Taxpayer Responsiveness and Statutory Incidence: Evidence from Irish Social Security Notches

Enda Patrick Hargaden*
Barra Roantree†

November 2, 2017

Abstract

This paper provides evidence that the statutory incidence of a tax — the party which the law declares responsible for the tax's payment — is a determinant of taxpayer behaviour. A unique feature of the Irish social security tax is that statute defines whether tax increases "should" be paid by the employer or the employee. Using administrative data, we exploit changes in these provisions to analyze whether behaviour depends on statutory incidence. We find it does: earnings responses are much stronger when statutory incidence falls on the employee. By decomposing earnings responses by employer and employee characteristics, we investigate potential channels to explain the differential levels of avoidance. The results are inconsistent with the tax-collection invariance result from standard economic models.

JEL classification: H22, H26, H55
Keywords: statutory incidence, bunching, notches, social security contributions.

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

†Institute for Fiscal Studies, and University College London. Email: barra_r@ifs.org.uk

1 Introduction

A prominent feature of textbook microeconomics is that 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. Rather, in the standard model, firms' and consumers' responses to the tax emerge from costs and preferences, so that the relative magnitudes of the elasticities of supply and demand alone predict both the size of the response and the distribution of economic incidence.

This paper investigates this prediction in a labour market setting. The social security tax in Ireland (Pay Related Social Insurance, or PRSI) contains provisions which shift statutory incidence between employees and employers. In particular, the law specifies multiple thresholds (notches) which change the fraction of PRSI that "should" be borne by employees. For example, crossing a €339 per week earnings threshold in 2007 increased an employee's statutorily defined 'expected' PRSI contribution by €8.48, but had no effect on the employer's expected contribution. Similarly, in the same year, crossing a €356 earnings threshold increased an employer's expected contribution by €8.01, but left employees (statutorily) unaffected. In standard models, provisions on statutory incidence like this should not predict avoidance behaviour; holding everything else constant, we expect responses to employee taxes to be indistinguishable from employer taxes. In the words of Salanié (2011, p. 41) "statutory incidence is not a useful economic concept."

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. Similar papers investigating statutory incidence have relied on considerably more fundamental changes in tax law. For example Kopczuk et al. (2016) find tax avoidance fell for state diesel taxes 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 *legally* falls on the employer or the employee.

We find earnings are much more responsive to increases in 'employee taxes' than to comparable increases in 'employer taxes'. For example, there is much stronger evidence of reported earnings adjusting to avoid a €7 increase in employees' contributions than an €8 increase in employers' contribution. Given the predictions of standard models, this is a surprising result. To investigate further, we empirically analyze the determinants of avoidance.

Using administrative tax return data, we decompose earnings responses to see if avoidance 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 employee-incentivized avoidance, but only a single predictor — the form of incorporation — is a robust determinant of employer-incentivized avoidance. Thus earnings are not only less sensitive to employer-focused taxes, they appear only to be determined by employee characteristics when the statutory burden falls on that employee.

Finally we find that evidence of tax avoidance is by far strongest when both employee and employer contributions are lowered by adjusting earnings. This is consistent with wages being set in a bargaining framework between employer and employee. Such a framework is a significant departure from both the standard model of tax incidence and more recent work on the elasticity of taxable income (cf. Kleven and Waseem (2013), Chetty et al. (2011), Saez et al. (2012)), which assume earnings are determined not by employer-employee bargaining but by effort and avoidance adjustments given a marginal product of labour.

2 Institutional Details

Social security in Ireland is funded primarily funded through the Pay Related Social Insurance (PRSI) system. PRSI is a social security tax, with legal obligations on both employees and employers to contribute to the tax. In particular, legislation specifies the shares of total contributions ("statutory incidence") that are to be borne by the employer and the employee.[1] Contributions pay for a wide number of benefits, such as increased unemployment insurance, the state pension program, and dental check-ups. 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 a true insurance system.

The tax base for PRSI is quite large, including as it does all full-time or part-time employees over the age of 16 earning over €38 (approx. 40 US Dollars) per week. As per legislation[2] it is the employer who is responsible for remitting the entire amount of the tax to the Irish revenue agency, however the employer is permitted "to recover" any amount paid

[1] The exact shares vary on the level of income, with the employer share ranging from 75–100% of the total contribution.

[2] Section 13 of the Social Welfare Consolidation Act, 2005.

on behalf of the employee.

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 people 'bunch' near these thresholds. The amount of bunching, the excess mass of people 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). A large amount of bunching indicates a large responsiveness.

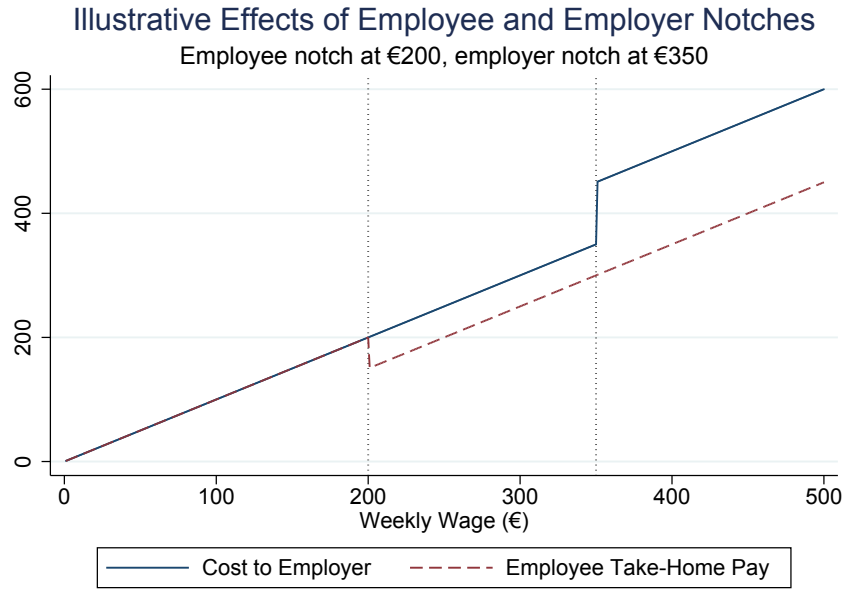

Figure 1: Graphical depiction of hypothetical €50 employee and €100 employer notches.

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 what fraction of the notch "should" be paid by the employer. 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.

The experimental design of this paper relies on the fact that the notched element of these thresholds apply differentially to the employers' share and the employees' share. 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 responsiveness. 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

	AX (Employee)		*AL (Employer)*	
Year	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 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 lump-sum penalty in the employee share. This is comparable to the notch at €200 in Figure 1. Notice that the tax increase applied to the employee share; the legal obligation/tax liability on the firm was entirely unaffected.[3]

In contrast, earning anything above €356 per week that 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 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. This situation is of course comparable to the notch at €350 in Figure 1.

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 unique equilibrium will be an adjustment of prices and quantities (hours or nominal wages) to avoid these penalties.

[3]Beyond the notional increase in liability from a 1 cent pay increase, of course.

3 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 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. 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 individual themselves. The income figure is formally called "Taxable Pay". The P35 form includes both the firm's and the individual's social security numbers, 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. Full details of the data are described in Hargaden (2016).

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 €8 tax liability by earning one cent above a certain notch threshold. It should be immediately clear that the employee is strictly better off earning an income 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 just below that threshold is a tax-advantageous income. How close is "just below"? It is not clear. We define, somewhat arbitrarily, "just below" as earning within €3 of the threshold per week without crossing it. Formally,

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

The dominated region variable is generated analogously for employers.

4 Empirical Analysis

The primary empirical question in this paper is whether taxpayer responsiveness 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. As alluded to above, we do indeed find differential responses. Secondly, to investigate potential channels to explain the different levels of avoidance, we investigate the determinants of response. This approach comprises regressions predicting whether an individual reports earnings just above versus just below a notch threshold. We find that individuals reporting incomes just below the threshold are not a random sample from the population, 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 as has been demonstrated in Hargaden (2016), taxpayer responsiveness declined considerably during the recession which resulted in the IMF assistance program in 2010. Consequently we focus our attention on a tight window to minimize the effect of any secular changes in bunching behaviour.

4.1 Bunching estimates

Our first empirical analysis on statutory incidence investigates if bunching differs between employee and employer notches. The work on bunching below thresholds that induce kinks or notches 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 (2015), 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 thick 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. Finally the dotted orange line is a third notch threshold, not mentioned before-hand, which applies to crossing the S1 threshold. As this is a different PRSI class and applies to a different set of employees, this is included largely as an avenue for future work.

Figure 2: Excess bunching graph in the first year of data

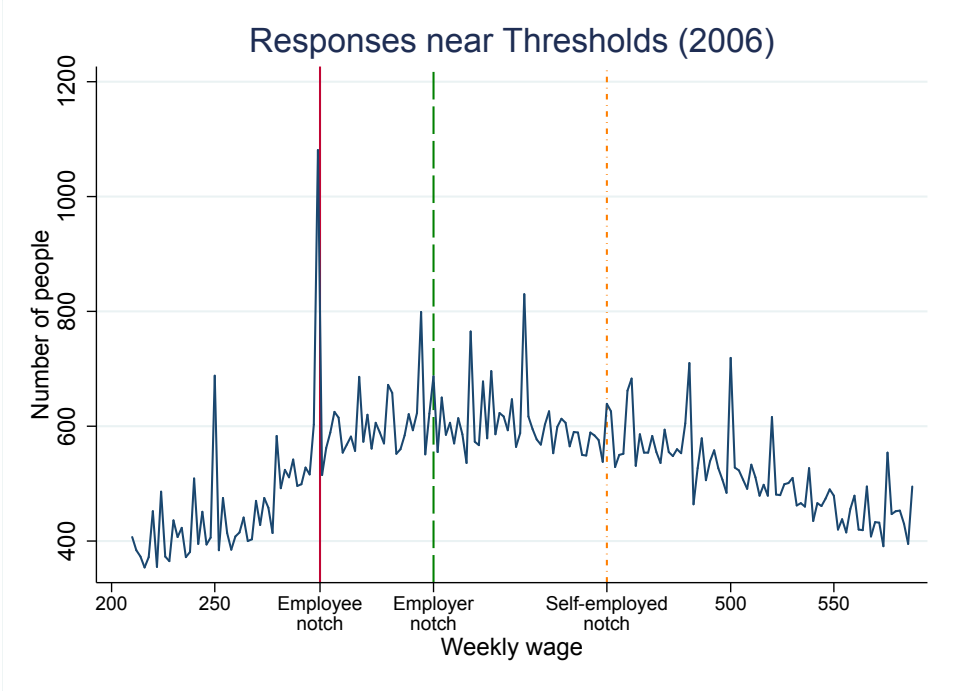

Figure 2 shows large spikes, or bunching, at several points in the income distribution. Some of these are simply at round numbers, such as at €250 per week. There is clear evidence of bunching just below the employee notch, indicating approximately 1000 people avoid the penalty associated with crossing that threshold, whereas the income distribution would suggest closer to 500 would be expected to earn within that €2 band. Similarly we see a considerable (but smaller) spike to the left of the employer notch, with approximately 200 ($800 - 600 = 200$) more people reporting earnings just below the threshold than would be expected looking at comparable bins.

Figure 3: Bunching (Levels) 2007

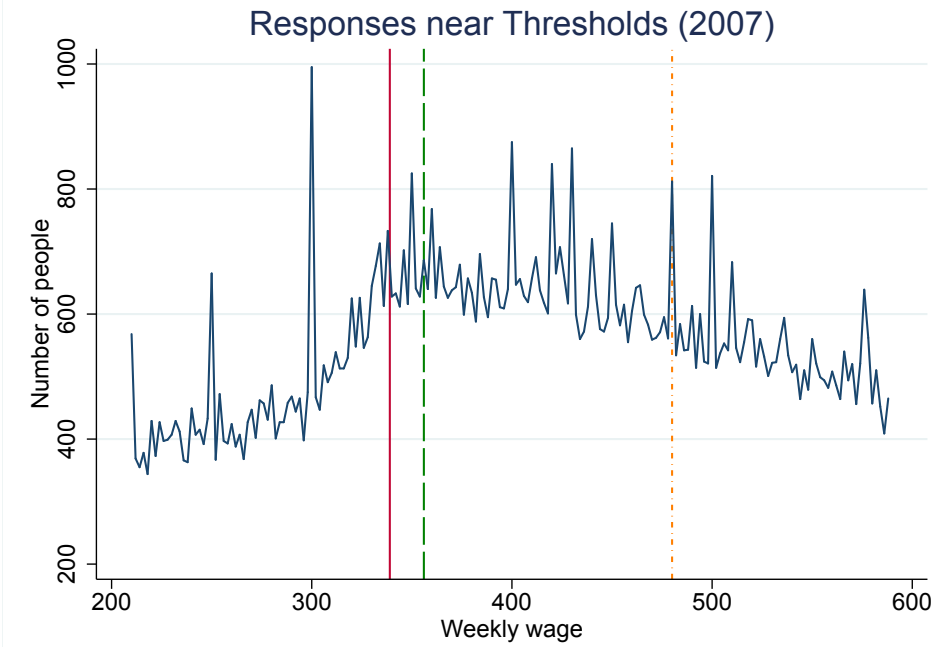

Figure 4: Bunching (Levels) 2008

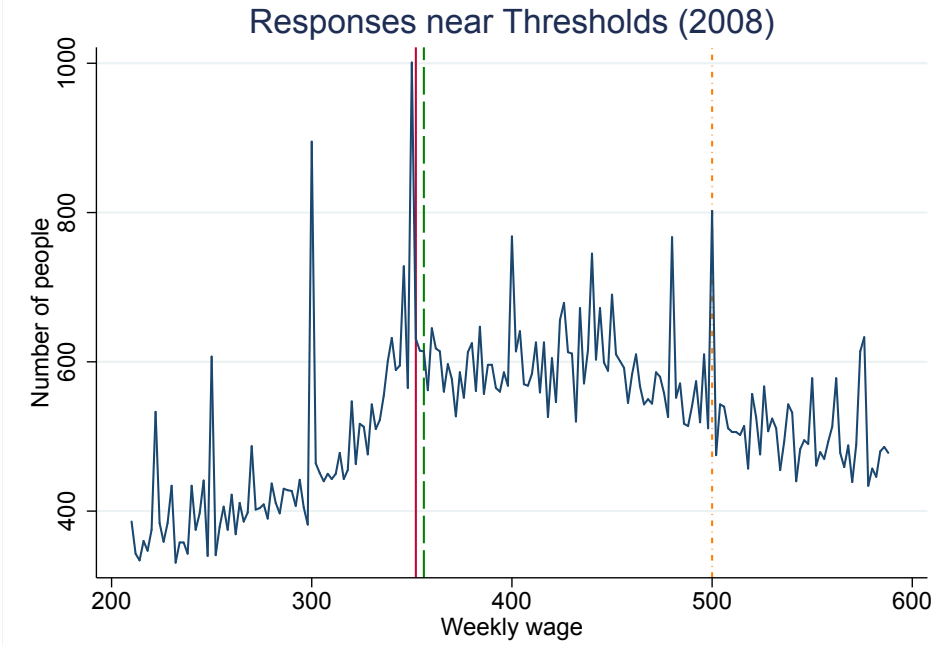

Figure 5: Bunching (Levels) 2009

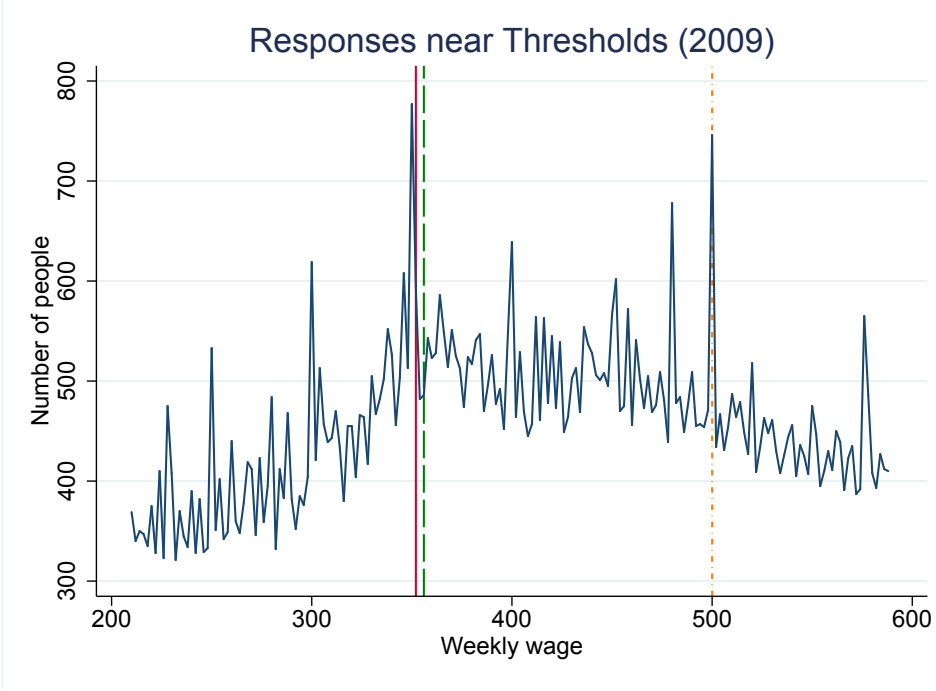

Figures 3 through 5 show bunching graphs for 2007–2009. The results are somewhat mixed for 2007, but quite striking for 2008. In that year, the AX and AL notches converged to within €4 of each other. In this case, both the employer and employee could lower that statutorily defined contribution with a relatively small adjustment in earnings. It is thus not surprising[4] that we see a large response to below the notch thresholds in this year. This effect persists in 2009, as demonstrated in Figure 5.

However, it is impossible to disentangle tax-inspired bunching from round-number bunching when looking at figures like those above. For example, the large spike at €300 in Figure 2 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 Figures 2 through 5 above, the figures below use a difference-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.) Figures 6 through 8 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,

[4]Perhaps unless one is thinking from a strictly neoclassical model.

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.

Figure 6: Bunching (Differences) 2007

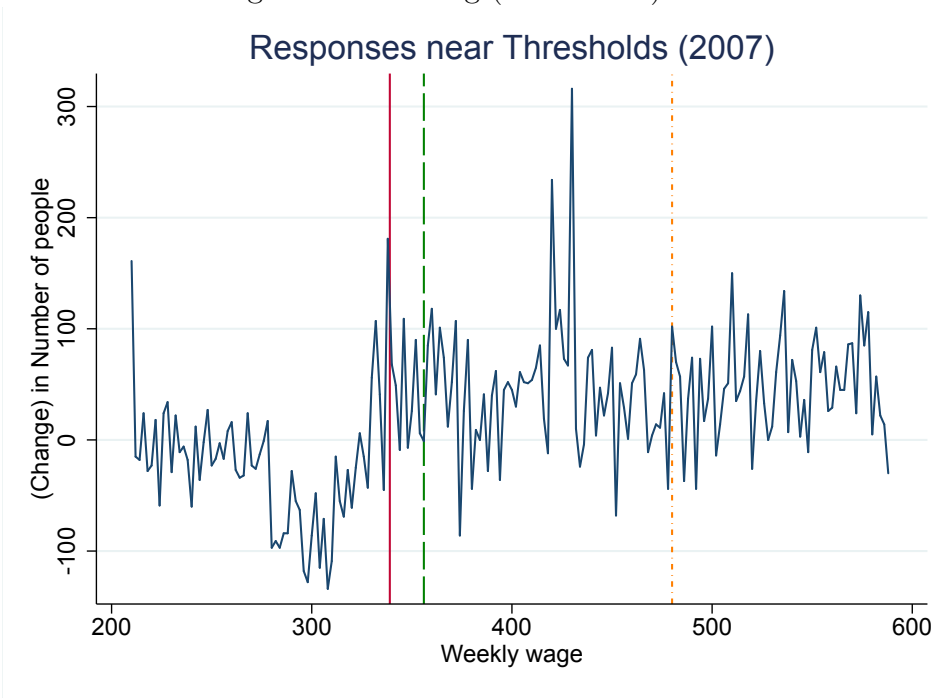

Figure 7: Bunching (Differences) 2008

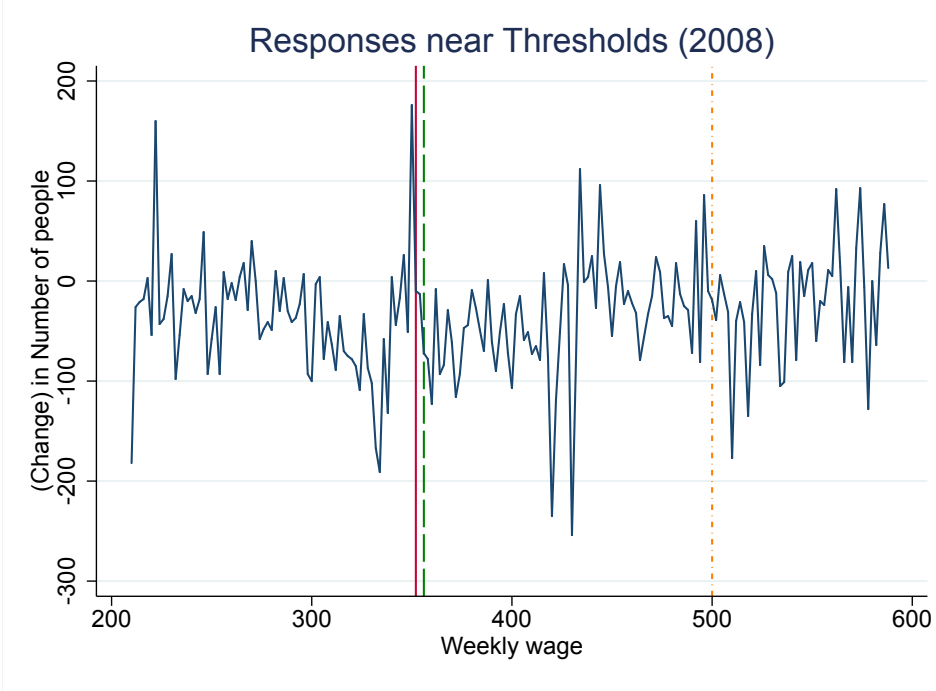

Figure 8: Bunching (Differences) 2009

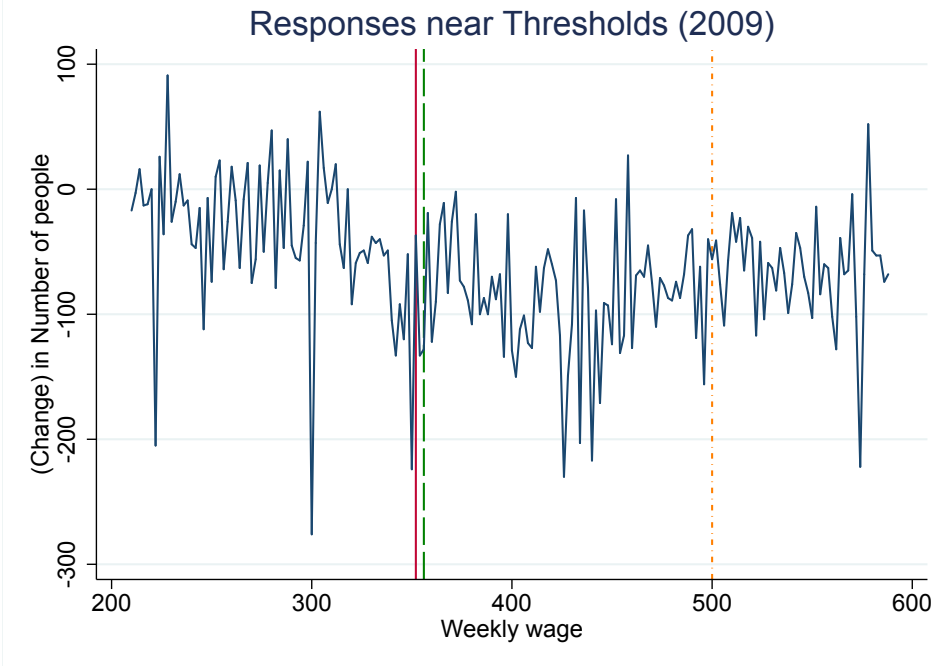

In 2007 we see a spike of (approximately 180) people responding to the tax treatment. This

is a cleaner form of identification for the tax effect, rather than a round number effect. The identification relies on the taste for round numbers being 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 the strongest evidence in the data, and it occurs when avoidance could arguably benefits both employer and employee. That is, when the adjustment of €4 could lower both the employee share and the employer share, we see a large spike in responses.

4.2 Determinants of Bunching

Given that we observe differences in the extent of bunching, it is pertinent to ask what is driving it. Who is doing the avoidance? Is it possible to decompose the determinants of bunching 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), Advani (2017), and Hargaden (2016), we focus analysis on a compact list of plausible predictors of reporting a tax-advantaged income.[5] These 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

[5]The conclusions from the analysis based on the complete set of possible covariates open to us follow what is reported below, but with a larger number of statistically insignificant results.

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 as the literature has found women have more elastic labour supply responses than men. 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 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 (2016) 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). Tables 4 and 5 tests this, running regressions for the early 2006–2007 period and for the complete 2006–2009 period.[6] The first column of results represents the coefficients from an OLS model 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 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 tax by 4.8% and the AL 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 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

[6]Permissions to distribute specific results from the 2008–2009 regressions are pending.

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

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

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

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

twice as many observations as in Table 4, it is not surprising that the likelihood of discovering a statistically significant result in increased. What is crucial is that the marginal effects — the coefficients — are not markedly different through time. For example, the statistical significance on 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. The conclusions from the first column in Table 5 continues 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.

4.3 Seemingly Unrelated Estimation

The evidence presented above suggest the mechanisms for avoidance 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; Urminsky, 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 hold 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 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

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

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

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 a willingness to report 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 convincing 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 tax avoidance, 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 tax avoidance. 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 tax avoidance. 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.

Table 8: Determinants of reporting above various notch thresholds

	(1) AX or AL	(2) AX, AL, A1, or S
Age (Decade)	-0.0081***	-0.0067***
	(0.0020)	(0.0017)
Irish	0.0087	0.0082
	(0.0063)	(0.0052)
Male	-0.0073	0.0000033
	(0.0049)	(0.0041)
EU 2004	0.036***	0.026***
	(0.0073)	(0.0062)
Fifty-two weeks	-0.017***	-0.0036
	(0.0050)	(0.0041)
Construction	-0.029**	-0.035***
	(0.0098)	(0.0079)
Hotels & Restaurants	-0.019**	-0.031***
	(0.0066)	(0.0059)
Public Sector	0.018*	0.021***
	(0.0076)	(0.0064)
Agriculture	-0.044**	-0.048**
	(0.017)	(0.015)
Public Body	0.029**	0.037***
	(0.011)	(0.0087)
Sole-proprietorship	-0.13***	-0.12***
	(0.0067)	(0.0058)
Self-employment income	-0.088***	-0.13***
	(0.015)	(0.011)
Constant	0.78***	0.76***
	(0.011)	(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 this set of 'best' predictors of tax avoidance, 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. These is achieved using 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*	-0.0023
	(0.0049)	(0.0046)
Public Sector	0.084***	0.076***
	(0.020)	(0.018)
Public Body	0.016	-0.032
	(0.029)	(0.028)
Sole-proprietorship	-0.14***	-0.18***
	(0.015)	(0.015)
Self-employment income	-0.12***	-0.16***
	(0.043)	(0.033)
Constant	0.72***	0.70***
	(0.017)	(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 demonstrates the determinants for the Employee AX notch in both the pre- and during-recession periods, and although the

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***	-0.014**
	(0.0046)	(0.0048)
Public Sector	-0.012	-0.028
	(0.019)	(0.018)
Public Body	0.048	0.074**
	(0.026)	(0.028)
Sole-proprietorship	-0.024	0.0074
	(0.014)	(0.016)
Self-employment income	-0.16***	-0.062
	(0.038)	(0.040)
Constant	0.79***	0.80***
	(0.016)	(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	

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 appear to be systematic differences within notches, through time. The structural relationship seems solid 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. Again this procedure will produce a test-statistic that follows a χ^2 distribution. This test produces a test-statistic that overwhelmingly rejects 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. Even when using the list of variables chosen from a large set of notches, the channels that determine tax avoidance are enormously different between the employee-notch and employer-notch.

5 Conclusion

Economists have known at least since Mill (1884) that taxes are shifted, in part, from the remitter to another person in the transaction. In the context of the difference between economic (or effective) and statutory incidence, Musgrave (1959, p. 231) mused that "Perhaps a more useful concept of shifting may be secured by measuring the difference between the actual change in distribution (or effective incidence) and in the incidence of legislative intent." This shifting of tax incidence, and thus the difference between who the legislative branch hopes pays a tax and who bears economic incidence, has risen in prominence in recent debates about who actually pays the corporate tax, for example (Fuest, Peichl and Siegloch, forthcoming).

This paper provides evidence that statutory incidence matters. Building on previous literature (Kopczuk et al., 2016) that finds avoidance is a function the remitter, we find avoidance is a function even of legislative intent. 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, who perhaps finds that cost more salient than costs on the employer.

Decomposing the earnings response by characteristics of the employer and characteristics of the employee, we find a statistically significant difference in the drivers of avoidance. Although both the employer- and employee-response seem consistently stable, with tests of inter-temporal changes finding little empirical support, there is overwhelming support that the responses are different between employers and employees.

Finally, we find the strongest evidence for tax avoidance when adjustments could lower both employee- and employer-focused taxes. This is consistent with a bargaining model where earnings adjustment is more likely if both parties lower their statutory burdens. In this bargaining context, statutory thresholds may plausibly serve as a basis point for negotiations, rendering them economically meaningful.

References

Advani, Arun, "Who does and doesn't pay taxes?," *Institute for Fiscal Studies Briefing Note*, 2017.

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," *Mimeo*, February 2015.

Chetty, Raj, 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.

Fuest, Clemens, Andreas Peichl, and Sebastian Siegloch, "Do Higher Corporate Taxes Reduce Wages? Micro Evidence from Germany," *American Economic Review*, forthcoming.

Hargaden, Enda Patrick, "Taxpayer Responses over the Cycle: Evidence from Irish Notches," *Job Market Paper, University of Michigan*, 2016.

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.

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.

Mill, John Stuart, *Principles of Political Economy*, D. Appleton, 1884.

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.

Musgrave, Richard A., "oThe Theory of Public Finance: A Study in Public Economy," 1959.

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.

_ , Joel Slemrod, and Seth H. Giertz, "The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review," *Journal of Economic Literature*, March 2012, *50* (1), 3–50.

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.