11 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 χ^2 distribution. The results overwhelmingly reject the null of equivalent coefficients. Comparing within-notch 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.

6 Conclusion

This paper provides theoretical and empirical evidence that statutory incidence matters in ways not previously thought. We investigate incidence for tax notches, discontinuities in schedules which are “ubiquitous across a wide range of tax and nontax settings” (Kleven, 2016). We show that the textbook incidence results do not hold for the case of notches.

For notches, taxes on net pay differ from dollar-equivalent taxes on gross costs. When notch thresholds are defined by the employee net pay, there can be no seamless transition to the standard 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. (2012) found firms were reluctant to treat young and old workers differently; 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. Decomposing the earnings response by characteristics of the employer and characteristics of the employee, we find a statistically significant difference in the drivers of responsiveness. Although both the employer- and employee-earnings response seem consistently stable, with little empirical support for inter-temporal effects, there is overwhelming evidence that the responses are different between employer taxes and employee taxes. As deadweight losses are functions of behavioural responses, policymakers wishing to minimize the distortions from taxation should place statutory incidence at the level that minimizes response.

References

- 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.
- 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.
- Hargaden, Enda Patrick**, “Cycles and Frictions in Taxpayer Behavior,” *Working Paper*, August 2018.
- Kleven, Henrik J. 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.
- Kleven, Henrik Jacobsen**, “Bunching,” *Annual Review of Economics*, 2016, *8* (1), 435–464.

- 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**, “How Sensitive Are Taxpayers to Marginal Tax Rates? Evidence from Income Bunching in the United States,” *Job Market Paper, Georgetown University*, 2015.
- Onji, Kazuki**, “The Response of Firms to Eligibility Thresholds: Evidence from the Japanese Value-Added Tax,” *Journal of Public Economics*, June 2009, 93 (5-6), 766–775.
- 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.
- , **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.